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

Award Number 1541215. PI: David Lifka, Cornell University; co-PIs: Tom Furlani, University at Buffalo; Rich Wolski, University of California, Santa Barbara

The goal of the Aristotle Cloud Federation is to optimize time to science by providing researchers with a portal that makes it easy to identify available resources in the federation and, additionally, NSF resources such as Jetstream and public cloud resources such as Amazon Web Services (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. We originally planned to use InCommon for this purpose; however, we realized that OAuth 2.0 (https://oauth.net/2/) was a better choice since it is used by much of the NSF community, particularly XSEDE users. OAuth 2.0 supports InCommon and other federated identities so it is ideal for ad hoc federations whose participant institutions may not support InCommon. We worked with our Aristotle partner HPE to add OAuth 2.0 support to the Eucalyptus console. We are currently testing an early release of this capability (12/16).

Challenge #2: Create 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 are currently building this system as part of Aristotle and will release it as an open source version for others to use to create their own federations. Our goal is 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 will be able to predict the availability of remote resources subject to their local control. Thus, a user who is bursting into a remote cloud in the federation will be able to predict when that remote cloud will need to reclaim its resources. In addition, DrAFTS allows users to predict the bid price that is necessary to guarantee a specific duration of execution in the AWS Spot Tier, thereby reducing the cost of bursting to Amazon by a factor of between 4 and 10. Thus, Aristotle will enable effective decision support for federation using new predictive capabilities made available to users.

Challenge #4: Track and insure the efficient use of federated resources with a new federated version of Open 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 use across all participating institutions. This necessitates creating a 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.