



**Advanced  
Research  
Computing**



# Advanced Research Computing

April 19, 2016

Advanced  
Research  
Computing

# Advanced Research Computing supports:

- HPC and Visualization Systems
- Collaborative Research
- Improved Data Access
- Leveraging National Research and Education Networks
- Large Data Set Management
- Big Data Development and Analysis

# Vice President for Information Technology and Chief Information Officer

Dr. Scott Midkiff

## Associate Vice President for Research Computing

Dr. Terry L. Herdman

Vijay Agarwala .....	Director, High Performance Computing
Nicholas Polys .....	Director, Visualization
Gary M. Worley .....	Director, Archival Data Management
Alana Romanella .....	Assistant Director, Development and Fiscal Administration
Mark K. Gardner.....	Network Research Manager
Justin Krometis .....	Computational Scientist
James McClure .....	Computational Scientist
Brian Marshall .....	Computational Scientist
Srijith Rajamohan .....	Computational Scientist
Bob Settlage .....	Computational Scientist
Lance Arsenault .....	Visualization and Virtual Reality Specialist
Nathan Liles .....	Software Engineer
Brent Woodruff .....	Systems Engineer

# Vice President for Information Technology and Chief Information Officer

Dr. Scott Midkiff

## Executive Director for Network Infrastructure and Services

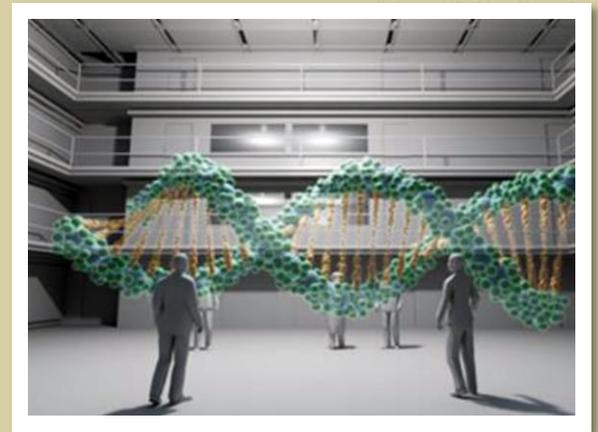
William Dougherty

Tim Rhodes.....	Technical Team Manager,	Systems Engineering & Administration
Christopher Snapp .....	HPC Systems Specialist,	Systems Engineering & Administration
Brandon Sawyers.....	HPC Systems Specialist,	Systems Engineering & Administration
Josh D. Akers.....	HPC Systems Specialist,	Systems Engineering & Administration
Jim Tibbs.....	Systems Engineer,	Systems Engineering & Administration
Wanda K Baber.....	Technical Team Manager,	Systems Engineering & Administration
Valdis Kletnieks.....	Systems Engineer,	Systems Engineering & Administration
J. Eric Wonderley.....	Systems Engineer,	Systems Engineering & Administration

# System – New River

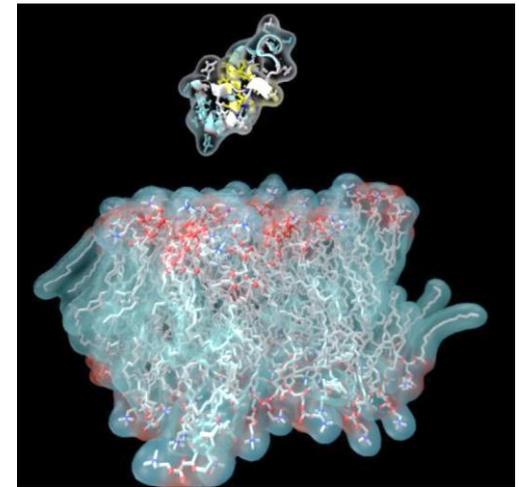
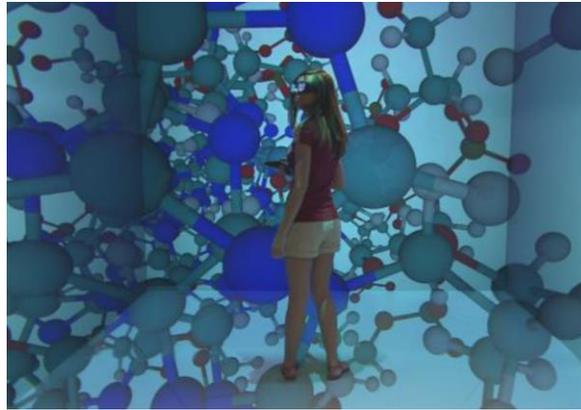
Our flagship HPC system

- 100-node 2400-core compute engine targeted at scalable and distributed workloads
- 16-node 384-core engine targeted at Big Data and Analytics
- 8-node 192-core compute engine targeted at code acceleration and data visualization
- 2-node 120-core compute engine targeted at very large data sets and graph analytics, with 3 TB memory per node



# System – BlueRidge

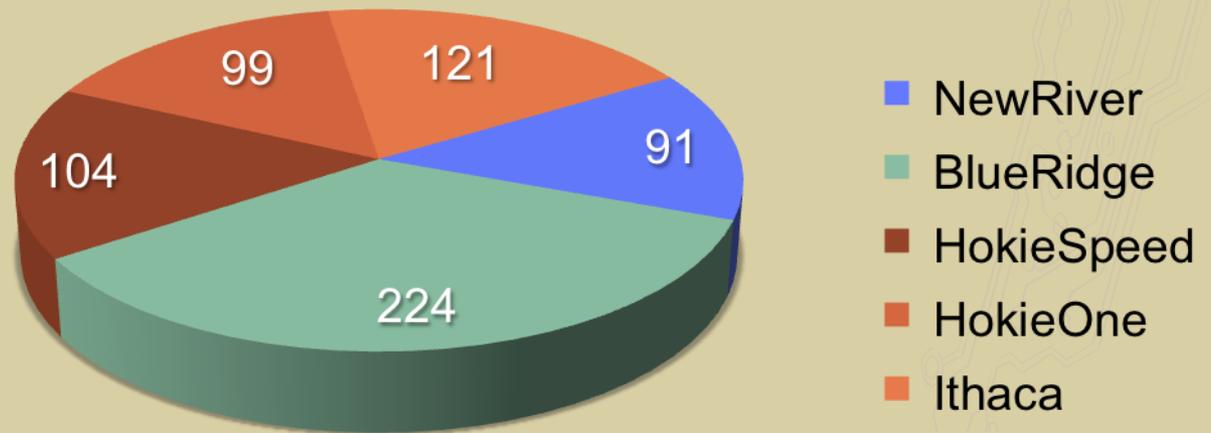
- Our "Workhorse" Cluster
- 408 Nodes (6528 cores) with 64 GB/node
- Remote GPU Visualization
- 130 Nodes each equipped with 2 Intel Xeon Phi (MIC) 5110P coprocessors



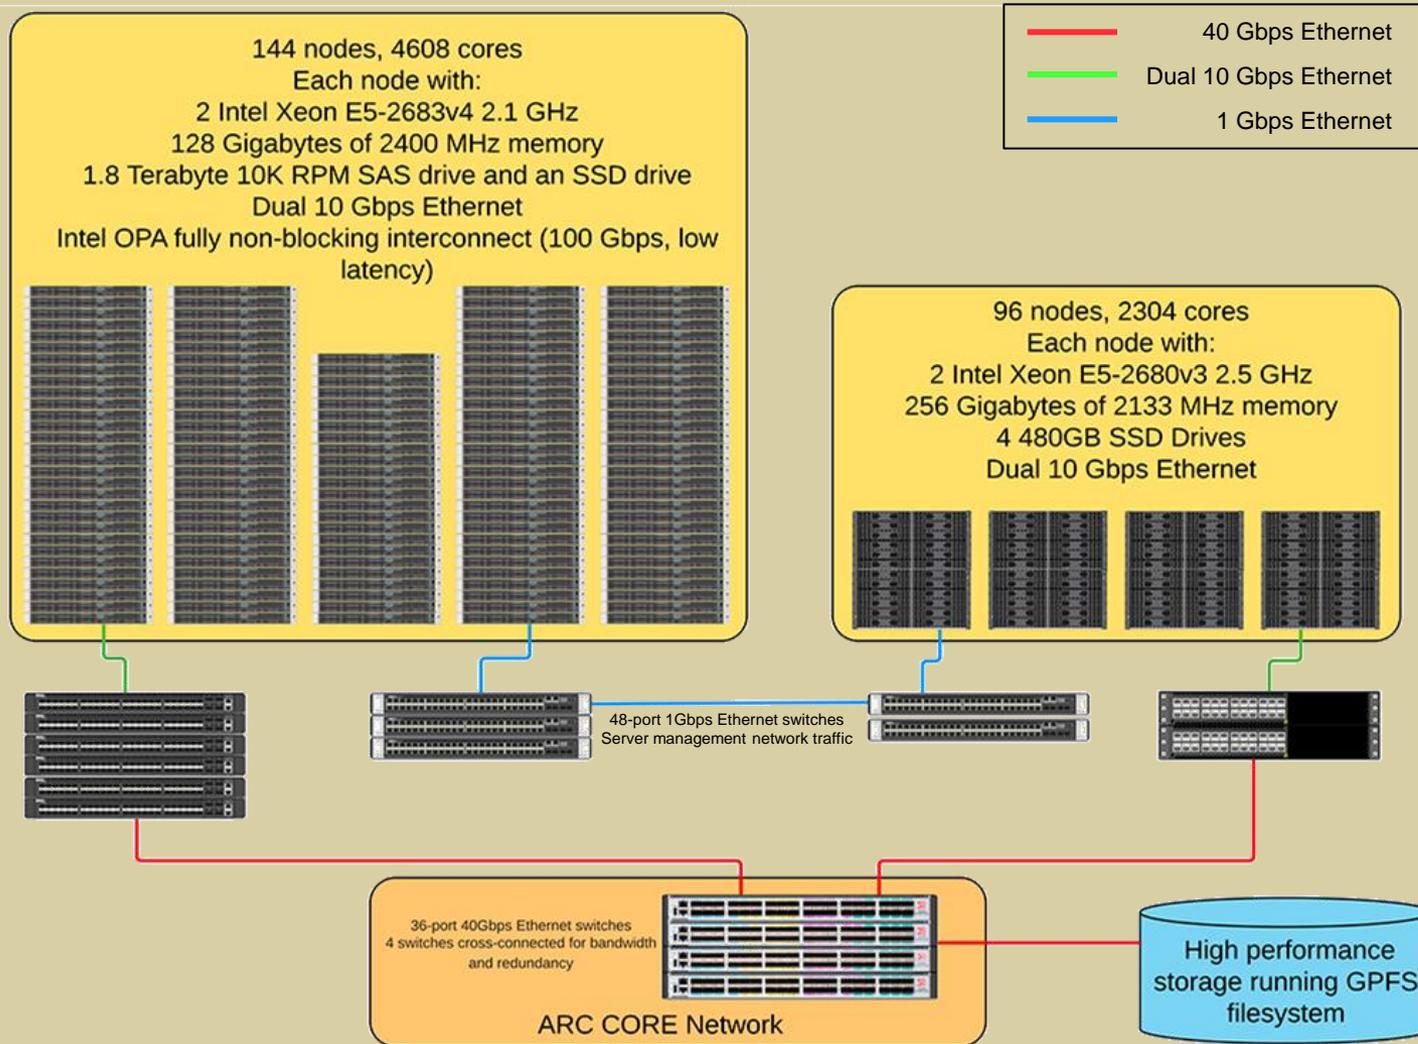
# HPC System Users

- 476 Registered User Accounts - last year
- 28% Used 2 or More Systems

## System User Distribution



# New Compute Engines and Storage – FY16



# Investment Computing Program

- Provides priority access to HPC Compute and Storage Systems.
- Alleviates the burden of power/cooling and more importantly system administration from researchers.
- Allows ARC to provide much better pricing by buying at larger scale.
- Allows more flexibility as a researcher now has access to an entire cluster rather than X nodes they purchased.
- Contact: James McClure
- MOU: <https://secure.hosting.vt.edu/www.arc.vt.edu/investment/>

# Investment Computing Program

- System - BlueRidge

FACULTY	DEPARTMENT
Kevin Wang	Aerospace & Ocean Engineering
Bhuvana Srinivasan	Aerospace & Ocean Engineering
Ioannis Koutromanos	Civil & Environmental Engineering
Linsey Marr	Civil & Environmental Engineering
Hongliang Xin	Chemical Engineering

- System - NewRiver

FACULTY	DEPARTMENT
D. Sarah Stamps	Geosciences
Ryan Pollyea	Geosciences
Nicholas Mayhall	Chemistry

# Visual Computing Group

- Desktop & Web Visualization
- Virtual Platforms
- Immersive Environments
- Collaboration, Consulting, and Support



# High-Performance Visualization

## Support and development

- Consult with faculty and student researchers to apply advanced visualization technology.
- Train users on how to develop advanced visualization techniques and effective visualization of research.
- Develop cutting-edge visualization solutions for domain experts/HPC users.
- Development additional grants and funding streams with domain experts to include visualization tools and HPC.

# High-Performance Visualization

## Facilities

- NewRiver, BlueRidge, and HokieSpeed provide large memory and GPUs for interactive and batch computation, analysis and remote visualization.
- **Visionarium Lab** includes reconfigurable display venues including Gigapixel and multi-touch tables.
- **VisCube Immersive Theater** provides rendering power, resolution and brightness and wireless to accommodate the size and complexity of modern applications.
- Cutting-edge software stack for numerous domains, emphasizing portability and ease-of-use.

# Networked Learning Initiative

- The mission of NLI is to provide faculty, staff, and students access to technology, knowledge, and resources for enhancing digital fluency and improving teaching and learning at Virginia Tech.
- From its inception in 1993, thousands of short courses and workshops have influenced the way faculty, staff, and students engage and learn in the 21st century.
- **ARC participation** last year represented 30 separate workshops with many offered on multiple occasions (workshops are hosted via NLI, ARC and XSEDE classes).
- **691 participants** attended these workshops.

# ARC & NLI Classes

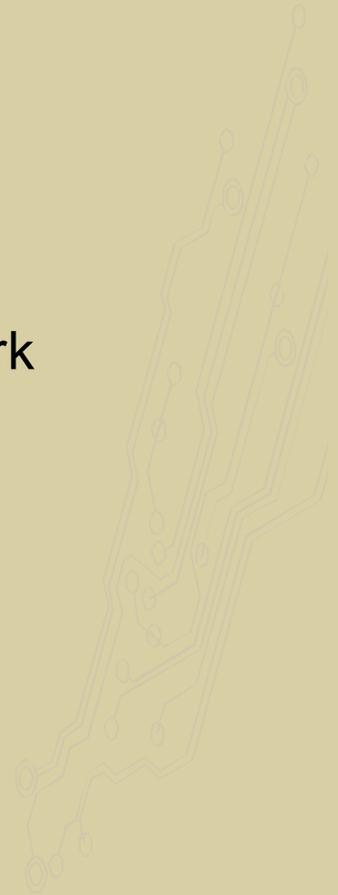
- Introduction to ARC Systems and User Environment (8)
- Introduction to HPC (7)
- Introduction to Python (4)
- Introduction to Visualization (3)
- Visualization Zoo
- Web 3D Publishing (2)
- Introduction to CUDA (2)
- Shared-Memory Programming in OpenMP (2)
- Advanced CUDA Performance Optimization Techniques (2)
- Parallel Performance Implicit Finite Element Methods
- 3D Virtual and Augmented Reality
- Parallel R I: Snow
- Programming Models for Multi-Threading
- Parallel R II: Rmpi and pbdR
- Parallel Programming Workshop
- Programming in MPI – Part 1 & 2
- Parallel Matlab II: PARFOR
- Parallel R Tutorial

# ARC & NLI Classes (cont.)

- Big Data I & II
- Parallel Matlab III: Single Program Multiple Data
- Hybrid Programming in OpenMP & MPI
- Faster Code for Free: Linear Algebra Libraries
- Managing Big Data Transfers & Sharing with Globus
- GPU Performance without CUDA: OpenACC
- Parallel Debugging and Optimization with Allinea
- Working with Big Data – Part 1 & 2
- Parallel Matlab I: Introduction and VT Resources
- Thread Building Blocks
- VSCE/XSEDE Supercomputing for Everyone Series: Performance Tuning Summer School
- VSCE/XSEDE Science Visualization
- Software Carpentry Workshop

# Other Initiatives

- Improved Analytics
  - User demographics
  - Project metadata
  - Data management
- Long-Term Archival Storage Framework
- Research Software
- New Data Center Planning



# HPC Investment Committee

- Scott Midkiff, VPIT / CIO – Chair
- Chris Barrett, BI
- Lay Nam Chang, COS
- Tom Dingus, VTTI
- Michael Friedlander, VTCRI
- Terry Herdman, ARC
- Ben Knapp, ICAT
- Roop Mahajan, ICTAS
- Theresa Mayer, VPRI
- Eric Paterson, AOE
- Cal Ribbens, CS

# Sponsors

*These organizations/groups have generously donated resources and money to make this event possible.*

- ✦ **Department of Aerospace & Ocean Engineering**
- ✦ **Biocomplexity Institute**
- ✦ **Carilion Research Institute**
- ✦ **Department of Chemical Engineering**
- ✦ **Department of Chemistry**
- ✦ **Department of Computer Science**
- ✦ **Department of Geosciences**
- ✦ **Interdisciplinary Center for Applied Mathematics**
- ✦ **Office of the Vice President for Information Technology**
- ✦ **Department of Mathematics**
- ✦ **NI&S: UNIX Administration Services**
- ✦ **College of Science**
- ✦ **Office of the Vice President for Research and Innovation**
- ✦ **Virginia Tech Transportation Institute**