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

Program Year 6: Quarterly Report 3
(No-Cost Extension Year Ending 9/30/2021)

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

6/25/2021

This is the Program Year 6: Quarterly Report 3 (PY6 Q3) of the Aristotle Cloud Federation team. This report is part of a one-year No-Cost Extension. We report on plans and activities for each area of the project Work Breakdown Structure (WBS).

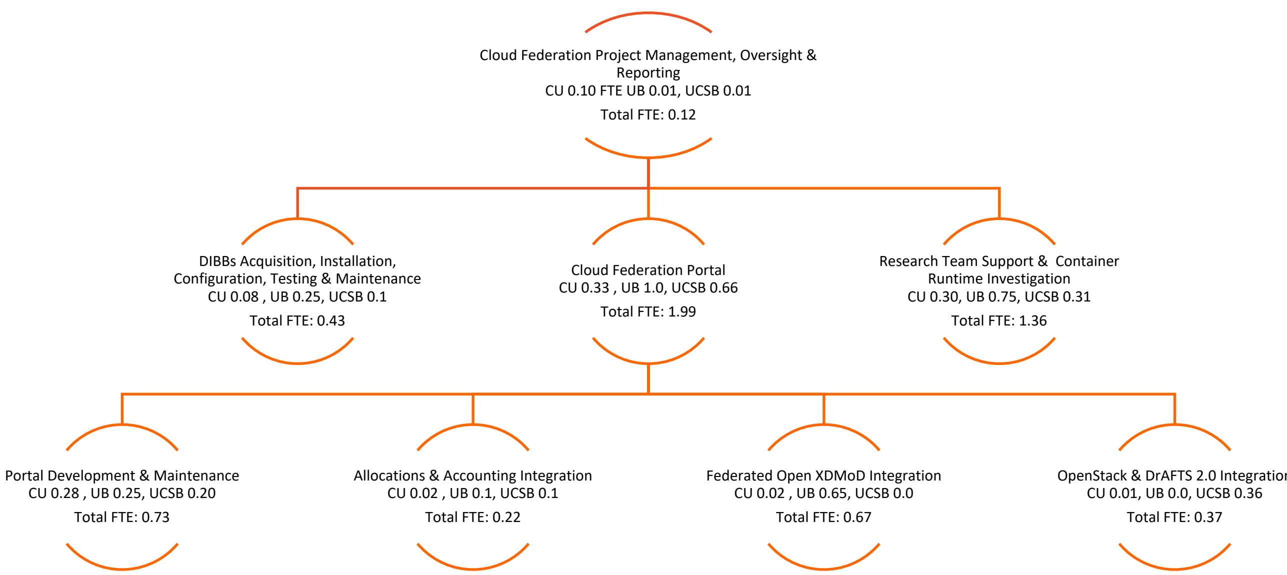

#1541215

Contents

1.0 Cloud Federation Project Management, Oversight & Reporting

1.1 Subcontracts

All subcontracts are in place. Nothing new to report.

1.2 Project Change Request

The UCSB Institute for Energy Efficiency (IEE) recently commissioned an experimental data center to support sustainability and the impact on climate change of digital infrastructure. Aristotle Science Team members studying the Food-Energy-Water Nexus are investigating the use of digital infrastructure to mitigate drought and climate change, particularly with respect to agriculture in the Central Valley of California. As part of these investigations, we are studying how the infrastructure that is being deployed to rural communities affects sustainability. In particular, if the ecological and climate change foot print associated with employing digital infrastructure exceeds the benefits brought by digital automation to climate change mitigation, the digital infrastructure will not improve sustainability. IEE has offered space in its experimental data center specifically to support the science teams.

To cite part of the Aristotle cloud in the IEE experimental data center requires approximately $50K in hardware expenditure. There are sufficient funds remaining in the effort to support this work, but they are currently available for staff support only. Thus, we requested permission to re-budget approximately $50K of staff funding to cover the siting costs in the IEE experimental data center to support the Aristotle Science Team's efforts. NSF project manager Amy Walton said we may reallocate costs from one line-item of the budget to another for allowable expenses without prior approval from NSF as long as the costs are not being reallocated from the line-item of Participant Support and it does not change the original scope of the project. We are proceeding with this reallocation of costs.

1.3 Project Execution Plan

The Project Execution Plan (PEP) was approved by NSF on 12/18/2015. We are operating as planned.

1.4 Project Planning Meetings

3/16/2021 status call:

- Our single node Singularity containers are working perfectly, but multi-node will not work on Stampede2 because the containers were built with Open MPI and Stampede2 doesn't allow Open MPI because of its Omni-Directional Interconnect. This is an important note to include in the lessons learned section of the technical report that we plan to write; that report will capture our experiences deploying containers on multiple platforms using multiple methods. Researchers had a similar experience with WRF which is highly sensitive to compilers and tools. We plan to focus on running benchmarks next and will include CONUS 12km and 2.5km.

3/30/2021 status call:

- We believe that a part of the research community (those who are not solely interested in running on the absolute fastest machine) would be willing to run slower in a cloud (10%? 15%?), if they got their answer 10% faster because they didn't have to wait in a queue. This, however, would require seamless access to the cloud for their application. Automation to this degree is difficult due to factors such as compilers. If this could be overcome, researcher productivity could be enhanced.

4/16/2021 status call:

- We are investigating other MPI applications that can leverage Terraform Kubernetes and use our parallel Python program with supporting packages. Creating a recipe for a container of this type is challenging.
- Our Slurm Virtual Cluster is up and running. The goal is to allow researchers to deploy a self-scaling environment—similar to a supercomputer with the Slurm job scheduler—in the cloud. We will use this method and others to deploy and run HPC benchmarks and, eventually, real science workloads in the cloud. We may develop a separate queue for X-Containers, as well as Singularity and Docker, with a specific image worker.

5/7/2021 status call:

- WRF was compiled on Stampede2.
- A former Aristotle REU student will be testing 3 container deployment methodologies this summer—Terraform and Kubernetes, Terraform and Ansible, and the Slurm Virtual Cluster.
- We are gaining a better understanding of how long it takes to containerize and deploy applications in the cloud vs. runtimes. Our focus is understanding what it takes to get things started and running in order to document general lessons learned for the community.
- We plan to update existing Terraform-based deployment methods and scripts to ensure ongoing usability and reproducibility.

6/4/2021 status call:

- We are documenting our Singularity-ready Slurm Virtual Cluster in the cloud and will include the location of the Ansible playbooks.
- We're also documenting Kubernetes lessons learned, including issues with Python+MPI. We would like to analyze job stats going out of Slurm on Stampede2; REMORA (REsource MOnitoring for Remote Applications) may be an option for seeing those performance metrics.

6/8/2021 status call:

- We ran Slurm on Stampede2 bare metal and are looking at the output. This will be compared to Slurm Virtual Cluster runs on Aristotle Red Cloud measuring floating point operations per second or turnaround time. Code issues can be due to memory parallelism rather than speed up parallelism. In addition, InfiniBand can create unpleasant network contention. We're also interested in understanding the impact of memory traffic on multi-tenant clouds. If L3 cache is polluted, performance will degrade. AWS claims they have a technique to minimize this impact, but the research community doesn't believe them. They believe that they need an XSEDE class machine in order to get floating point per second performance. Our thesis is the cloud doesn't have to win, but it has to be close, i.e., 10-15% slower with less queue time. Whether that is possible, is yet to be determined.

NSF
#1541215

2.0 Container Runtime Investigation

2.1 Accomplishments this Quarter

The Container Runtime Investigation team has been working on many different aspects of the project in parallel: orchestration of containers and HPC-style virtual clusters using containers in the cloud, running experiments of HPC applications both bare metal and containerized on supercomputing resources, as well as gathering recommendations and experiences to support our assessment of running HPC applications in containers. Experiments thus far for bare metal and containerized runs have been performed on Stampede2 using the Weather Research and Forecasting (WRF) model at software version 4.2.2, specifically the CONUS benchmarks 12km single node and the CONUS 2.5km multi-node. Additionally, scientific experiments are being performed to support the Pryor group's use case.

Due to the architecture of the system and requirements of the WRF Science Use Case code, we have chosen to compile with the Intel 18.0.2 compiler and leverage Intel MPI to make efficient use of the Omni-Path Interconnect for runs inside and outside of the container. This necessitated the development of a new container for WRF 4.2.2 including the Intel compilers and libraries. These have only recently become distributable in a container image due to Intel's release of the oneAPI HPC Toolkit as a free resource, including container images. Previously, we and the greater CI community relied on open source and free tools such as the GNU compilers and Open MPI for containerizing HPC applications, but these do not take full advantage of Intel CPUs, and may not work on certain architectures (as in our case).

Using the base container image provided by Intel, we built an Intel-compiled WRF 4.2.2 Docker container that is configured to use Intel MPI, which we have made freely available in a GitHub repository and as an image on Docker Hub. Singularity and X-Containers will also be developed, and we are exploring options to host these publicly. This container is novel, as far as our literature review has demonstrated, and thus will be very useful to the HPC and atmospheric science communities both to use directly and as an example for future work.

We have also continued development on our multi-cloud automated deployment method with Kubernetes using Terraform for deploying multi-VM MPI-capable HPC clusters with containers. Our deployment has been tested and used for experiments on Google Cloud, and is currently being extended to AWS and Azure in an easily configurable manner. While Kubernetes deployments have been performed before to deploy HPC applications (including WRF) in the cloud, all available literature seems to use the Open MPI implementation and GNU compilers. This means that our deployment method is novel as well, enabling a closer comparison between environments for container runtimes in the cloud and on HPC resources.

For OpenStack clouds, the team has been building on the XCRI Virtual Cluster toolkit to deploy and test the self-scaling multi-node cluster in order to perform experiments with Singularity containers in Red Cloud, the Cornell portion of the Aristotle Cloud Federation. Containerized benchmarks have been tested successfully, and features are being verified and expanded to support more HPC applications such as WRF.

2.2 Container Runtime Investigation: Future Plans

Using the X-Containers runtime has revealed some missing features. We are awaiting some implemented features from the team developing the runtime so that we can run in more environments. Pending these developments, we have Docker containers ready to test the runtime by importing into X-Containers. There have been several new developments within the Terraform framework; we are leveraging these capabilities some of which will greatly simplify the process of deploying an HPC-style cluster. Further, additional features are being implemented to support the ability to perform multi-node and large-scale (including large data) deployments of WRF 4.2.2 for both benchmarking and science use case purposes. All of our deployment methods and examples will be made open source and freely available to the wider community on GitHub, including documentation and convenient scripts for initiating, working with, and destroying the cluster.

In addition to the public version of the WRF 4.2.2 container, we are developing a private version that is nearly identical, but includes proprietary science code also compiled with Intel. Our goal is to perform large-scale multi-node runs of both the CONUS benchmarks and science codes using the WRF 4.2.2 containers in Google Cloud, AWS, Azure, Red Cloud, and on Stapmede2 (both bare metal and containerized) to provide a comprehensive analysis of these container runtimes across environments. We have begun collecting metrics from the runs already performed, and are expecting a full set of metrics for these runtimes and platforms to be included in our technical report. We also plan on pursuing publication of the novel work accomplished to communicate the results more broadly.

3.0 DIBBs Acquisition, Installation, Configuration, Testing & Maintenance Report

3.1 Hardware Acquisition

- Cornell and UB had no hardware acquisitions this quarter.
- Cornell's Slurm Virtual Cluster is running on Aristotle Red Cloud; developing queues for different deployment methods are under consideration.
- UCSB is integrating the last remaining hardware components into their Aristotle cloud. That process should be completed by July 1st. Per the discussion with PI Lifka and program manager Walton, we are planning to deploy 2 racks of Aristotle hardware in the Institute for Energy Efficiency's Experimental Data Center. The goal of this work is to allow the science teams to study the energy usage and carbon footprint of the ecology and agricultural systems, end-to-end. We will purchase racks, a network switch, and power distribution units. We have developed the specification for what needs to be installed and will start the purchasing and installation process this month.

3.2 Installation, Configuration, and Testing

- UB is finalizing the installation of the new OpenStack platform that VEXXHOST installed for us. Currently, we are working on the authentication part of it so that is fully integrated within our OpenID environment. We had to upgrade all the SSD drives (32 total) in our Ceph OSDs with larger drives to support a BlueStore backend. We ran into some delays on obtaining hard drives due to part shortages and manufacturing delays. Once the authentication part is completed, we will be in the burn in and testing phase. We expect this to last a few weeks and then we will move onto the migration stage where we will bring the VMs from the old cloud to the new. We expect to go into full production with the new cloud in mid-July.

#1541215

3.3 Federated Identity Management

Researchers use single sign-on at any member site.

3.4 Cloud Status by Site

The chart below shows each site's production cloud status.

	Cornell	Buffalo	UCSB
Cloud URL	https://redcloud.cac.cornell.edu	https://lakeeffect.ccr.buffalo.edu/ (access only to federation)	https://openstack.aristotle.ucsb.edu/
Status	Production	Production	Production
Software Stack	OpenStack	OpenStack	OpenStack
Hardware Vendors	Dell	Dell, Ace	Dell, HPE, DXC
DIBBs Purchased Cores	*616	**792	***740
RAM/Core	8GB	up to 8GB	9GB Dell, 10GB HPE
Storage	Ceph (1.6PB)	Ceph (768TB)	Ceph (720TB)
10gb Interconnect	Yes	Yes	Yes
Largest Instance Type	28core/192GB RAM	24core/192GB RAM	48core/119GB RAM
Globus File Transfer/End Points	Yes/CAC Home Directories	Yes/CCR Home and Projects Directories	Yes/POSIX UCSB Aristotle
Globus OAuth 2.0	Yes	Yes	Yes
Total Cores (DIBBs purchased cores + existing cores) = 2776	* 616 additional cores augmenting the existing Red Cloud (1316 total cores).	** 792 total cores (UB Lake Effect Cloud and CCR cloud will be one pool after upgrade).	***740 cores in UCSB Aristotle cloud (956 total cores, Aristotle is separate from UCSB campus cloud)

4.0 Cloud Federation Portal Report

Content updates to the project portal are ongoing (https://federatedcloud.org).

4.1 Software Requirements & Portal Platform

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

4.2 Integrating DrAFTS into the Portal

The *Aristotle AWS Pricing Tool* remains operational. This tool helps users compare Aristotle resources to the various AWS alternatives based on performance, cost, and price-performance. Visit https://federatedcloud.org/using/drafts.php to learn more or go directly to the Aristotle AWS Pricing Tool at http://169.231.235.92:5000/.

NSF
#1541215

4.3 Integrating Open XDMoD into the Portal

4.3.1 Application Kernels (AK) Containerization in the Cloud

AK containers are used in all Aristotle OpenStack instances within the XDMoD performance monitoring module. UB plans to run benchmarks on public clouds in July.

4.3.2 XDMoD Cloud Integration

All 3 sites are running Federated Open XDMoD 9.0.
.

4.4 Allocations & Accounting

The federation's database schema is now available to the broader cyberinfrastructure community (https://federatedcloud.org/using/buildyourown.php).

5.0 Research Team Support

5.1 Science Use Case Team Updates

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

The OUTSTEP Community Platform (https://outsteps.org/) is hosted on Aristotle and serves over 100 members. UB professor Varun Chandola and Penn State professor Alfonso Mejia are using the data collected on the platform to analyze network interactions between the different members.

An NSF proposal focused on developing a community for Lower Great Lakes sustainability research was submitted that makes use of our initial analyses. Results are pending after the reverse site visit. We are currently working on a paper that describes the digital platform and its capabilities.

Use Case 2: Global Market Efficiency Impact

Dominik Roesch and University of Utah collaborators Jonathan Brogaard and Matthew Ringgenberg completed a new version of their paper "Does Floor Trading Matter" that was presented at the *8th Annual Conference on Financial Market Regulation* and will be presented at the 49th European Finance Annual Meeting in Milan, Italy in August (https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3609007).

Target publication is the *Journal of Finance*.

NSF
#1541215

Use Case 3: Application of the Weather Research and Forecasting (WRF) Model for Climate-Relevant Simulations on the Cloud

PI: Sara C. Pryor, Cornell

New WRF Simulations:

Currently our effort is focused on attempting the implementation of a multi-node WRF simulation capability. I believe partial success has been achieved.

Analysis of Prior Simulations:

Most of our effort was focused on analysis of our existing simulations centered on three key themes:

1. *Use of machine learning approaches for wind gust detection and quantification.* I appointed a new REU for the summer and he will be continuing and extending the work on a new deep learning-based approach with a straightforward (physics-based) post-processing parameterization of wind gusts and statistical approaches that employ linear methods. Our initial research (reported in Coburn and Pryor, in review and also presented at the AGU conference on Natural Hazards and Machine Learning, May 2021) focused on Newark, Boston and Chicago O'Hare. Julian (the new UG research assistant) has begun to download of data from a distributed network of major airports across the US (see Table below) to examine generalizability of our findings. Our work in this arena benefits from the Aristotle architecture in that we are using bootstrapping with 1000 iterations of the artificial neural networks to establish confidence intervals on the network forms. This benefits from the use of multi-processor parallel processing.

Site	ID	Latitude (˚N)	Longitude (˚W)	Elevation (m)
Atlanta, Georgia	KATL	33.630	-84.442	307.9
Denver, Colorado	KDEN	39.833	-104.658	1650.2
Dallas, Texas	KDFW	32.898	-97.019	170.7
Los Angeles, California	KLAX	33.938	-118.389	29.6
Orlando, Florida	KMCO	28.434	-81.325	27.4
Fort Snelling, Minnesota	KMSP	44.883	-93.229	265.8
Phoenix, Arizona	KPHX	33.428	-112.004	337.4
SeaTac, Washington	KSEA	47.444	-122.314	112.8

2. *Analysis of our first set of simulations of wind farm wakes from the east-coastal offshore lease areas using ultra-high resolutions with WRF.* The simulations were performed on NERSC-Cori but were ported onto Aristotle for analysis. A presentation of the results was given at the Wind Energy Science Conference: Barthelmie R.J., Shepherd T.J. and Pryor S.C. (2021): Offshore wakes in the U.E. east coast lease areas. *Wind Energy Science Conference (WESC)* Hannover, Germany May 2021. A manuscript on this work is currently in review. Our work in this arena benefits from the Aristotle architecture in three key ways; (i) Availability of large RAM. The WRF output is multi-dimensional and large. (ii) Availability of multi-processors for analysis speed using parallel processing. (iii) Availability of large disk volumes so all WRF output can be hosted for analysis.

3. *Analysis of an ensemble of simulations of a historically important derecho (deep convection) event from June 2012.* This event comprised a line of self-organizing thunderstorms that moved over Washington DC causing massive economic damage and substantial loss of life. We have performed an 11-member ensemble of simulations (on Aristotle) with varying lateral boundary conditions, with and without nudging and varying microphysics schemes and are now evaluating those simulations to identify the simulation configuration that yields greatest fidelity. We are evaluating the simulation relative to ground-based measurements of:

    - Wind gusts
    - Presence/absence of HAIL/GRAUPEL
    - Precipitation accumulation.
    - Pressure (and temperature) for the cold pool (downbursts and outflow from the thunderstorms). And, also RADAR composite reflectivity (for the presence of thunderstorms) and RADAR precipitation accumulation.

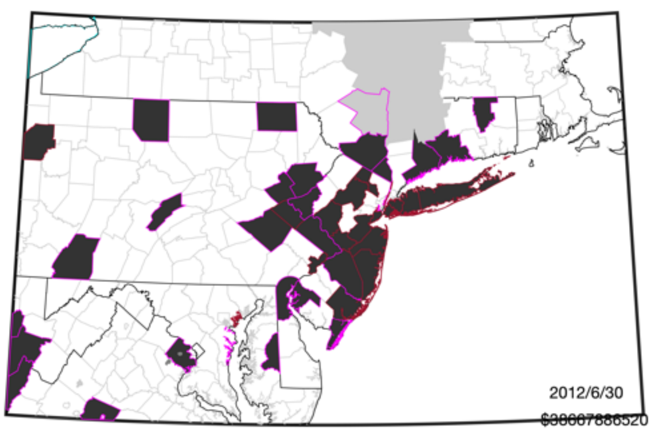

Figure 1. Data from the NOAA storm reports for 29 June-1 July 2012. Counties in black suffered economic losses of > $1000 (US$2012) and those outlined in red showed > 3 deaths.

WRF CI Accomplishments/Plans

CI Accomplishments:

We have been working on the development of cyberinfrastructure to support the Pryor group's computational needs: containerizing the WRF model and associated scientific software to perform large-scale analysis of Offshore Wind Farms, developing scalable cloud deployment methods, and performing various experiments. These experiments aim to verify and validate configurations for running WRF with supplemental science applications across resources, and to determine optimal configurations for performance on various systems in and out of a container. We have developed containers for various versions of WRF, most recently 4.2.2, and compiled with GNU and Intel compilers to perform tests. The most recent version of the EWP wind farm parameterization module, which is necessary for the Pryor group's work, is configured to work with WRF 4.2.2 and compiled with Intel; we have developed a container to support this. Experiments have been performed in Aristotle and on the Stampede2 supercomputer to support this work, with large-scale runs in development to be deployed on other platforms. The major development hurdles for this software set have been the difficulty in standardizing a working configuration for WRF across systems, and supporting the heavy computational demands of nested domains in WRF, which are needed for accurate simulations of wind turbine effects. We have resolved a standardized configuration by leveraging the Intel oneAPI HPC Toolkit within our containers to match the Intel compiler and Intel MPI modules on Stampede2 as closely as possible.

CI Plans:

We are still testing our implementation of the WRF and EWP combined codes with a profiler on various systems to ensure the computational demands of the application runs will be met on various systems. The Pryor group has also been providing scientific data and run configurations to aid in testing the software, containers, and deployments. We plan to continue combining our efforts to perform large-scale science runs. As we perform a variety of experiments, we are tracking a list of performance, I/O, memory, and other metrics that will be shared to help others in the atmospheric science and HPC communities to leverage our public implementations for further research.

Use Case 4: Transient Detection in Radio Astronomy Search Data

We migrated the container of radio astronomy software that combines pipeline components developed for pulsar and fast radio burst (FRB) detections to an XSEDE HPC resource. The Singularity container will be used by an astronomy graduate student on the large memory nodes of Bridges to process a large dataset as part of an upcoming XSEDE allocation request. The data originated from the Breakthrough Listen project that aims to detect potential signals from extraterrestrial civilizations but are also suitable for pulsar and FRB studies. This particular data set is for the direction to the Galactic Center, which harbors a supermassive black hole and likely has may neutron stars orbiting it.

We will be continuing to improve clustering methods (such as Friends-of-Friends) in the FRB pipeline and complete the comparison with other detection methods implemented in PRESTO with the modulation index calculation.

Use Case 5: Water Resource Management Using OpenMORDM

Bernardo Trinidade (now-graduated PhD student in the Reed Group) has provided CAC cloud systems engineer Bennett Wineholt a measurement script to run OpenMORDM with step timing to quantify overhead, which Wineholt may investigate.

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

"The Predicted Metabolic Function of the Gut Microbiota of *Drosophila melangogaster*" by Nana Y.D. Ankrah, Brandon E. Barker, Joan Song, Cindy Wu, John. G. McMullen II & Angela E. Douglas was published in *mSystems Journal* in May (https://journals.asm.org/doi/full/10.1128/mSystems.01369-20). They reconstructed and analyzed 31 metabolic models for every combination of the five principal bacterial taxa in the gut microbiome of *Drosophila*. This revealed that metabolic interactions between *Drosophila* gut bacterial taxa are highly dynamic and influenced by cooccurring bacteria and nutrient availability. Their results generate testable hypotheses about among-microbe ecological interactions in the *Drosophila* gut and the diversity of metabolites available to influence host traits.

NSF
#1541215

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

Citrus Frost Prevention
Lindcove Research and Extension Center, Exeter, CA:

At Lindcove, the Citrus Under Protective Screening (CUPS) facility experienced a major malfunction that will require significant repair. The Aristotle team gathered approximately 8 months of monitoring telemetry data that it will analyze while the repair is underway. Lindcove operations anticipates availability for tree planting next spring. The plan is for the science team to provide an analysis of the now inoperative CUPS to help inform the construction of the repaired facility.

UCSB Edible Campus:

The UCSB Edible Campus Food Security project allowed the Aristotle Science Team lead to tour the facility ahead of an anticipated reopening in late June. The goal of the visit was to plan for the installation of soil moisture sensors that the science team will use to implement precision irrigation on the farm.

Sedgwick Reserve:

Covid-19 restrictions on student visits to the Sedgwick Reserve remain in place (but are likely to be lifted soon). As a result, the science team has had only limited physical access for emergency maintenance of the equipment. During the hiatus, the team has prepared a new set of camera traps for deployment at Sedgwick. While no specific date for a visit can be scheduled at present, the team is working on the assumption that access will be possible by September 1st at the latest.

New Collaborations:

As the Aristotle project sunsets, the science team has been engaging in various outreach discussions with other collaborators who wish to leverage the Aristotle infrastructure and its software artifacts in the coming months.

These discussions have yielded 3 new collaborations:

1. The first is with the Woods Hole Oceanographic Institute and UCSB Marine Sciences with the goal of studying California kelp forest dynamics using remote sensing. The UCSB science team will be applying the techniques and technologies developed for Aristotle science projects to the task of gathering fusing, and analyzing data from the sea bed and atmosphere off the coast of California.

2. The second collaboration is with DOE researchers who are building ground-based cloud observation networks to study atmospheric cloud formation and lifecycle dynamics. The science team will be using Aristotle artifacts to deploy two new observing networks remotely—one at the Sedgwick Reserve and the other on Long Island in New York.

3. The third collaboration is with the Mathematic Department at Cal State Long Beach who wish to automate and scale their calculus curriculum. The Aristotle science team's expertise in using, maintaining, and scaling campus clouds for academic use is the basis for this education focused collaboration.

NSF
#1541215

These new collaborations have generated six new funding proposals: three to NSF, one to DOE, and two to an educational foundation that will hopefully become another avenue to carry forward the artifacts generated by the Aristotle project.

6.0 Community Outreach and Education

6.1 Community Outreach

- Technologies and techniques developed by the Aristotle project team are available on the portal and we will respond to future inquiries in order to share our experiences as requested by the CI and research communities (https://federatedcloud.org/using/technologies.php).
- The "New York zone" of Jetstream 2 (1,024 computer cores and 869TB storage) were installed at Cornell and performance tests were passed at the end of April. We will be installing OpenStack on the hardware during the summer and will connect it to the Aristotle Cloud Federation to enhance community access. This addition will be used to explore the federation of clouds and to make OpenStack enhancements that will be disseminated to the broader research community. We also recently tested and deployed Slurm-based MPI clusters on Aristotle and Jetstream. The availability of the New York zone will help to facilitate faster development and dissemination of cloud tools for research.
- Aristotle Use Case Scientist Sara C. Pryor and Rebecca Barthelmie's *TED-Ed* video provides the general public an animated description of how wind turbines work: https://ed.ted.com/lessons/how-do-wind-turbines-work-rebecca-j-barthelmie-and-sara-c-pryor

6.2 Education

- Information on how to create, access, and set up Linux and Windows instances; set up networks and share files; and, use the OpenStack Command-Line Interface is available at https://federatedcloud.org/using/gettingstarted.php.
- Documentation on how to build your own federation in the cloud has been completed and is available at https://federatedcloud.org/using/buildyourown.php. This includes hardware and software recommendations, Single Sign-On with Globus Auth, Aristotle account information, Globus user credential lookup sample code, and our database schema.