

*CC\*DNI DIBBs: Data Analysis and Management Building Blocks for Multi-Campus Cyberinfrastructure through Cloud Federation (NSF Award #1541215). PI: David Lifka, Cornell University; co-PIs: Thomas Furlani, University at Buffalo; Rich Wolski, University of California, Santa Barbara. DIBBs18 white paper.*

## **Challenges: DIBBs for multi-campus cyberinfrastructure through cloud federation**

The goal of the Aristotle Cloud Federation is to reduce time to science by building a cloud federation, which is managed/discovered through a portal that enables users to easily identify available resources in the federation and, additionally, NSF resources such as Jetstream and public cloud resources such as AWS. The portal (<https://federatedcloud.org>) will provide the information necessary for researchers to make an informed decision on the availability of resources combined with the “cost” in terms of allocation units or actual dollars in the case of public cloud providers. Seven strategic science use cases from intentionally diverse disciplines will demonstrate the potential of federated cloud as a campus bridging paradigm.

### **Challenge #1: Implement single sign-on for federated cloud resources.**

To make it as seamless as possible to move from one federated cloud resource to another, single sign-on is essential. OAuth 2.0 supports InCommon and other federated identities so it is ideal for ad hoc federations whose participant institutions may not support InCommon. OAuth 2.0 support is currently available at each site; next, we will integrate OAuth 2.0 with the new user dashboard which will provide tokenized access to all federation sites.

### **Challenge #2: Create, implement, and release a federated allocation and accounting system.**

In order to demonstrate the value of sharing institutional resources to their leaders, it is important to have an accounting system that can track a user’s usage on any federated resource. The ability to burst to other federated resources and trade allocation units before having to pay a public cloud provider allows institutions to only capitalize hardware their users can keep busy and to leverage partner institutions investments. We built a new federated accounting database encompassing resources, research teams, allocations, usage by site, and in the future, exchange rates across all sites in the federation. The dashboard in the user portal will enable signed-in project managers to manage their team members, view their allocation balance, usage across sites, access all federation sites with a single sign-on, see Open XDMoD metrics, and see availability across all federated cloud sites. Our next challenge is to refine and implement the allocation and accounting system and, when production ready, release it as an open source version for others to use to create their own federations. We plan to run the allocations and accounting system in AWS in order to ensure there is no single point of failure in accessing the system.

### **Challenge #3: Predict the availability of cloud resources.**

To identify opportunities for resource federation, we are developing predictive capabilities that allow users to decide whether and how to make use of resources across the federation. Using QBETS and DrAFTS (Durability Agreements from Time Series—a new prediction methodology), Aristotle users are able to predict the costs associated with using the AWS Spot Market for instances (at a considerable discount). Thus, a user who is bursting into a remote cloud in the federation will be able to predict how much doing so will cost for the burst of the duration when using spot instances. The challenge is for Aristotle to enable effective decision support for federation using new predictive capabilities made available to users.

### **Challenge #4: Track/ensure efficient use of federated resources with a federated version of XDMoD.**

Open XDMoD was initially designed to work as a comprehensive resource management system for individual HPC centers. However, in order for federations such as Aristotle to successfully share resources among institutions, it is necessary to have a global view of usage across all participating institutions. This necessitates creating a new federated version of Open XDMoD that collects usage and performance data from independently operating HPC centers and clouds and provides a global view of the usage across all institutions. Federated Open XDMoD will have applicability outside of the Aristotle project. Federated Open XDMoD is in beta testing and will be released to production in July 2018.