

# Spring 2019 ITS-CI Faculty Advisory Committee Meeting

Gwen Jacobs, Michelle Choe, Sean Cleveland,  
Adriana Comerford, Maria Dumanlang,  
Jennifer Geis, Ron Merrill,  
David Schanzenbach

April 5, 2019



# Outline

- Cray warranty expired March 15, 2019
- Cluster Reorganization
  - Status
  - Timeline
  - New Equipment
- Long Term Storage
- Potential new hardware expansion
- Proposed Policy Changes



# Cluster Reorganization

- Virtualization of login, data transfer (DTN) and mgmt (xCAT) nodes
- Upgrade of operating system to CentOS 7.x
- Tighten security
  - Multi-factor authentication
- Upgraded connectivity from 10G to 100G
- Added a lower cost user paid storage option
  - ~ 1 PB of NAS storage at \$100/TB/5 yr contract
- Upgrade of Lustre Scratch filesystem
  - Upgrade from 2.5.x to 2.12.x
- Physical reorganization of nodes and improved power/cooling
- Update of CI resource usage policies



# Cluster Reorganization - Status

- Equipment has been purchased and recently deployed
- Virtual login nodes and DTNs have been created
- xCAT has been configured and boots nodes with CentOS 7.x
  - Wide deployment blocked on Lustre upgrade
- Lustre upgrade blocked on scheduling with user base



# Current Timeline

- May 1 - Public user migration/testing CentOS 7
  - Start with 40 std nodes and 1 Im node
- June 1 - Move another 80 std, 1 gpu and 3 Im nodes
  - PI node moves?
- July 1 - Lustre down
  - All machines moved to CentOS 7
  - CentOS 7 cluster up with NAS scratch
- August 1 - Lustre upgraded and reformatted and mounted on CentOS 7



# New Equipment

- ITS has invested \$300K+ for upgrades
  - 2x Firewall for WAN access for login and compute nodes
  - 2x vHost for non-compute infrastructure
  - ~144 TB NAS for home and vHost hosting
  - ~1 PB NAS for purchase
  - 2x 100G Ethernet switches
  - 3x Active Chill doors
  - 4x 60A Breakers and 4x 60A PDUs
- PI purchased GPU nodes
  - Peter Sadowski - 2 ( 16 GPUs)
  - EPSCoR - 2 ( 12 GPUs)
  - Rui Sun - 3 ( 24 GPUs)



# Long Term Storage

- NAS based storage with ~1 PB of storage
- NFS access from all HPC resources
- ZFS with 10 Raidz-3 arrays (12 disks per array 9+3)
- Snapshots available on a per allocation basis
- Pricing for 5 year contract is at cost of the hardware
- Sales currently scheduled to begin May 1, 2019

<b>Contract period (Years)</b>	<b>Storage amount (Terabytes)</b>	<b>Total Cost</b>
5	1	\$100.00
4	1	\$112.00
3	1	\$96.00
2	1	\$72.00
1	1	\$40.00



# Potential new hardware expansion

- MRI
  - 24 CPU based compute nodes (960 cores)
  - 6 GPU nodes (24 Nvidia Tesla v100)
  - Flash NAS scratch file system (80 TB)
  - HDR100 infiniband
  - 25 Gb Ethernet
- Jetstream
  - 16 CPU nodes (2,048 cores)
  - 3 GPU nodes (9 GPUs)
  - 144TB of flash storage
- Mid-scale
  - 25 CPU based compute nodes (1000 cores)
  - 600 TB of NAS storage
- Node purchases slated to resume for fall purchase window





# Proposed Policy Changes

- User quotas
  - home 50 GB, up to 300 GB by request (replaces ~/apps)
  - scratch 5 TB for NAS & 50 TB Lustre
- Inactive accounts
  - Locked after 120 days
  - Removed 60 days later (user notified prior to removal)
- Ohana users removed after 1 year
- Free compute resource allocations
  - Teaching Allocation (condo style)
  - Startup Allocation (credit/elevated priority)
  - Production Allocation (reservation/condo style)
- Service Units will be discontinued
  - Lack of adoption



# Discussion

