My Desktop
Prepare & Submit Proposals
New! Prepare Proposals (Limited proposal types)
Prepare Proposals in FastLane
Proposal Status
Awards & Reporting
Notifications & Requests
Project Reports
Award Functions
Manage Financials
Program Income Reporting
Grantee Cash Management Section Contacts
Administration
Lookup NSF ID

Preview of Award 1541215 - Annual Project Report

Cover |
Accomplishments |
Products |
Participants/Organizations |
Impacts |
Changes/Problems

Cover

Federal Agency and Organization Element to Which Report is Submitted:	4900
Federal Grant or Other Identifying Number Assigned by Agency:	1541215
Project Title:	CC*DNI DIBBs: Data Analysis and Management Building Blocks for Multi-Campus Cyberinfrastructure through Cloud Federation
PD/PI Name:	David A Lifka, Principal Investigator Thomas R Furlani, Co-Principal Investigator Richard Wolski, Co-Principal Investigator
Recipient Organization:	Cornell University
Project/Grant Period:	10/01/2015 - 09/30/2021
Reporting Period:	10/01/2019 - 09/30/2020
Submitting Official (if other than PD\PI):	N/A
Submission Date:	N/A
Signature of Submitting Official (signature shall be submitted in accordance with agency specific instructions)	N/A

Accomplishments

*** What are the major goals of the project?**

The major goals established for the Aristotle project were:

- Implement a scalable and sustainable multi-institutional cyberinfrastructure cloud federation model that provides data analysis building blocks in support of multiple research disciplines requiring flexible workflows and analysis tools for large-scale data sets. Federation sites are Cornell University, University at Buffalo (UB), and University of California, Santa Barbara (UCSB).
- Support seven strategic science use cases from intentionally diverse disciplines (earth and atmospheric science, finance, chemistry, astronomy, civil engineering, genomics, and agriculture) to demonstrate the potential of a federated cloud as a campus bridging paradigm. Explore data analysis techniques and their applicability to different disciplines. Document tools, workflows, challenges, and best practices for each use case.
- Encourage and reward data analysis resource sharing with a new allocations and accounting model that provides a fair exchange mechanism for resource access between and across multiple institutions. Develop and build a new tool for cloud metrics into Open XDMoD and develop DrAFTS (Durability Agreement From Time Series) statistics to make online forecasts of performance.

*** What was accomplished under these goals (you must provide information for at least one of the 4 categories below)?**

Major Activities:

Architect, install, and maintain PY5 Aristotle cloud and storage assets

Specific Objectives:

- Cornell added 6 servers to their Ceph storage pool to increase their storage capacity to 1.6PBs, 20 NVIDIA T4s and 4 V100s to support GPU computing (non-DIBBs funds), and upgraded Red Cloud's network to increase bandwidth.
- UB added servers, switches, and PDUs to support production and development clouds, a data transfer node with 100GB interface to improve data transfers, and NVIDIA T4 and V100 GPUs.
- UCSB added 6 servers (288 cores, 4608TB RAM) and 768TB Ceph storage.
- A total of 3,064 cores (includes campus resources), 41 GPUs, and 3.09PB storage is available at Cornell, UB, and UCSB.
- Dartmouth installed their OpenStack 16.1 cloud and will be connecting it to the federation.

Develop application kernel containers for performance monitoring and benchmarking

- Created 7 AK containers for HPCC, HPCG, IOR, MDtest, NAMD, NWChem, and Enzo that automatically detect the most suitable executable, number of cores to use, and then execute the application with provided input parameters and output the results. AK containers are used in all Aristotle OpenStack instances within the XDMoD performance monitoring module.

Investigate the performance and considerations for running Docker, Singularity, and X-Containers

- Developed application containers that are convertible from Docker to Singularity and verified functionality on a variety of cloud platforms.
- Implemented reproducibility of builds and environments within containers by leveraging the Nix package manager.
- Determined a set of recommendations for the conversion of containers from Docker to Singularity for XSEDE HPC machines.
- Created a multi-cloud automated deployment method using Terraform and Ansible for deploying multi-VM MPI-capable clusters with Docker.
- Developed an automated deployment method for Kubernetes using Terraform and created MPI-capable clusters (tested on GCP).
- Developed a cluster benchmarks container in both Docker and Singularity and tested on multiple platforms.

Investigate cloud marketplace

- Concluded that the time resolution and data granularity of RightScale's cost tracking tool is not well-suited to analyzing the usage of particular applications or runs, making it difficult to profile individual science workflows.
- Decided to pursue the cost analysis of multi-VM MPI-capable clusters which appear to be of the most interest to researchers.
- Published work-to-date on the creation of multi-VM MPI environments in multiple clouds at *PEARC '20* and will continue this work with full cost analysis details in a forthcoming paper.

Create price comparison tool to help users decide which Aristotle or AWS instances to run on

- Developed DrAFTS 2.0 tool that runs the TOP500 LINPACK Benchmark on all Aristotle instance types and all AWS instance types, and sorts the results by cost and performance so users can ask questions such as which AWS instance type is most equivalent to an instance type in Aristotle, if I want to spend 20% more and go 30% faster, which instance type do I use, etc. Completion of this tool is anticipated during the no cost extension period.

Update portal content and share

- Updated the portal user guide, publications, use case products, and Aristotle news on a regular basis and shared the Aristotle skeleton and database with the research community.

Support 7 science use case teams

1. **A Cloud-Based Framework for Visualization & Analysis of Big Geo Data** (Varun Chandola, UB). Created a new platform called OUTSTEPS to support a network of sustainability scholars, practitioners, stakeholders and students in the Lower Great Lakes. OUTSTEPS runs on top of Aristotle and allows participants to interact and share information and will be merged with webGlobe to share geospatial datasets and workflows. The OUTSTEPS team (8 universities and 100 practitioners) recently submitted a Civic Innovation Challenge planning grant to NSF to develop new ways to increase community resilience to natural disasters such as COVID-19.
2. **Global Market Efficiency Impact** (Dominik Roesch, UB). Roesch, 3 PhD students, CS students, and collaborators from UCLA, U. of Utah and others are running VMs that hold extensive finance data on Aristotle. The availability of this data and a financial framework resulted in two *Journal of Financial Economics* publications: "Asset pricing: A tale of night and day" and "Tick size, liquidity for small and large orders, and price informativeness."
3. **Application of the Weather Research and Forecasting (WRF) Model for Climate-Relevant Simulations on the Cloud** (Sara C. Pryor, Cornell). Using the WRF model, the Pryor team assessed the possible impact of climate variability, inter-annual variability in power production, and the conditions in which wind turbines operate. They also evaluated the sensitivity of wind farm wake effects and power production whose results have important applications to projections of annual energy production from new wind turbine arrays constructed in the wind shadow from existing wind farms. These and other analyses resulted in 9 journal publications and 4 conference presentations.
4. **Transient Detection in Radio Astronomy Search Data** (James Cordes, Cornell). Implemented the workflow for the Friends-of-Friends (FOF) algorithm in parallel form. Completed and tested a single container of radio

astronomy software that combines the pipeline components developed for Pulsar and other transient detections that can be deployed either on the cloud with Docker or on an XSEDE HPC resource with Singularity. This work was included in a *PEARC'20* publication. An REU student will continue to work towards a new FRB detection during the academic year.

5. **Water Resource Management Using OpenMORDM** (Patrick Reed, Cornell). Performed WaterPaths runs at scale on Aristotle Red Cloud using up to 224 CPU cores on 8 worker VMs organized into an on-demand MPI cluster with Docker container-based software deployment. This accomplishment was published in *Environmental Modellling and Software*. The paper contains benchmarks for hybrid MPI and OpenMP workflows showing high levels of performance for Aristotle relative to Stampede and a local HPC cluster.
6. **Mapping Transcriptome Data to Metabolic Models of Gut Microbiota** (Angela Douglas, Cornell). Created a VM image resulting in publication of "Syntrophic splitting of central carbon metabolism in host cells bearing functionally different symbiotic bacteria" in *ISME Journal.* Integrated data from all prior analyses and submitted "The predicted metabolic function of the gut microbiota of Drosophila melanogaster" to *Cell Reports*. Developed new algorithms, a method that employs SteadyCom with discrete time steps (*Semi Dynamic SteadyCom*), and a modeling framework that extends the COBRA Toolbox for MATLAB to make large scale simulations more scalable and reliable (*hCOBRA)* so scientists can better navigate the complexity of metabolic modeling and the recursive structure of simulations with priority effects.
7. **Multi-Sourced Data Analytics to Improve Food Production and Security** (Kate McCurdy, Sedgwick; Elizabeth Grafton-Cardwell, UC Lindcove; Chandra Krintz, UCSB). Provided farm operation managers citrus frost prevention data collected and analyzed by Aristotle and analyzed differential irrigation data. Built edge-based computational infrastructure to support Citrus Under Protective Screens (CUPS), a potential remedy for citrus greening disease which is now threatening CA. Developing a hybrid machine-learning and CFD model for CUPS. An REU student created a telemetry data visualizer that works as a local tool or as a service hosted in Aristotle; it will support CUPS and the Edible Campus farm which has been instrumented with a weather sensor data acquisition capability running on Aristotle. "CSPOT: portable, multi-scale functions-as-a-service for IoT" is one of 3 publications this reporting period.

See Aristotle quarterly reports for details: https://federatedcloud.org/reports/.

Significant Results:

Key outcomes or
Other achievements:

*** What opportunities for training and professional development has the project provided?**

Cross-Training & Knowledge Sharing

Expertise was shared between sites every two weeks on Aristotle team conference calls to ensure timely cross-training and knowledge sharing, and in-depth follow-up calls occurred to solve specific technology implementation issues and to share lessons learned.

Cross-site discussions and collaborations between use case scientists and Aristotle support teams facilitated the sharing of research tools and cloud computing technologies, e.g., sharing insights on Docker security and

networking, Red Hat TripleO implementation, most cost-effective GPU selection, etc.

Knowledge gained from events such as the monthly Zoom-based "Cornell Cloud Forum 2020" conference series were beneficial to federation staff and the wider community, and included sessions on cloud enablement success stories, cloud community governance, research on the public clouds, science breakthroughs in the cloud, cloud strategy, and trust and security.

We continued to provide advice and guidance to Dartmouth College who is deploying their first OpenStack cloud for research. Dartmouth joined the federation as its first member. They are providing us with feedback on Aristotle's cloud implementation documentation that we constantly refine to ensure that we are sharing best practices with the research community. The Aristotle team also hosted Binghamton University's CIO and research IT staff to educate them on federation and OpenStack technologies.

Cornell partnered with Indiana University's Jetstream 2 project and will deploy 1,024 computer cores and 896TB storage. We are currently reviewing hardware specifications. The "New York zone" will be used to explore federation clouds and to make OpenStack enhancements that will be disseminated to the broader research community. We will also draw on our Aristotle experience to create campus software so campuses can set up their own clouds.

Undergraduate & Graduate Student Development

Five REU students made valuable contributions to Aristotle science use cases and gained first-hand experience using clouds for data analysis.

REU student Matthew Farnese improved the pipeline used to identify Fast Radio Bursts (FRB) by testing the functionality of the Friends-of-Friends algorithm, extending the existing pipeline methods for use with multiple data formats, creating new plotting methods, and adding features to make the pipeline more user-friendly.

In response to the research question "How can Machine Learning frameworks and/or computational fluid dynamics (CFD) solvers be provisioned reliably as a GPU-accelerated resource?" Priyanka Dilip created Cornell Red Cloud images containing TensorFlow, MATLAB, PyTorch, and Jupyter. These images were developed in an Anaconda environment and Dockerized for Forest Large Eddy Simulations and cryo-electron microscopy. She also containerized an OpenFOAM CFD solution, developed Dockerfile and images for RapidCFD, created GitHub test applications for Tensorflow+Keras and PyTorch, and wrote documentation to help others create server images and use GPUs more effectively.

Cornell undergraduate student Jeffrey Lantz wrote a guide on how to get started with Kubernetes, updated a Terraform-Ansible tool, and created a Terraform-Kubernetes tool. He then used High Performance LINPACK Benchmarks to compare the cost and efficiency of the two tools. Lantz concluded that the Terraform-Kubernetes tool is faster to deploy and less costly to deploy as a cloud computing cluster.

Mentored by Cornell professor Sara C. Pryor, REU Sherri Tan worked on predicting the occurrence and magnitude of wind gusts at major U.S. airports. She downloaded datasets of predictors such as upper air variables using Python scripts, and then aligned those with datasets for the observed occurrence and magnitude of wind gusts at the airports using MATLAB. Tan then used MATLAB's functions and toolboxes for generalized linear regression, stepwise regression, and deep learning to build predictive models and calculate descriptive statistics.

Finally, UCSB undergraduate Kerem Celik developed a telemetry data visualizer for the Citrus Under Protective Screening project (CUPS) and the UCSB Edible Campus Program farm.

During spring 2020, the Aristotle cloud was among the resources supporting a Computer Science class taught by Wolski on *Cloud Computing, Edge Computing, and IoT (CS293B - https://sites.cs.ucsb.edu/~rich/class/cs293b-cloud/)*. The learning objective was to build a working prototype of either an application that uses edge and cloud computing technologies or a new edge/cloud technology itself. For example, several students developed video analysis applications to help identify when people were visiting the Edible Campus farm without an essential farming activity to perform.

At UB, Mohammed Zia used the Aristotle cloud service to host JupyterHub with the nbgrader module enabled on an Ubuntu 18.04 image for his *Programming and Database Fundamentals for Data Scientists (EAS 503)* class. The nbgrader module allows him to create and grade programming assignments in JupyterHub, a streamlined method that enables more frequent distribution of assignments (increased from 4 to 13). Aristotle is cost-effective because Zia can choose the instance type based on the expected class size. There are usually 120 students in the Fall semester and 50 students in the Spring semester. Zia selects the instance type he needs based on the class size, and after the semester is over, deletes the instance. Ease of deployment and the ability to handle all of the student's interactive Python needs in one place are very beneficial.

Many of the Aristotle science use cases impact graduate and PhD student skills development. For example, Aristotle use case lead Dominik Roesch taught UB finance PhD and CS students to use the OneTick Time-Series financial framework with the underlying data hosted on Aristotle; the students subsequently launch new investigations such as an analysis of liquidity for large and small orders.

PhD, postdocs, and graduate students at Cornell, UB, and UCSB learn about cloud and edge computing technologies when collaborating with Aristotle science use case staff, particularly when working on application containerization, library functions, modeling software, multiple VM execution, deployment automation, and algorithmic issues. A PhD in CS was awarded to Nevena Golubovic for her work on an Aristotle use case; this work was reflected in her dissertation: "Scalable analytic systems for multi-tier IoT deployments with applications in agriculture" (https://escholarship.org/uc/item/36b678sn).

Use case scientist Nana Ankrah attended the Annual Biomedical Research Conference for Minority Students (ABRCMS) where he was a conference presentation judge and had many opportunities to provide formal and informal mentoring for undergraduate students. Ankrah also participated in the 2019 SACNAS National Diversity in STEM conference where he gave a presentation on Aristotle-related research titled "Gut microbiome metabolism and host health."

How-To Documentation & Training

How-to user guides continue to be developed on GitHub and, when ready, publicly released on the Aristotle portal (https://federatedcloud.org/using/gettingstarted.php). Container investigation products (methods, tools, images, etc.) will be added to the how-to user guides as the investigation progresses. Build-your-own federation documentation is nearly complete and will be updated for community use (https://federatedcloud.org/using/buildyourown.php).

*** How have the results been disseminated to communities of interest?**

Scientific Meetings, Publications, and Conferences

Researchers and the Aristotle technical team presented results at scientific meetings where they referenced the project and its contribution to their success. They also published papers acknowledging the project.

The "Products" section of this report lists 26 selected publications, conference presentations, and videos delivered in the past year (https://federatedcloud.org/about/publications.php). For example, the *Humans Advancing Research in the Cloud (HARC'19)* publication entitled "Red Cloud and Aristotle: campus clouds and federations" shared our experience with federated cloud components, science use cases, and container technologies with the cyberinfrastructure community (https://federatedcloud.org/papers/RedCloudAndAristotle.pdf) and a 2020 *Journal of Applied Meteorology and Climatology* publication by Sara C. Pryor et al.—"Wind theft from onshore wind turbine arrays: Sensitivity to wind farm parameterization and resolution"—is one of many science impacts enabled in part by Aristotle computational resources and our use case support team (https://journals.ametsoc.org/jamc/article/59/1/153/346167).

Aristotle PI and co-PIs are actively involved in professional societies and conferences such as the Practice & Experience in Advanced Research Computing (PEARC) conferences, Coalition for Academic Scientific Computing (CASC) meetings, SC Conferences, and NSF workshops all of which have afforded opportunities to share progress on the project. For example, at *PEARC'20* the Aristotle team made presentations on "Reproducible and Portable

Workflows for Scientific Computing and HPC in the Cloud" and participated in a sold-out tutorial with Indiana University colleagues on "Deep Dive into Constructing Containers for Scientific Computing and Gateways," and at the *2020 NSF CSSI Meeting*, the Aristotle team presented a poster on Aristotle goals, new federation members, federated cloud metrics, new technologies, and advances in science (https://federatedcloud.org/papers/WolskiPoster2020.pdf).

PI David Lifka has a leadership role in the eXtreme Science and Engineering Discovery Environment (XSEDE) project and keeps XSEDE management abreast of developments in cloud computing. In addition, Aristotle portal lead Mehringer is the XSEDE training lead and Aristotle container project lead Rich Knepper is the XSEDE manager for Community Resource Integration (CRI). Each of these roles afford opportunities to share cloud computing insights and experiences with cyberinfrastructure colleagues.

At the *SC19 Conference* in Denver, Cornell and UB featured the Aristotle project in their exhibits, conveyed science impacts, and shared the latest news including "OpenStack cloud implementation toolkit under development by XSEDE and Aristotle" and "Dartmouth joins Aristotle to explore the federated cloud computing model" (https://federatedcloud.org/science/CornellNewsHighlightsSC19.pdf). The project team also met with industry partners, universities, and OEMs and ISVs to brief them on cloud federation. Aristotle infrastructure team lead and Dell XL HPC Consortium chair Resa Reynolds presented "Federated Clouds: The Aristotle Project" at the *Dell XL HPC Meeting* in Scottsdale.

At the *Arm Research Summ*it in Austin, co-PI Rich Wolski's keynote presentation "Devices-as-Services and the Internet as a Platform of Things" featured Aristotle use cases and the Serverless Platform of Things (C-SPOT) technology that runs on Aristotle as well as public clouds (https://www.youtube.com/watch?v=97FbH4Kob_o). This technology was also shared at the *ACM Symposium on Edge Computing* in Washington, DC.

In Seattle, Aristotle use cases were highlighted at the *2020 American Association for the Advancement of Science Annual Meeting* by UCSB professor Chandra Krintz who presented her work on SmartFarm, an open source, hybrid cloud approach to agriculture analytics.

Aristotle Portal

The Aristotle portal provides scientists and the cyberinfrastructure community extensive information on the project, including detailed use case accomplishments/plans/products, publications, news and events, user guides and other documentation, and in depth quarterly project reports (https://federatedcloud.org/).

Communicating to a General Audience

News stories about Aristotle use case scientists broaden the public understanding of the value of IT and engineering in solving societal problems and the impact of mentoring programs such as the NSF REU program.

- October 2019 *Daily News* story highlights the Edible Campus Program at UCSB which is integrating Aristotle developed and hosted technologies to monitor weather and other factors that impact the farm's productivity and food security: https://dailynexus.com/2019-10-17/edible-campus-program-unveils-new-farm-on-west-campus/
- November 2019 *Data Center Dynamics* news story highlights Aristotle DrAFTS 1.0 analysis of the change in AWS Spot pricing: https://www.datacenterdynamics.com/en/analysis/amazons-spotty-pricing/
- February 2020 *Cornell Chronicle* story generated from an Aristotle acknowledged publication describes how the U.S. can meet 2030 wind-energy goals by quadrupling turbines: https://news.cornell.edu/stories/2020/02/quadrupling-turbines-us-can-meet-2030-wind-energy-goals);
- September 2020 *HPCwire* story highlights how Cornell students were immersed in the latest cloud technologies thanks to the NSF's REU program and includes quotes about their experiences and lessons learned: https://www.hpcwire.com/off-the-wire/cornell-students-immersed-in-latest-cloud-technologies-thanks-to-nsf-research-experiences-for-undergraduates-program/.

*** What do you plan to do during the next reporting period to accomplish the goals?**

Infrastructure improvements and sharing

- Continue to improve the installations at each site to provide a stable environment to complete the investigation of containers and select science use case cloud access.
- Improve OpenStack operations and share lessons learned

Aristotle portal dashboard and content updates

- Maintain the portal dashboard and continue to make regular content updates (how-to guides, etc.)

Application kernel performance comparisons

- Deploy containerization kernels on a public cloud provider(s).
- Benchmark and compare respective measurements from Aristotle, UB's HPC cluster, and several XSEDE resources to the public cloud(s).

DrAFTS 2.0 price comparison tool

- Complete the tool and share it with the community to help users decide which Aristotle or AWS instance to run on based on cost and performance.

Docker, Singularity, and X-Containers performance/user considerations investigation

- Complete benchmarking runs of an MPI sample science application (most likely WRF) on bare metal and containerized with Singularity on an HPC machine, and containerized with Docker on multiple clouds.
- Explore and report on Kubernetes orchestration and considerations for creating and managing workflows with each runtime.
- Explore interoperability between Docker, Singularity, and X-Containers.
- Complete a technical report detailing the performance and considerations for running Docker, Singularity, and X-Containers.
- Share the technical report, lessons learned, and any practical outcomes with the wider scientific community.

Cloud marketplace investigation

- Continue work on the creation of a reproducible container deployment for the WRF weather modeling Aristotle use case which is extensible to the broader atmospheric community.
- Write a paper that captures costing information for comparison of different resources by the user community.

Aristotle cloud access

- Provide use case scientists with Aristotle cloud access (with limits) during the no cost extension period. In-depth consulting support will be provided only if supported by non-DIBBs funding.

Use case scientist plans

1. **A Cloud-Based Framework for Visualization and Analysis of Big Geo Data** (Varun Chandola, UB). Use Aristotle to host the OUTSTEPS integrated community platform and the analytical ecosystem (webGlobe). OUTSTEPS is a multi-institution collaboration focused on sustainability activities around the Lower Great Lakes region.
2. **Global Market Efficiency Impact** (Dominik Roesch, UB). Roesch and U. of Utah collaborators plan to use Aristotle and the financial framework to investigate whether human traders still matter at a time in which trading is dominated by computers. For that we exploit the exogenous closure of the NYSE Floor due to COVID-19. Other projects include a LASSO regression predicting short-term stock returns using the whole cross-section of international stocks and an investigation of CumEx trading (labeled as the "Biggest Tax Heist Ever" by the *NYTimes*). All these projects require large datasets and computational power.
3. **Application of the Weather Research and Forecasting (WRF) Model for Climate-Relevant Simulations in the Cloud** (Sara C. Pryor, Cornell). The Pryor team have recently begun very high-resolution simulations of deep convection and the associated environmental hazards (see Letson et al. 2020). One example is derechos, fast-moving, damaging deeply convective systems, associated with tornadoes, wind gusts, very heavy precipitation and hail. These events have traditionally been very difficult to simulate even with the advanced Weather Research and Forecasting (WRF) model due to the physical complexity and multi-scale nature of these phenomena and resulting high compute demands. Thus, little work has been performed to assess model fidelity

for these transient events and/or identify optimal WRF configurations. The Pryor team also was recently awarded a new grant to examine the wind resources and optimize wind farm layouts for the offshore wind turbines along the U.S. east coast. This work strongly leverages our previous simulations on Aristotle and will entail additional WRF simulations with wind farm parameterizations enabled. The goal is to define the optimal density of wind farms to optimize system-wide power production and minimize the levelized cost-of-energy.

4. **Transient Detection in Radio Astronomy Search Data** (James Cordes, Cornell). We will continue to use Aristotle for the prototyping of FRB search codes and the development of deployable Docker and Singularity containers for deployment on large-scale XSEDE resources and public clouds.
5. **Water Resource Management Using Python Rhodium Framework** (Patrick Reed, Cornell). David Gold is continuing the work of Bernardo Trindade on MORDM and WaterPaths. It would be interesting to build on our initial scaling analysis for WaterPaths and use Aristotle to support our water supply infrastructure investment and portfolio management analyses. The allocation will be used to more fully evaluate the core steps of our workflow: (1) generating alternatives with multi-objective optimization, (2) ensemble Monte Carlo evaluation of candidate alternatives in a broader sampling of states-of-the-world, (3) elicit acceptable limits and requirements on the range of performance attained on key measures of performance, (4) global sensitivity analysis to discover what factors and scenarios cause system failures. The workflow can be tested on a range of broader case studies such as the Sedento Valley benchmarking system, N.C. Research Triangle, and the Tampa Bay Water Authority. We can use a range from modest to larger allocations subject to the difficulty and expense of the workflow.
6. **Quantifying the Metabolic Determinants of the Fate of Probiotic Bacteria Administered to an Animal Gut Microbiota** (Angela Douglas, Cornell). We are currently working on (1) integrating functionality from the Funflow library into hCOBRA to allow caching of workflow results in a distributed setting, (2) adding R-based analyses to our hCOBRA workflows through the HaskellR library. We are expanding our analysis to investigate priority effects in a 5-member Drosophila gut microbe community comprising three Acetobacters and two Lactobacilli, which will require continued computational resources.
7. **Multi-Sourced Data Analytics to Improve Food Production and Security** (Kate McCurdy, Sedgwick Reserve; Elizabeth Grafton-Cardwell, UC Lindcove; Chandra Krintz, UCSB SmartFarm). The science team plans to continue to use Aristotle to model agricultural conditions at Sedgwick, UCSB Edible Campus, and Lindcove in order to optimize food production and land use. Aristotle time will be used to run large scale image processing algorithms and CFD simulations that are parametrized by sensor data gathered from the various sites in real time. Aristotle will also be used to share large scale data sets (both sensor data and the result of simulation runs) as well as image and video corpuses between the science team researchers and their collaborators.

Plans to Disseminate Results

The Aristotle team will continue to engage the CI community through presentations and dialogue at *CASC, PEARC 21, SC21* and future scientific meetings such as the *International Symposium on Microarchitecture*, October 17-21, 2020. Aristotle PI/co-PIs will respond to project inquiries and keep leadership at XSEDE and relevant NSF projects abreast of new developments. The portal will highlight project results, and be updated regularly with news, events, and how-to user guides. Aristotle users are also spreading the word on what applications are effective in the cloud and best practices. Interest at scientific meetings in their cloud experience has been high as many research groups throughout the U.S. realize that cloud computing has a complementary role to play in their cyberinfrastructure portfolios and they need to plan for that.

Products

Books

Book Chapters

Inventions

Journals or Juried Conference Papers

Bernardo C. Trindade, David F. Gold, Patrick M. Reed, Harrison B. Zeff & Gregory W. Characklis (2020). Water pathways: An open source stochastic simulation system for integrated water supply portfolio management and infrastructure investment planning. *Journal*. 132 . Status = PUBLISHED; Acknowledgment of Federal Support = Yes ; Peer Reviewed = Yes ; DOI: 10.1016/j.envsoft.2020.104772

Fatih Bakir, Rich Wolski, Chandra Krintz & Gowri Ramachandran (2019). Devices-as-services: Rethinking scalable service architectures for the internet of things. *Juried Conference Paper*. . Status = PUBLISHED; Acknowledgment of Federal Support = Yes ; Peer Reviewed = Yes

Frederick Letson, Rebecca J. Barthelmie & Sara C. Pryor (2019). RADAR-derived precipitation climatology for wind turbine blade leading edge erosion. *Journal*. . Status = PUBLISHED; Acknowledgment of Federal Support = Yes ; Peer Reviewed = Yes ; DOI: 10.5194/wes-2019-43

Frederick Letson, Tristan J. Shepherd, Rebecca J. Barthelmie & Sara C. Pryor (2020). Modelling hail and convective storms with WRF for wind energy applications. *Journal*. 1452 . Status = PUBLISHED; Acknowledgment of Federal Support = Yes ; Peer Reviewed = Yes ; DOI: 10.1088/1742-6596/1452/1/012051

Kee H. Chung, Albert J. Lee & Dominik Rosch (2020). Tick size, liquidity for small and large orders and price informativeness: Evidence from the Tick Size Pilot Program. *Journal*. 136 (3), 879. Status = PUBLISHED; Acknowledgment of Federal Support = Yes ; Peer Reviewed = Yes ; DOI: 10.1016/j.jfineco.2019.11.004

Kyle Carson, John Thomason, Rich Wolski, Chandra Krintz & Markus Mock (2019). Mandrake: Implementing durability for edge clouds. *Juried Conference Paper*. . Status = PUBLISHED; Acknowledgment of Federal Support = Yes ; Peer Reviewed = Yes

Michael Zhang, Chandra Krintz & Rich Wolski (2020). STOIC: Serverless teleoperable hybrid cloud for machine learning applications on edge device. *Juried Conference Paper*. 1. Status = PUBLISHED; Acknowledgment of Federal Support = Yes ; Peer Reviewed = Yes ; DOI: 10.1109/PerComWorkshops48775.2020.9156239

Michael Zhang, Chandra Krintz, Rich Wolski & Markus Mock (2019). Seneca: Fast and low cost hyperparameter search for machine learning models. *Juried Conference Paper*. . Status = PUBLISHED; Acknowledgment of Federal Support = Yes ; Peer Reviewed = Yes

Nana Y. Ankrah, Rebecca A. Wilkes, Freya Q. Zhang, Dantong Zhu, Tadeo Kaweesi, Ludmilla Aristilde & Angela E. Douglas (2020). Syntrophic splitting of central carbon metabolism in host cells bearing functionally different symbiotic bacteria. *Journal*. 14 1982. Status = PUBLISHED; Acknowledgment of Federal Support = Yes ; Peer Reviewed = Yes ; DOI: 10.1038/s41396-020-0661-z

Nevena Golubovic, Rich Wolski, Chandra Krintz & Markus Mock (2019). Improving the accuracy of outdoor temperature prediction by IoT devices. *Juried Conference Paper*. 117. Status = PUBLISHED; Acknowledgment of Federal Support = Yes ; Peer Reviewed = Yes ; DOI: 10.1109/ICIOT.2019.00030

Peter Vaillancourt, Bennett Wineholt, Brandon Barker, Plato Deliyannis, Jackie Zheng, Akshay Suresh, Adam Brazier, Rich Knepper & Rich Wolski (2020). Reproducible and portable workflows for scientific computing and HPC in the cloud. *Juried Conference Paper*. 311. Status = PUBLISHED; Acknowledgment of Federal Support = Yes ; Peer Reviewed = Yes ; DOI: 10.1145/3311790.3396659

Rich Wolski, Chandra Krintz, Fatih Bakir, Gareth George & Wei-Tsung Lin (2019). CSPOT: portable, multi-scale functions-as-a service for IOT. *Juried Conference Paper*. 236. Status = PUBLISHED; Acknowledgment of Federal Support = Yes ; Peer Reviewed = Yes ; DOI: 10.1145/3318216.3363314

Sara C. Pryor, Tristan J. Shepherd, Melissa Bukovsky & Rebecca J. Barthelmie (2020). Assessing the stability of wind resource and operating conditions. *Journal*. 1452 . Status = PUBLISHED; Acknowledgment of Federal Support = Yes ; Peer Reviewed = Yes ; DOI: 10.1088/1742-6596/1452/1/012084

Sara C. Pryor, Tristan J. Shepherd, Patrick J.H. Volker, Andrea N. Hahmann & Rebecca J. Barthelmie (2020). “Wind theft” from onshore wind turbine arrays: Sensitivity to wind farm parameterization and resolution. *Journal*. 59 (1), 153. Status = PUBLISHED; Acknowledgment of Federal Support = Yes ; Peer Reviewed = Yes ; DOI: 10.1175/JAMC-D-19-0235.1

Stratos Dimopoulos, Chandra Krintz & Rich Wolski (2019). Towards distributed, fair, deadline-driven resource allocation for cloudlets. *Juried Conference Paper*. 7. Status = PUBLISHED; Acknowledgment of Federal Support = Yes ; Peer Reviewed = Yes ; DOI: 10.1145/3366614.3368102

Stratos Dimopoulos, Chandra Krintz & Rich Wolski (2020). Fair scheduling for deadline-driven, resource-constrained, multi-analytics workloads. *Juried Conference Paper*. 261. Status = PUBLISHED; Acknowledgment of Federal Support = Yes ; Peer Reviewed = Yes ; DOI: 10.1109/ICNC47757.2020.9049802

Terrence Hendershott, Dmitry Livdan & Dominik Rosch (2020). Asset pricing: A tale of night and day. *Journal*. 136 (3), . Status = PUBLISHED; Acknowledgment of Federal Support = Yes ; Peer Reviewed = Yes ; DOI: 10.1016/j.jfineco.2020.06.006

Tristan J. Shepherd, Rebecca J. Barthelmie & Sara C. Pryor (2020). Sensitivity of wind turbine array downstream effects to the parameterization used in WRF. *Journal*. 59 (3), 333. Status = PUBLISHED; Acknowledgment of Federal Support = Yes ; Peer Reviewed = Yes ; DOI: 10.1175/JAMC-D-19-0135.1

Licenses

Other Conference Presentations / Papers

Chandra Krintz (2018). *Adventures and opportunities in cyber-physical systems and research*. 2018 International Conference on Computer Aided Design. San Diego, CA. Status = PUBLISHED; Acknowledgement of Federal Support = Yes

Sayed Mohammed Arshad Zaidi, Varun Chandola, Melissa R. Allen & Bhudhendra L. Bhaduri (2018). *Anomaly detection in Energy-Water Nexus: webGlobe – A cloud-based geospatial analysis framework for interacting with climate data*. American Geophysical Union Fall 2018 Meeting. Washington, DC. Status = OTHER; Acknowledgement of Federal Support = Yes

Tristan J. Shepherd, Rebecca J. Barthelmie & Sara C. Pryor (2019). *Assessment of wind turbine impact on future climate in GCM-driven WRF simulation*. North American Wind Energy Academy WindTech Conference. Amherst, MA. Status = OTHER; Acknowledgement of Federal Support = Yes

Terrence Hendershott, Dmitry Livdan & Dominik Roesch (2018). *Asset pricing: A tale of night and day*. American Finance Association 2019 Annual Meeting, Jackson Hole Finance Conference & Desmarais Global Finance Research Centre. Vienna, Jackson Hole, WY & Montreal. Status = PUBLISHED; Acknowledgement of Federal Support = Yes

Rich Wolski, Chandra Krintz & Wei-tsung Lin (2018). *CSPOT: A serverless platform of things*. University of California, Santa Barbara Computer Science Technical Report. Santa Barbara, CA. Status = PUBLISHED; Acknowledgement of Federal Support = Yes

Rich Wolski, David Lifka & Tom Furlani (2020). *Data analysis and management for multi-campus cyberinfrastructure through cloud federation*. 2020 NSF Cyberinfrastructure for Sustained Scientific Innovation (CSSI) PI Meeting. Seattle, WA. Status = OTHER; Acknowledgement of Federal Support = Yes

Matt Baughman, Simon Caton, Christian Haas, Ryan Chard, Rich Wolski, Ian T. Foster & Kyle Chard (2019). *Deconstructing the 2017 changes to AWS spot market pricing*. ScienceCloud@HPDC. Boulder, CO. Status = PUBLISHED; Acknowledgement of Federal Support = Yes

Eric Coulter, Jeremy Fisher, Stephen Bird, Sanjana Sudarshan, Peter Vaillancourt & Suresh Marru (2020). *Deep dive into constructing containers for scientific computing and gateways*. Practice & Experience in Advanced Research

Computing (PEARC'20). Virtual. Status = OTHER; Acknowledgement of Federal Support = Yes

Rich Wolski (2019). *Devices-as-Services and the Internet as a Platform of Things (Keynote)*. Arm Research Summit 2019. Austin, TX. Status = OTHER; Acknowledgement of Federal Support = Yes

Rich Wolski (2019). *Devices-as-services: Rethinking scalable service architectures for the Internet of Things*. 2019 USENIX Annual Technical Conference. Reston, VA. Status = PUBLISHED; Acknowledgement of Federal Support = Yes

Susan Mehringer (2018). *Federations: Lessons learned*. Coalition for Academic Scientific Computation Fall Meeting. Alexandria, VA. Status = OTHER; Acknowledgement of Federal Support = Yes

Dominik Rosch (2020). *Financial market frictions*. Cornell Johnson School. Ithaca, NY. Status = OTHER; Acknowledgement of Federal Support = Yes

Richard Knepper, Resa Reynolds, Eric Coulter & Steve Bird (2019). *Gathering requirements for a campus cloud toolkit*. Practice and Experience in Advanced Research Computing (PEARC19). Chicago, IL. Status = OTHER; Acknowledgement of Federal Support = Yes

Shami Chatterjee (2018). *New results in radio astronomy: Fast radio bursts and transients*. International Astronomical Union General Assembly. Vienna. Status = OTHER; Acknowledgement of Federal Support = Yes

Tristan J. Shepherd, Rebecca J. Barthelmie & Sara C. Pryor (2019). *Quantifying array-array wind turbine impact on future climate in GCM driven WRF simulations*. North American Wind Energy Academy WindTech Conference. Amherst, MA. Status = OTHER; Acknowledgement of Federal Support = Yes

Tristan Shepherd, Bennett Wineholt, Rebecca J. Barthelmie & Sara C. Pryor (2019). *Quantifying weather and climate simulation reproducibility in the cloud*. 99th American Meteorological Society Annual Meeting: 5th Symposium on High Performance Computing for Weather, Water, and Climate. Phoenix, AZ. Status = OTHER; Acknowledgement of Federal Support = Yes

Richard Knepper (2019). *Red Cloud and Aristotle: Campus clouds and federation*. Humanware Advancing Research in the Cloud (HARC) Workshop. Chicago, IL. Status = PUBLISHED; Acknowledgement of Federal Support = Yes

Resa Reynolds (2019). *Red Cloud and the Aristotle Cloud Federation*. Dell XL HPC Consortium Fall Meeting. Scottsdale, AZ. Status = OTHER; Acknowledgement of Federal Support = Yes

Rich Wolski, Chandra Krintz, Fatih Bakir, Wei-tsung Lin & Gareth George (2019). *Rethinking scalable services for the Internet of Things*. Big Data and Extreme-Scale Computing Workshop (BDEC2). San Diego, CA. Status = OTHER; Acknowledgement of Federal Support = Yes

Chandra Krintz (2018). *SmartFarm: IoT systems that simplify and automate agriculture analytics*. 8th International Conference on Internet of Things (IoT 2018). Santa Barbara, CA. Status = PUBLISHED; Acknowledgement of Federal Support = Yes

Chandra Krintz (2020). *SmartFarm: computing research for the next-generation of precision agriculture*. American Association for the Advancement of Science Annual Meeting. Seattle, WA. Status = OTHER; Acknowledgement of Federal Support = Yes

Richard Knepper (2019). *Sustainability Research Impacts*. Cornell University Sustainability Leadership Summit. Ithaca, NY. Status = OTHER; Acknowledgement of Federal Support = Yes

Shami Chatterjee (2018). *The dwarf galaxy host of a repeating fast radio burst*. International Astronomical Union General Assembly. Vienna. Status = OTHER; Acknowledgement of Federal Support = Yes

Rebecca J. Barthelmie, Sara C. Pryor & Tristan J. Shepherd (2018). *Wakes from wind turbine array*. National Renewable Energy Laboratory. Golden, CO. Status = OTHER; Acknowledgement of Federal Support = Yes

Sara C. Pryor, Tristan J. Shepherd, Melissa Bukovsky & Rebecca J. Barthelmie (2019). *Wind energy scenarios for climate change mitigation*. The Scenarios Forum 2019. Denver, CO. Status = OTHER; Acknowledgement of Federal Support = Yes

Other Products
Audio or Video Products.

Rich Wolski (2019). *Devices-as-Services and the internet as a Platform of Things.* Keynote presentation at the Arm Research Summit 2019. See the video at https://www.youtube.com/watch?v=97FbH4Kob_o&feature=youtu.be

Other Publications

Patents

Technologies or Techniques

The Aristotle project team and science use case researchers developed the following technologies and techniques:

Automated Deployment Methods – developed a multi-cloud automated deployment method using Terraform and Ansible for deploying multi-VM MPI-capable clusters with Docker. Also, developed an automated deployment method for Kubernetes using Terraform and created MPI-capable clusters.
Centaurus – created a cloud service for K-means clustering. Centaurus is currently being used by former REU student Nevena Golubovic (who went on to complete her PhD) to correlate California power grid usage data with water usage data. Golubovic is with a start-up that received a grant from the California Energy Commission. Two Aristotle REUs to date have gone on to complete their PhDs and a third has been accepted into a PhD program. (https://sites.cs.ucsb.edu/~ckrintz/papers/centaurus-journal18.pdf).
Convertible Application Containers – developed application containers that are convertible from Docker to Singularity and verified functionality on a variety of cloud platforms.
CSPOT – developed a portable, multi-scale Functions-as-a-Service (FaaS) system for implementing IoT applications (https://www.cs.ucsb.edu/sites/cs.ucsb.edu/files/docs/reports/master_2.pdf).
Federated Open XDMoD with Cloud Metrics – added cloud metrics to Open XDMoD 9.0: average cores reserved, average memory reserved, average root volume storage reserved, average wall hours per session, total core hours, number of active sessions, number of sessions ended, and number of sessions started. These metrics can be grouped or filtered by instance type, project resource, and VM size (core/memory). See XDMoD cloud metrics on Cornell's Red Cloud (https://openxdmod.cac.cornell.edu /#tg_usage?node=statistic_Cloud_none_cloud_avg_cores_reserved).
hCOBRA – developed a modeling framework that extends the COBRA Toolbox for MATLAB to make large scale simulations more scalable and reliable so that scientists can better navigate the complexity of metabolic modeling and the recursive structure of simulations with priority effects.
Radio Astronomy Container - completed a single container of radio astronomy software that combines the pipeline components developed for Pulsar and other transient detections that can be deployed either on the cloud with Docker or on an XSEDE HPC resource with Singularity.
Semi Dynamic SteadyCom – developed a method that employs SteadyCom with discrete time steps.
Telemetry Data Visualizer – REU student Kerem Celik created a tool to visualize telemetry data from Citrus Under Protective Screening (CUPS) and the Edible Campus farm at UCSB. The visualizer can be downloaded and run on a MAC or it can be run as a Docker-based software service. A paper may be forthcoming. UCSB is also considering using this tool for a big history study on climate change.
WaterPaths Container – developed an on-demand MPI cluster with Docker container-based software deployment.

Thesis/Dissertations

Websites
Aristotle Cloud Federation
https://federatedcloud.org

The Aristotle Cloud Federation portal was updated regularly to feature new web content and user guide documentation as well as Aristotle science use case accomplishments, plans, and products. Quarterly National Science Foundation project reports are also available on the portal; they provide detailed descriptions of the project's activities, challenges, and accomplishments.

Participants/Organizations

Research Experience for Undergraduates (REU) funding

Form of REU funding support:	REU supplement
How many REU applications were received during this reporting period?	35
How many REU applicants were selected and agreed to participate during this reporting period?	5
REU Comments:	

What individuals have worked on the project?

Name	Most Senior Project Role	Nearest Person Month Worked
Lifka, David	PD/PI	1
Furlani, Thomas	Co PD/PI	1
Wolski, Richard	Co PD/PI	3

Full details of individuals who have worked on the project:

David A Lifka
Email: lifka@cornell.edu
Most Senior Project Role: PD/PI
Nearest Person Month Worked: 1

Contribution to the Project: Programmatic oversight of the Aristotle Cloud Federation project ensuring deliverables outlined in the program execution plan are met on schedule.

Funding Support: No funding support from other projects used for this award.

International Collaboration: No
International Travel: No

Thomas R Furlani
Email: thomas.furlani@roswellpark.org
Most Senior Project Role: Co PD/PI
Nearest Person Month Worked: 1

Contribution to the Project: Programmatic oversight of the UB subaward, and continuous interaction with the entire Aristotle technical team.

Funding Support: No funding support from other projects used for this award.

International Collaboration: No
International Travel: No

Richard Wolski
Email: rich@cs.ucsb.edu
Most Senior Project Role: Co PD/PI
Nearest Person Month Worked: 3

Contribution to the Project: Architected deployments at UCSB for production and Aristotle services. On-boarded new science use cases.

Funding Support: No funding support from other projects used for this award.

International Collaboration: No
International Travel: No

What other organizations have been involved as partners?

Name	Type of Partner Organization	Location
Dartmouth	Academic Institution	Hanover, NH

Full details of organizations that have been involved as partners:

Dartmouth

Organization Type: Academic Institution
Organization Location: Hanover, NH

Partner's Contribution to the Project:
In-Kind Support
Collaborative Research

More Detail on Partner and Contribution: Dartmouth is collaborating with the Aristotle team to test our Federated Cloud resources and documentation

---

What other collaborators or contacts have been involved?
Nothing to report

Impacts

What is the impact on the development of the principal discipline(s) of the project?

The Aristotle project is advancing the knowledge of federated cloud computing and its potential role as a campus bridging paradigm. By building and deploying a federated cloud model with the necessary allocations, accounting, and cloud metrics, Cornell University, University at Buffalo, and UC Santa Barbara are exploring how cloud resources can be effectively shared between campuses and their impact on researchers who require flexible workflows and analysis tools for large-scale data sets. The project serves as an important model for campus cyberinfrastructure that others may follow and provides elasticity by sharing resources, data, software, and tools that may not be available locally.

What is the impact on other disciplines?

Aristotle use case scientists are strategically exploring problems of increasing complexity and corresponding increases in data and, as a result, are advancing scientific knowledge. Data challenges from a diversity of disciplines (earth and atmospheric sciences, finance, chemistry, astronomy, civil engineering, genomics, and agriculture) are being addressed with collaborators from other academic institutions, public agencies, and research labs, as well as citizen scientists. The sharing of data infrastructure building blocks capacity and the movement of instances across institutional boundaries has the potential to create wider science collaborations and increased data sharing. The creation and performance testing of containerized applications can provide portability to from local cloud resources to other campus clouds, NSF clouds, or public clouds, thereby, producing a positive impact on research productivity in terms of agility and cost-effectiveness.

What is the impact on the development of human resources?

Aristotle is pioneering the concept of federated cloud computing for research which may ultimately increase the availability of on demand resources, data, and analysis tools that engage underrepresented groups. In addition, virtual laboratories in the cloud can enhance classroom learning. For example, Aristotle cloud resources were used to solve video analysis application problems in a Computer Science class (CS293B) taught at UC Santa Barbara in spring 2020 and were used to host JupyterHub with the nbgrader module in a Programming and Database Fundamental for Data Scientists class (EAS503) at the University at Buffalo in the fall and spring. Aristotle is cost-effective because the professor can choose the instance type based on the expected class size, and after the semester is over, delete the instance.

The availability of campus-to-campus cloud computing may also spurn the development and dissemination of ready-to-launch VMs and containers with training software and tools preloaded. This could reduce the redundant development and preparation of educational material development and onsite computer labs administration, resulting in an increased focus on individual student learning needs.

Nineteen REU students have participated in the Aristotle project to date. This year's REU student experiences were highlighted in *HPCwire*: https://www.hpcwire.com/off-the-wire/cornell-students-immersed-in-latest-cloud-technologies-thanks-to-nsf-research-experiences-for-undergraduates-program/. Two former Aristotle REUs have gone on to complete their PhDs, and a third just entered a PhD program.

What is the impact on physical resources that form infrastructure?

The federated cloud model may impact the physical resources that form infrastructure by reducing the number of computer labs required for learning. Campus clouds may be installed so that researchers can cost effectively use local cloud resources and, when more capacity is needed, move their application container to the most suitable campus, public, or NSF cloud resource. Ultimately, federated clouds will likely become complementary resources to high-end supercomputers, e.g., performing on demand iterative tasks, streaming IoT data, etc. Integrating federated clouds (campus, labs, etc.) into the national cyberinfrastructure ecosystem may increase the sharing of resources and, subsequently, reduce physical resource expenditures at the local level.

What is the impact on institutional resources that form infrastructure?

Aristotle will maximize institutional resources through federation with other institutions by (1) offloading variable computational and data analysis demands from local infrastructure, (2) starting coarsely parallel computations on demand, (3) providing heterogeneous instance types (CPUs, GPUs, etc.) and sizes to allow for unpredictable computational demand.

What is the impact on information resources that form infrastructure?

Aristotle's federated cloud model will facilitate (1) sharing of high-value processed datasets of general interest and separate data resources, (2) generation of reproducible pipelines in the form of VMs or VM configurations, and containers (3) access to multiple data sources, some of which are already in public and private clouds.

What is the impact on technology transfer?

Technologies generated by this project are open source, therefore technology transfer licensing, patent applications, etc. are not applicable. The project's collaborative relationships for R&D include Amazon Web Services, Google Cloud, Microsoft Azure, Red Hat, Globus, Dell. and HPE.

What is the impact on society beyond science and technology?

Aristotle use cases have the potential to impact wind turbine companies interested in harnessing the energy of the atmosphere and converting it into carbon-free electricity; policymakers regulating high-frequency trading; policymakers making water resources management decisions; manufacturers producing sustainable insect pest management products; farmers increasing yields and protecting the environment by accessing on demand soil, water, and crop sensor data to aid decision-making; and, orange growers trying to protect U.S. citrus trees from the Huanglongbing bacteria (citrus greening disease) that has devasted FL and is threatening CA.

Changes/Problems

Changes in approach and reason for change
Nothing to report.

Actual or Anticipated problems or delays and actions or plans to resolve them
Nothing to report.

Changes that have a significant impact on expenditures
Nothing to report.

Significant changes in use or care of human subjects
Nothing to report.

Significant changes in use or care of vertebrate animals
Nothing to report.

Significant changes in use or care of biohazards
Nothing to report.