Preview of Award 1541215 - Annual Project Report

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 - 09/30/2020
Reporting Period:	10/01/2015 - 09/30/2016
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?**

- 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 (CU), 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 food science) 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 that features QBETS statistics to make online forecasts of future performance and allocations levels available to users.

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

Major Activities:

Order the initial set of storage assets for each site.

Specific Objectives:

- Performed hardware/software requirements analysis & gathered quotes *(fall 2015*).
- Installed hardware at CU; received hardware at UB and UCSB *(winter 2016)*.

Develop portal code allowing InCommon authentication by users to access their Eucalyptus keys for the federated cloud resources at each site.

- Implemented InCommon.

Develop Phase 1 portal content.

- Defined portal software requirements to meet authentication and Open XDMoD functionality *(fall 2015)*.
- Built portal on CU Red Cloud instance; installed SSL certificate; added InCommon authentication (winter 2016).
- Launched federatedcloud.org on 6/27/16 with systems info, etc.

Install, configure and test storage assets, networks, allocations and accounting mechanisms across federated sites. Monitor Eucalyptus product support & if necessary plan for switch to OpenStack.

- Stress tested hardware at CU; installed/stress tested hardware at UB; began installation at UCSB. Evaluated CEPH performance *(spring 2016)*.
- All storage assets in production *(summer 2016):* 168 cores added to CU Red Cloud for total of 376 cores and 226TB SAN; 144 cores added to UB Lake Effect for total of 312 cores and 336TB SAN; 140 UCSB Aristotle cores installed with 228TB CEPH. The largest instance type is 28 cores/256GB RAM.
- Leveraged CU Red Cloud allocations and accounting software while developing redistributable building blocks. Plan to use Globus Authentication instead of InCommon for a better interface to XSEDE and the Euca console *(winter 2017)*.
- Developed a schema (see pdf) for the project, user, accounting, and usage database. Created a temporary database and built stored procedures and API calls so users could begin setting up instances. Data from UB and UCSB will be ingested into the main database by use of a REST API developed by UB. CU is testing this API for use in the federation.
- Negotiated with HPE for multi-year Eucalyptus support and SAN drivers.

Implement Open XDMoD data collection across the federation.

- Developed a framework for Federated XDMoD in which data from independent XDMoD installations is

collected and displayed by a central federated XDMoD instance. Application is to general HPC use as well as Aristotle, with Aristotle being the first installation. See Federated XDMoD Requirements: 1.0 Draft (pdf).
- Redesigned the XDMoD data warehouse and ETL infrastructure to store and report on non-traditional and innovative cyberinfrastructure (cloud, Hadoop, etc.). See XDMoD Requirements Document: Job Reporting for Cloud and Other Non-Traditional HPC Resources (pdf).

Integrate QBETS Metrics and Predictions into Open XDMoD and associated data collection for portal.

- Wolski developed a stand-alone version of QBETS for batch queue systems in addition to the AWS spot-price prediction capabilities. UB ran simulations to explore the use of QBETS as a wait time prediction tool through comparison with historical wait time data from XSEDE. QBETs provided reliable wait time upper bounds indicating potential capability as a tool for XSEDE/Aristotle users to achieve optimal time to science by selecting the resource with the smallest wait times.
- QBETS will be integrated into Open XDMoD *(winter 2017).*

PIs agree upon initial allocation for strategic science use cases.

- PIs agreed upon early allocations and are using a CU Red Cloud node with fixed use-or-lose core hours available to all with a 10TB storage limit, pending full deployment of the accounting system.

Create help desk to track questions/gather information.

- Implemented Aristotle help desk (help@federatedcloud.org) using CU RT *(winter 2016).*

Support staff at each site will begin working with 7 science use case researchers and collaborators to develop plans to use federated cloud when it goes online. Prioritize support team activity. Allow early user testing.

1. **A Cloud-Based Framework for Visualization & Analysis of Big Geo Data** (UB Chandola). ***Requirements*:** 16 node cluster (8 cores/node, 28GB), 25,000 CPU hours, 1TB storage. ***Work to date:*** Developed the capabilities to create Spark analytical clusters on-demand. Deployed a distributed machine learning method for change detection in sustainability data and a browser-based, cloud-driven interactive 3D globe based environment for analysis. Ran scalability tests on Red Cloud at CU before migrating the application to Lake Effect at UB.
2. **Global Market Efficiency Impact** (UB Roesch, Tiu & Wolfe). ***Requirements***: 1 node cluster (8-16 cores/node, 64-128GB, up to 40,000 CPU hours, 2-10TB storage. ***Work to date****:* Developed the capability to migrate OpenNebula images to the cloud. Deployed an image on Lake Effect and installed the Thomson Reuters Tick History (TRTH) database, a record of the world's major asset classes and markets at miniscule time intervals. Conducted preliminary research which will be scaled up to understand market efficiency at a global level.
3. **High Fidelity Modeling and Analytics for Improved Understanding of Climate-Relevant Aerosol Properties** (CU Pryor). ***Requirements****:* 4GB+ per core with distributed memory; 10TB base data and 25TB data output for each year processed. ***Work to date*:** Brazier met with the research team to define requirements; Lee built/documented a parallel version of WRF-Chem; and single-process and distributed

memory instances were migrated for initial testing/development as a cluster of VMs.

4. **Transient Detection in Radio Astronomy Search Data** (CU Cordes). ***Requirements:*** 4GB per core; 30TB data growing at 5TB/year; processing takes 1.5-2 core hours per beam, ~210,000 beams, ~500,000 core hours per complete reprocessing. ***Work to date:*** Brazier/Chatterjee defined the architecture. A VM with basic software was created and user access configured. PRESTO pulsar search/analysis software (used for transients detection) was installed. Data decimation planning (factor of 4 in time resolution/2 in frequency) and modification of file handling began for data deresolution. Work continues building a pipeline to process the data.
5. **Water Resource Management Using OpenMORDM** (CU Reed). ***Requirements:*** 32 compute nodes (512 cores, 8GB-128GB/core); 8TB storage, 900GB local, 95TB Lustre/scratch space. ***Work to date****:* Trindale is building an image for OpenMORDM so the analytic model can be tested in the cloud. Plan to get the software stack up next quarter and develop a usage guide. CU continues to work on CloudLaunch to launch/tear down clusters.
6. **Mapping Transcriptome Data to Metabolic Models of Gut Microbiota** (CU Douglas). ***Requirements****:* ~200GB storage years 1-3; 8TB years 4-5 for metagenome data analysis; 64GB per node, 4+cores/node; 8,500+ core hours. ***Work to date****:* Barker investigated NixOS (cloud-focused OS) as a way to develop reproducible scientific workflows, but will use Ubuntu and Windows for now. He also investigated Supercloud as a possible means to scale out to other clouds. A Windows instance was created for MATLAB and a Samba server was set up to allow file sharing across Windows and Linux. A transcriptomic pipeline is being developed using Docker. This work is documented in the wiki, including scalable file access for cross-cloud and many-instance file/app access.
7. **Multi-Sourced Data Analytics to Improve Food Production** (Sedgwick Reserve McCurdy/UCSB). ***Requirements:*** 8-32 nodes (4-8 cores/node, 4-8GB/node), ~96 core hours/per week/per node. ***Work to date:*** Aristotle back end processing is functional for collecting moisture sensor data to improve seedling development of CA Live Oaks; testing continues. A grad student and undergraduate student developed an animal-identification system that runs Google's TensorFlow in parallel on Aristotle to auto-identify animals from camera-trap images (200,000 photos/month).

Contact help@federatedcloud.org to request access to detailed online monthly and quarterly Aristotle NSF reports. Science Team Advisory Committee meeting minutes are available at https://federatedcloud.org/science/advisorycommittee.php.

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 effective cross-training and knowledge sharing, and in

depth follow-up calls occurred to explore key technologies, both available and emerging, and to share lessons learned.

Undergraduate & Graduate Student Development

At UCSB, an undergraduate student from the EUREKA! Program (an UCSB program to encourage underrepresented minorities to explore scientific research) has been working with the Aristotle staff to learn cloud computing. He is part of the science team developing the animal identification application.

At UB, a Computer Science and Engineering PhD student has been supported by the grant and is working on developing resources to allow the execution of UB science use cases in the cloud. The student also developed a novel distributed machine learning algorithm for spatio-temporal change detection which is key for the sustainability analysis.

How-To Documentation & Training

Training for CU local researchers and a webinar for remote researchers will be delivered in fall 2016 or winter 2017. Documentation will continue be made available as it is created and tested. Documentation currently available includes Eucalyptus Basics, Euca2ools, Image Management, Shared File Systems, and How to Create a Spark Cluster.

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

Conferences & Presentations

- Coalition for Academic Scientific Computation (CASC), Arlington, VA: *Aristotle Project Overview* – David Lifka (10/14-10/16/2015)
- Aristotle Science Team Advisory Committee: *Spring '16 Meeting* – Adam Brazier and Aristotle team; minutes available at portal (4/1/2016).
- National Science Foundation Invited Presentation, Arlington, VA: *Building a Federated Cloud Model: Aristotle Cloud Federation* – David Lifka (4/1/2016)
- Best Practices in Data Infrastructure Workshop, Pittsburgh, PA: *Building a Federated Cloud Model* – Tom Furlani (5/17-5/18/2016)
- XSEDE '16 BOF, Miami, FL: *Current Practices and Future Opportunities in Processing XSEDE Allocations Awards and Usage Data* – Tom Furlani (7/18/2016)
- XSEDE '16 BOF, Miami, FL: *Topics in Cloud and Virtualization in XSEDE* – David Lifka and colleagues (7/20/2016)
- Aristotle Science Team Advisory Committee*: Summer '16 Meeting* – Adam Brazier and Aristotle team; minutes available at portal (7/27/2016)
- MITRE, McLean, VA: *Advances in Virtual Globe Technology Using NASA World Wind, an Open Source Geobrowser* – Varun Chandola (8/10-8/11/2016).

News Coverage

ACM TechNews, CIO, Cloud Strategy, Cloudwards, Cloud Strategy, Data Center Talk, Government Computing News, HPCwire, InsideBIGDATA, Network World, NextGen Communications, Primeur, Scientific Computing, Scientific Computing World, and *Web Host Industry Review* were among the news outlets that covered the Aristotle Cloud Federation. Stories included *Will the cloud change scientific computing?* (retweeted to 40+ outlets with > 25,000 followers).

Industry Case Studies

Case studies included "Cornell's Red Cloud: A Model for the Future of Research" and "Red Cloud: A Hybrid Cloud Model for Academia."

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

Infrastructure & Portal Plans

- Order/install/configure second year storage assets at each site.
- Test network performance; install 10G switch at UB.
- Implement OAuth 2.0 support for single sign-in credentials (OAuth 2.0 availability is scheduled for ~Dec. 2016 Eucalyptus 4.4 release); Aristotle motivated HPE to add OAuth 2.0 support to Eucalyptus).
- Continue to share resources and transition from a local to a federated accounting and allocations system currently under development.
- Update user documentation to be federation-specific as well as to reflect new software and tools, e.g., plan to add and document single node (up to 28 core instance) and multi-node MATLAB Distributed Computing Server 2016A capability at CU.
- Develop how-to guides for sites who wish to build similar systems.
- Enhance use case documentation, explaining benefits/challenges and new techniques and best practices.
- Hold training for local researchers, webinar for remote researchers, focusing on more advanced training topics.

Metrics & Usage Plans

- Implement a working version of Federated Open XDMoD
- Implement cloud metrics in XDMoD which will require re-engineering the XDMoD data warehouse (re-engineering is underway).
- Refactor the DrAFTS and QBETS analytics web services.

Science Use Case Plans

1. **A Cloud-Based Framework for Visualization & Analysis of Big Geo Data** (UB Chandola). Make a webGlobe-based user interface available to the research community to conduct sustainability-based research in the cloud. Create a science gateway to allocate resources to the scientific community to use this tool.
2. **Global Market Efficiency Impact** (UB Roesch, Tiu & Wolfe). Conduct a full scale analysis of global market efficiency using Thompson Reuters Tick History data.
3. **High Fidelity Modeling and Analytics for Improved Understanding of Climate-Relevant Aerosol Properties** (CU Pryor). Performance and scaling testing on WRF-Chem installation. Run a limited model to better understand atmospheric particle concentrations, one of the largest uncertainties in understanding climate change
4. **Transient Detection in Radio Astronomy Search Data** (CU Cordes). Scale PALFA data down to a manageable size. Design and build the basic pipeline to process specified data to search for single pulses.
5. **Water Resource Management Using OpenMORDM** (CU Reed). Run the water resource management software stack and continue to investigate whether many containers can be spun up across multiple clouds, starting with AWS. The ultimate goal is to potentially develop an AWS service for US municipalities.
6. **Mapping Transcriptome Data to Metabolic Models of Gut Microbiota** (CU Douglas). Complete the initial transcriptomic analysis pipeline with batch parallelization handled by the newly developed Docker container runner. Model newly created symbiotic models using custom modifications to the OptCom algorithm on a scalable Windows instance.
7. **Multi-Sourced Data Analytics to Improve Food Production** (Sedgwick Reserve McCurdy/UCSB). Fully implement an automated irrigation monitoring system for a grape block using Aristotle to host and analyze soil moisture sensor data. Initial tests indicated over watering by a factor of nearly 3x. Use multiple Aristotle instances to try to scale the TensorFlow machine learning application for animal identification. Host processed images in Aristotle storage.

Emerging Technology Plans

CloudLaunch: CU plans to continue to develop CloudLaunch, a technology that consists of extensions to a SLURM scheduler which would allow researchers to use cloud resources with traditional HPC methods. The goal is for each site to offer a "login node" where researchers can log in, compile codes, check output, and submit, cancel and check job status. When a job is submitted on this login node, the requested number of virtual "node instances" would be spun up and the job run on them as if it were a dedicated HPC cluster (possibly a cluster of containers). When the job completes, SLURM would shut the virtual node instances down making the cloud resources available for other types of work.

DrAFTS (Durability Agreements from Time Series): UCSB and the Aristotle team plan to continue DrAFTS testing and optimization. DrAFTS was developed by Wolski under the Aristotle project. It is a system for predicting the "bid price" that an AWS user should bid in the spot market to ensure a minimum duration of execution before AWS terminates the instance. A user who chooses the bid price gets a statistical guarantee, i.e., an SLA on the lifetime before AWS will terminate the instance. DrAFTS utilizes QBETS internally. It has been tested on individual job launches and it verified when a user wants to run one job. Each of these experiments was a week long and launched 100 jobs (one at a time, in sequence) so, statistically it is operational for individual users with occasional needs. At the request of the Globus project, UCSB collaborated with Ian Foster's team and integrated DrAFTS with Globus Genomics. Results to date are promising.

Supercloud: The Aristotle team, in conjunction with Cornell Computer Science developers, plans to continue to investigate the capabilities of Supercloud which enables users to migrate VMs across different platforms. There is interest in using Supercloud's Xen-Blanket wrappers to allow images created under Eucalyptus to run on OpenStack clouds, Azure, and other platforms. The Aristotle team contacted Indiana University and arranged a proof of concept. A VM was successfully migrated from Jetstream to Red Cloud and back using Supercloud.

Plans to Disseminate Results

The Aristotle team will continue to engage the HPC community through presentations and dialogue through Coalition for Academic Scientific Computing (CASC) meetings and conferences such as SC16, XSEDE '17, International Workshop on Analytics for Big Geospatial Data, etc.

CU plans to submit a proposal to lead the first annual NSF Data Infrastructure Building Blocks PI Workshop (target January 2017) for all DIBBs PIs/Co-PIs.

Aristotle PI/Co-PIs will continue to respond to inquiries regarding the project, e.g., CU briefed Chris Hill, an MIT scientist and co-leader of the research, education and outreach committee of the Massachusetts Green High Performance Computing Center.

The Aristotle team will continue to seek guidance from and keep the Aristotle External Advisory Committee abreast of our activities and plans. Committee members are:

- Amy Walton, NSF
- Dmitrii Calzago, HPE
- Ian Foster, ANL/University of Chicago
- Steve Johnson, Weill Cornell Medicine
- Sanjay Padhi, AWS
- Ben Rosen, Dell
- Craig Stewart, Indiana University
- John Towns, XSEDE/NCSA
- Rick Wagner, SDSC.

Supporting Files

Filename	Description	Uploaded By	Uploaded On
AllocationsAccountingDatabaseSchema.pdf	Allocations Accounting Database Schema in response to "What was accomplished under these goals? Specific Objectives:"	David Lifka	09/15/2016
FederatedXDMoDReqsDraft1.0.pdf	Federated XDMoD Requirements in response to "What was accomplished under these goals? Specific Objectives:"	David Lifka	09/15/2016
XDMoDReqsDocJobReporting.pdf	XDMoD Requirements Document for Job Reporting in response to "What was accomplished under these goals? Specific Objectives:"	David Lifka	09/15/2016

Products

Books

Book Chapters

Inventions

Journals or Juried Conference Papers

Wolski, R. and Brevik, J. (2016). Providing Statistical Reliability Guarantees in the AWS Spot Tier.. *24th High Performance Computing Symposium (HPC 2016)*. 13. Status = PUBLISHED; Acknowledgment of Federal Support = Yes ; Peer Reviewed = Yes ; OTHER: Society for Computer Simulation International, 2016

Wolski, R., Brevik, J., Chard, R., and Chard K. (2017). Probabilistic Guarantees of Execution Duration for Amazon Spot Instances.. *IEEE International Conference on Cloud Engineering. (IC2E 2017)*. . Status = UNDER_REVIEW; Acknowledgment of Federal Support = Yes ; Peer Reviewed = Yes ; OTHER:

Licenses

Other Conference Presentations / Papers

Yang, Z., Du, A.Y., Das, S., Ramesh, R., Furlani, T., von Laszewski, G., and Qiao, C. (2016). *Providing Statistical Reliability Guarantee for Cloud Clusters.*. Global Communications Conference (IEEE GLOBECOM '16. Washington, DC.. Status = PUBLISHED; Acknowledgement of Federal Support = Yes

Other Products

Case Study.

Case Study: Cornell's Red Cloud: A Model for the Future of Research (April 2016)

https://www.hpe.com/h20195/v2/GetPDF.aspx/4AA6-5086ENW.pdf

Case Study.

CloudLaunch case study

https://www.cac.cornell.edu/technologies/CloudLaunch.pdf

Requirements Document.

Federated XDMoD Requirements: Draft 1.0 (April 2016)

Requirements Document.

XDMoD Requirements Document – Job Reporting for Cloud and Other Non-Traditional HPC Resources: Draft 1.0 (May 2016).

Other Publications

Patents

Technologies or Techniques

The CU science support team is developing an API that can be used to run parallel, distributed programs on diverse clouds that support Docker Swarm. The API will be flexible, allowing both interactive and dynamic workflows, but will initially focus on allowing users to run embarrassingly parallel jobs and distributed MPI programs. Core features, including data collection, are expected to be completed in fall 2016 or winter 2017 and will be documented. This capability will be used by several strategic science use cases.

Thesis/Dissertations

Websites
Aristotle Cloud Federation
https://federatedcloud.org

Project Website

Supporting Files

Filename	Description	Uploaded By	Uploaded On

Filename	Description	Uploaded By	Uploaded On
ProbabilisticGuaranteesofExecutionDuration.pdf	Copy of Paper "Probabilistic Guarantees of Execution Duration"	David Lifka	09/15/2016
ProvidingStatisticalReliabilityAWSSpotTier.pdf	Copy of Paper "Providing Statistical Reliability Guarantees in the AWS Spot Tier"	David Lifka	09/15/2016
ProvidingStatisticalReliabilityCloudClusters.pdf	Copy of Paper "Providing Statistical Reliability Guarantee for Cloud Clusters"	David Lifka	09/15/2016

Participants/Organizations

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@cac.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

What other organizations have been involved as partners?
Nothing to report.

What other collaborators or contacts have been involved?

Aristotle science use case researchers have collaborators at US and international universities as well as government agencies (NASA, NCAR, etc.) and several US cities.

The Aristotle project team collaborated with Eucalyptus engineers at HPE and the Scientific Computing (SciCo) group at AWS.

Impacts

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

The Aristotle project has the potential to have a transformative effect on the way computational resources are made available to support scientific research by allowing universities to easily share computational resources such that the computational and data resources available to a researcher is greater than that available through his or her campus alone.

In the area of scientific visualization, cloud-enabled systems such as webGlobe present a novel way of allowing users to interact with scientific data through their browsers.

What is the impact on other disciplines?

The Aristotle project could serve a model for the development of future cyberinfrastructure that provides scientists with a robust, delocalized cyberinfrastructure to conduct research.

What is the impact on the development of human resources?
Nothing to report.

What is the impact on physical resources that form infrastructure?
Nothing to report.

What is the impact on institutional resources that form infrastructure?

Aristotle will maximize institutional HPC resources through federation with other institutions by:

- Offloading variable computational demands from local infrastructure
- Starting coarsely parallel computations on demand
- Bursting to process new data
- Providing heterogeneous instance types and sizes to allow for unpredictable computational demands.

What is the impact on information resources that form infrastructure?

The Aristotle model will facilitate:

- Sharing of high-value processed datasets of general interest and disparate data resources
- Generation of reproducible pipelines in the form of VMs or VM configurations
- Access to multiple data sources, many of which are already in public and private clouds.

What is the impact on technology transfer?
Nothing to report.

What is the impact on society beyond science and technology?

Aristotle may prove to be a model that is widely adopted by academic institutions to support research and collaboration in many fields, including the social sciences and the arts.

Aristotle use cases have potential societal benefits such as a better understanding how humans and ecosystems are impacted by aerosols in the atmosphere; how US water supply systems and river basics can be better managed and protected; how microbial therapies can support health and resolve clinical conditions; and how sensor data can increase food protection, enhance the growth of at-risk species, and reduce water usage in drought prone areas.

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.