

VIRTUAL TOWN HALL JUNE 2023





To Participate:

- Type Questions into Q&A
 - You have the option of selecting "Anonymous" when submitting the question
- Raise your Hand to Speak

AGENDA



- Introducing new Research Computing & Data (RCD) People
- The Condo Model is alive & well
- New Services since last Town Hall
- Maintenance Week of July 10-14
- Migration from PBS to SLURM
- REDCAT Faculty advisory body
- Indigo highspeed storage/retiring zfs
- Slow responses week of July 24-28
- Open discussion and Q&A





The Condo Model is Alive and Well



Double Research by 2035 Initiative 2: Create world-class research infrastructure

Key actions:

- Increase total research space from 861,000 to 1,100,000 square feet by 2028 and 1,300,000 square feet by 2035.
- Align new space additions plus renovations with the cluster hiring strategy.
- Add research administrative staff to provide the foundation and management of the growth.
 - Upgrade high-performance computing (HPC) systems and add specialized HPC research faculty.
- Add additional centrally operated research cores.



The Condo Model is Alive and Well

- Is Palmetto available for everyone at Clemson to use, at no charge? YES
- Can faculty purchase nodes to own, for priority use and to share with external collaborators? **YES**
- Does RCD provide data center space, system administration, user support and training at no charge? YES
- Does RCD provide racks, network, power and cooling at no charge? **YES**
- So, what has changed?
 - Defining on-going funding will be addressed as part of RBB
 - Node prices are increasing because
 - Many more CPU/GPU cores packed into each node
 - We no longer subsidize by purchasing GPU cards



NEW SERVICES SUMMARY

	Palmetto Documentation	
bout		
Getting Started	Palmetto Documentation	on
counts >		
mpute >	Welcome to the documentation for the Palmetto Cluster!	
b Management	Palmetto supports the research mission of <u>Clemson University</u> through innovative High Performance Computing (HPC) and Storage solutions. Our <u>Research Computing & Data</u> team helps students and faculty by providing user support, training, and hardware maintenance.	
onnecting >		
torage >		
ata Transfer		
cknowledgments	Recommended Pages	
xamples >	Here are a few pages we recommend taking a look at:	
AQ		
nboarding Sessions	Getting Started	About the Cluster
oftware >	We're excited that you are here and interested in using the Palm	Palmetto cluster racks at the data center
pport		
	lacs red clams	on odu/nali
		ULLEUU/ Vall



GITLAB

https://xdmod.palmetto.clemson.edu (on campus or VPN needed)

MATTERMOST – chat w. us & each other

Kubernetes/K8 in beta



REDCAT Advisory Body

Research Data Computational Advisory Team

- Two-way conversation between faculty users and RCD team
- Open meetings w. representative steering committee
- Term-limited positions
- Foster Clemson computational community activities
- MONTHLY commitment; help w. administrivia

If interested in helping to define and/or serve: Contact gemmill@Clemson.edu





MAINTENANCE JULY 10-14

- Operating System and Driver Updates

- Updates to NVIDIA GPU Drivers to support newer CUDA versions
- Update HDR IB Drivers for bug fixes
- New Feature Deployment: GPU Direct Storage Access Indigo Data Lake direct from node GPUs
- System Patching
 - Compute nodes locked to Rocky Linux 8.6
 - Update to Rocky Linux 8.8 elsewhere for latest security patches
 - Update system firmware on cluster hardware
- /scratch, /scratch1, and /fastscratch will be re-initialized
- Move all virtualized services/systems to new VMWare Cluster
- Indigo Filesystem Add additional Ethernet/IB bandwidth to cluster
 - Will have a total of 800Gbit of IB Connectivity, and 400Gbit of Ethernet Connectivity
- Deploy the initial SLURM-based cluster for testing
 - Palmetto will be undergoing a *gradual transition* to the SLURM job scheduler



MAINTENANCE JULY 10-14

Engagement Team Tasks:

- Benchmarking of applications
 - Applications will be benchmarked pre and post maintenance to ensure no regressions introduced
- Moving the scientific software stack to the VAST all-flash system:
 - Improve performance/ Increase Capacity
 - Should be no change in the modulefile naming system
 - User-side minimal changes in workflow
 - Unused modulefiles/software will be removed
 - Software to be removed will be listed prior to maintenance
- Open OnDemand will be updated to the latest stable version



Why SLURM?

- Larger community and support group
 - More parity with peer institutions making workflows more portable
- More actively maintained than PBS
 - New features added regularly
 - Faster response to bug/security issues
- Plug-in support
 - Adds additional functionality
- More efficient and better control over scheduling better fair sharing
- Less management overhead, simpler to maintain





- Supports accounting based on unit allocation
 - Allows the possibility of providing compute as a service
- Priority access to nodes owned within condominium model maintained
 - Retains existing PBS behavior for node owners
- Automated fair share for user jobs
- Tools available to convert PBS scripts to SLURM
- Built in tools to check the status of job(s)





- Stage1 Initial Deployment (Beta) July 2023
 - 20 Phase29 Nodes moved to SLURM
 - Separate Login Node
- Stage2 Early Access (EA) December 2023
 - Remaining phase29 nodes moved to SLURM
 - Owner resources moved to SLURM on request
- Stage3 Production Summer Est. 2024
 - All nodes Sky Lake or newer moved (All nodes on HDR IB) to SLURM
 - Cluster now known as Palmetto 2
- Stage4 Est. 2025
 - Palmetto 1 cluster converted to SLURM



SLURM – Transition Assistance

- New Documentation on using SLURM
 - Will be based on current documentation
- Reference SLURM scripts
 - Examples of what PBS scripts will look like as a SLURM script
- Additional Office hours to help troubleshoot
- Tools to convert PBS scripts to SLURM
- Workshops/Training on using SLURM



STORAGE – What is Indigo?

- Approximately 5PB of all NVMe flash with erasure code resiliency
 - Total Capacity varies due to Data Reduction and Deduplication
 - All paid storage backed-up, snapshots also available
 - Large speed-up over current BeeGFS and ZFS filesystems
- High resiliency architecture to reduce downtimes
 - On-line software updates, hardware fault tolerant
- High Speed Interconnects for better performance and added features
 - Infiniband connectivity provides GPU-Direct Storage, NFSoRDMA
 - Ethernet Connectivity to Campus resources, and C1 compute resources



STORAGE – Indigo Features

- Unified storage space for all filesystems
 - NFS, SMB, S3 access to same directories
 - 1TB of project space for researchers at no charge
 - Scratch space provided (available now at /scratch)
 - 5TB space limit and 5 million file limit per user
- Data analytics using Starfish provided by request



STORAGE – Indigo in, ZFS out

- ZFS project space no longer being sold
 - Indigo will replace ZFS
- New rates for Indigo storage coming
 - Will be sold as service by the TB by the month
- ZFS renewals, including partial dataset renewals, result in entire dataset being moved to Indigo
 - Renewals will be charged at Indigo prices
 - Active agreements for remaining partial dataset will be honored
- Starting in July 2024, ZFS filesystems will be migrated to VAST as ZFS equipment is decommissioned
 - Any existing storage agreements honored in VAST



SLOW week –Conference "Practice & Experience in Advanced Research Computing" (PEARC)

July 23-30 RCDE responses (Help Desk Tickets) will be slow – staff are at a training conference

Systems will be being monitored and infrastructure team is available

Discussion & Questions?

