

# NATIONAL CENTER FOR SUPERCOMPUTING APPLICATIONS

Mark Fredricksen

System Administrator, UIUC Campus Cluster



# NCSA History



- One of five original NSF-funded supercomputing centers
- Supported by state of Illinois, University, NSF and others
- Hub for research for University faculty, staff, students and their collaborators from all over the world
- Provides computing, data, networking and viz resources and expertise

# NCSA History



- Mosaic / TELNET
- Smarr / Seidel / Black holes / LIGO
- John Kogut (left) doing quantum chromodynamics calculations
- More than \$1.3B in funding
- TeraGrid / XSEDE / LSST / DES
- Invaluable contributions to science, computational science, HPC

# NCSA Machines



- Blue Waters

- In first 3 years of operations (March 2013-16) Blue Waters provided 12.1 billion core-hours to scientists and engineers across the country
- Can complete more than 1 quadrillion calculations per second (sustained) and peak at more than 13 times that
- HIV capsid simulation
- Chemistry, physics, cosmology, all fields of science and humanities!

# NCSA Machines

- Campus Cluster
  - Allows for professors to purchase nodes
  - Wide variety of hardware
    - Nodes have been added over its 5 year life
    - Purchasers have a variety of needs
  - The professor's authorized users have priority access to their nodes
  - When the nodes are not in use the nodes are made available for general use for up to 4 hours

- iForge

- Production-grade HPC cluster for industry's applied R&D, production, commercialization, and post-commercialization computing and data analytics use cases
- Yearly upgrades to incorporate the latest compute, memory, data, and networking technologies
- Multipetabyte GPFS filesystem w/ custom metadata acceleration; 80 Gb/sec WAN connectivity
- Many hundreds of industry users in the aerospace, agriculture, consumer goods, energy, life sciences, health care, and technology sectors

# NCSA Machines

- ROGER
  - CyberGIS
  - Openstack/Hadoop/Standard Compute
  - Stateful Nodes
    - Higher memory needs
    - Cheaper storage
    - Less efficient management

# NCSA Machine Management

- xCAT
  - An extreme Cluster Administration Toolkit
  - Open source
  - xCAT 2 started by IBM
  - Allows for image creation and PXE booting
  - Scripts that can be run during image install, or post install, to customize the image
  - Tools allow for parallel copying, running commands, and power control

# NCSA Authentication

- LDAP
  - Campus LDAP server
    - Ties authentication in to the same user name/password used for all campus activities
    - Unique to Campus Cluster
- Kerberos
  - NCSA Kerberos server
    - Additional Access Control
    - Off-campus/Non-affiliated users
- OTP
  - Far more secure
  - Requires the user to activate a token
  - Limited access points, jump hosts

# NCSA Storage

- GPFS
  - Parallel file system
  - Robust
  - Performance tweaks
- Lustre
  - Parallel file system
  - Open source
  - Fast
  - Harder to tell what went wrong
- Local Drives
  - Fast
  - Volatile

# NCSA Networking

- Infiniband
  - Fast
  - Proprietary
  - Expensive
- Ethernet
  - Inexpensive
  - Off-site & supplemental communication

# NCSA Resource Management

- Torque/MOAB
  - Priority
  - Fair Share
  - Wall time
  - Job arrays for large numbers of jobs
  - Made with HPC in mind
- Condor
  - Simpler
  - No wall time
  - Easy to submit large numbers of jobs

# High Performance/Throughput

- HPC

- Multiple simultaneous machines
- Fast networking
  - Expensive
- Parallel file system

- HTC

- Many machines, but not simultaneous
- Many small jobs
  - Scheduler to handle

# NCSA Campus Cluster

- Campus Cluster Issues

- Many investors = Many needs
- Initially set up as other HPC systems at NCSA
  - Expansion
  - Usage
- Pricing model
- Lack of software support
- Aging infrastructure

- Campus Cluster Benefits

- Lots of resources at a low buy in
  - \$7270-\$9950 non-GPU node cost
  - \$15903-\$18328 GPU nodes
- Looking to offer computing as a service for lower up front costs
- Heavy users/high peak users get to make use of unused resources

