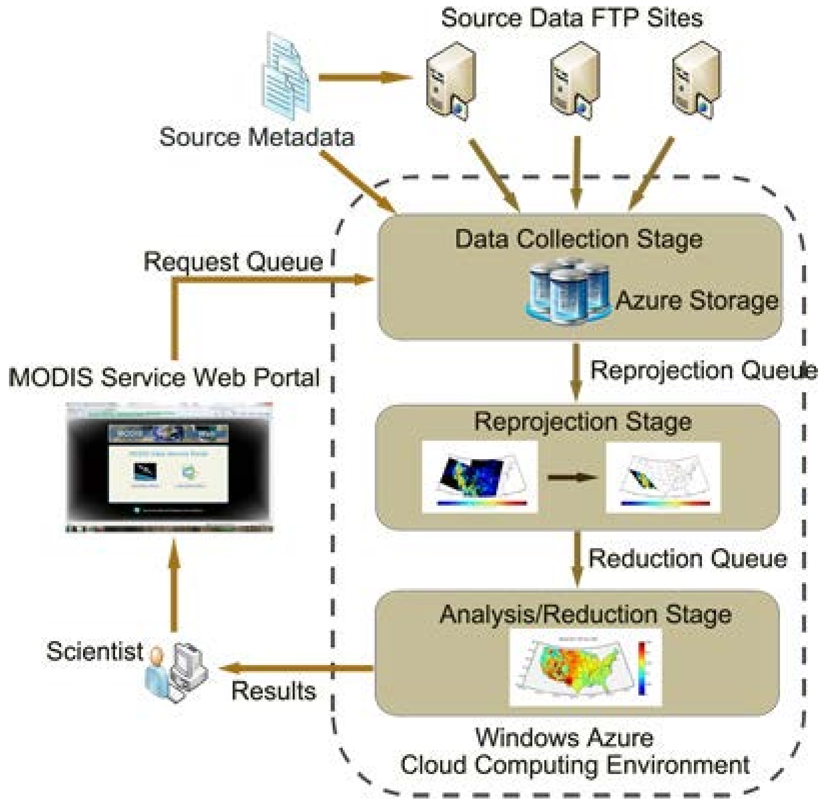


The focus of our group is to ensure predictable eScience based on public cloud resources such as Amazon Web Services (AWS). We are currently engaged in projects which seek to advance hydrologic science and water resource management by leveraging cloud computing for modeling large watershed systems. We are also creating a collaborative platform for environmental scientists through the FLUXNET, AmeriFlux and International Soil Carbon Networks.

# Cloud Computing Group


## Marty Humphrey

Associate Professor
humphrey@cs.virginia.edu
http://www.cs.virginia.edu/~humphrey

Department of Computer Science
University of Virginia
Charlottesville, VA
434.982.2258


University of Virginia
School of Engineering & Applied Science

## Cloud Computing

It is clear that the cloud is very successful and will continue to make a major impact on computational activities for aspect of society. However, the particular software architecture that has emerged for scalable and fault-tolerance cloud applications is significantly restricted: the work assumed necessary to fulfill client requests is simplistic and uniform, computational servers are homogeneous, and the financial aspects of cloud applications are treated independent of software requirements and implementations. Furthermore, while applications conforming to this pattern are scalable, the actual run-time scaling (and de-scaling) is only via direct human interaction or via rudimentary automated scaling mechanisms (such as adding a new server when the average load on existing server machines increases beyond a threshold value). The end result is that the cloud application can often fail to meet performance requirements related to complex time-dependent client requests. Our group is working to address the performance challenges faced by public clouds.

## Virtualization for eScience

The use of virtual machines for eScience has been advocated within the enterprise to replace aging machines and as the underlying technology of cloud computing. However, both scenarios can lead to inadequate performance. Within the enterprise, with incorrect planning or under unexpected heavy or even moderate load, there might not be enough physical capacity for every virtual machine to achieve reasonable performance. In cloud-computing-based scenarios, the "renters" are largely subject to the informal service promises of the cloud provider. We are pursuing a novel unified framework to ensure predictable eScience. The foundation of the approach is to wrap an eScience application in a performance container framework and dynamically regulate the application's performance through the application of formal feedback control theory. The application's progress is monitored and ensured such that the job meets its performance goals without requiring exclusive access to physical resources even in the presence of a wide class of unexpected disturbances.

## RECENT RESEARCH DEVELOPMENTS

- Paper, "Empirical Evaluation of Workload Forecasting Techniques for Predictive Cloud Resource Scaling". Proceedings of 9th IEEE International Conference on Cloud Computing (Cloud 2016).
- Paper, "Toward Optimal Resource Provisioning for Cloud MapReduce and Hybrid Cloud Applications". Proceedings of 8th IEEE International Conference on Cloud Computing (Cloud 2015).
- Working closely with LBL and USDA, collaborative platform published at http://iscn.fluxdata.org/

## RECENT GRANTS

- U.S. Dept. of Energy – UVA Support for AmeriFlux Network
- U.S. Dept. of Agriculture – UVA Participation in Coordination and Support of the National Soil Carbon Network
- NSF – Computing in the Cloud

**SEAS Research Information**
Pamela M. Norris,
Executive Associate Dean for Research
University of Virginia
Box 400232
Charlottesville, VA 22903
EngineeringResearch@virginia.edu
434.243.7683

University of Virginia | School of Engineering & Applied Science