
SciComp Compute Coordinator Meeting

Aug 3, 2023

Brad Sawatzky

Data Center Outage (Aug 10?)

- Duration
 - Aug 10 (ish) – TBD!
 - » 6am – EOB
 - Work needs to happen during work day due to safety regulations
 - » Problem, slow work may require overflow into 2nd day... (NOT the plan though)
- Work being done:
 - 2nd temp chiller installation for some failover/maintenance capability
 - electrical infrastructure maintenance and improvements
- Systems affected
 - tape library
 - Farm nodes
 - LQCD
- Unaffected Systems
 - filesystems stay up
 - ifarm node status unclear?
- *Watch for notice/emails to scicomp list, etc.*

JLab (internal) SciComp Review 2024

- Tentative dates:
 - First week of Feb, 2024
 - Formal Charge and dates distributed soon
 - *Let me know if you see a conflict/issue with Feb!*
- Nominal 2-day review following historical pattern
 - [Dec 2021 Review Indico](#)
[Access Key: "JLAB2021"]
 - » See [Review Report](#) and [Responses](#) documents on that page
- Will require updates on
 - Scientific Computing Systems
 - » incl. OSG usage/plans
 - EPSCI Report
 - Hall Reports/Projections
 - Data Science Report
 - EIC, Theory Reports/Projections
- Possible special topics
 - SRO status and plans
 - AI/ML progress
 - HPDF status and pre-planning(?)

Projected Rates / Compute Resource Projections

- [2023 Hall Projections Spreadsheet](#)
 - Link to 'summary' spreadsheet
 - Includes some new updates from Hall D
- Thia, Doug, Patrizia, would like to have greater confidence in the projected data rates
 - Impacts multi-Hall operation and long-lead procurement cycles
 - Avoid 'data crunches' requiring exceptional resource juggling and surprise expenses
- Data / Compute resources are finite
 - Goal is NOT to penalize overruns but to re-emphasize that data-center resources are finite. Think of them like total accelerator beam loading or cryo.
 - Effective projections matter to allow CST to be ready for your experiment!
- **Hall Coordinators: Please confirm numbers on the next couple slides**
 - Drawn from data submitted previously
 - » Make sure it still looks correct
 - These will be sent to Thia / Doug to document expectations for the future.
 - » Updates will happen annually. You're not locked in to 5 year old projections!
 - *Excessive* over-runs may require discussion between Hall leadership and Thia to rebalance priorities in data center
 - » May require a Hall adjust trigger, detector readout, or dial back current

Projected rates / Compute resource projections

- **Hall A**

- Mostly concerned about GEp-V in Hall A (2024–25)
 - » 1200 MB/sec to tape (peak); 600 MB/sec average
- GEn-RP has similar number of wire-chambers as GEp-V (2024)
 - » 400 MB/sec to tape (peak); 200 MB/sec average
 - » Are the numbers above plausible given recent experience?
- MOLLER (2025–28)
 - » 150–200 MB/sec sustained (90% 'uptime' assumed)

- **Hall B**

- Projections seem to be inline with historical running
 - » 800–1000 MB/sec (peak); 400–500 MB/sec average

Projected rates / Compute resource projections

- **Hall C**

- NPS in Hall C (2023–24)

- » 300 MB/sec (peak); 150 MB/sec sustained
 - » N.B.: These are extremely sensitive to trigger configuration and NPS read-out settings.

- **Hall D**

- updated Farm usage numbers that are closer to history

- » ** future projections still seem low to me; continued monitoring is greatly appreciated!*

- **All Halls**

- SRO-related tests can be very high-rate

- » Please check around and keep SciComp in the loop

Hall ESX Virtual Machine Cluster

- VMs within the Hall experimental enclave will be possible within weeks
 - » High-uptime infrastructure explicitly targeted to support Hall beam operations, etc.
 - » Update policy, downtime scheduling, defined by Hall Compute Coords as usual.
- Slow control systems
 - » EPICS softIOCs
 - » Windows/Rockwell cryo controls
 - cmagnets, skylla10
 - » PXE boot hosts/services
- Data-base hosts
 - » RCDB, CCDB hosts
- “Remote CH” support hosts
- VM hosts will functionally operate within the Hall subnets (no firewall issues)
 - Direct access within respective subnets
 - 2-factor hop (via. hallgw) as with any existing Hall hosts
- Take advantage of VM flexibility
 - snapshotting / backups
 - auto-failover on HW issues
 - advantages wrt “cloning”, load balancing, etc
- Keep this in mind for future deployments, HW upgrades, etc

Ongoing Containerization Support

- CST will be developing some more formal containerization support for Users
 - Brad, Bryan, Wes, Laura will talk next week to lay groundwork
 - Documentation
 - » ie. Easy 'on-ramp' / how-to
 - Infrastructure
 - » Docker, Podman, etc?
 - What else?
 - » We want your input/advice on what your Halls use, friction points, etc
- *Please email brads@jlab.org with suggestions, advice, input!*
- Among other benefits, containers can provide
 - 'Plug and play' software configurations
 - Version / configuration snapshots
 - Ability to run 'custom' software frameworks on other datacenters, computers, laptops
- Also necessary / encouraged to streamline upcoming Farm transition from RHELX → ALMA9
 - *(No date for this yet...)*

Miscellaneous Meeting Notes

- Issue with /work/halld/ quota filling due to job submission glitch
 - issue is being addressed
- Some discussion into making it easier to see /work usage numbers (ie. what subdir is big consumer, etc)
 - ie. more granular than what is under <https://scicomp.jlab.org/scicomp/workDisk>
 - Bryan will look into improving reporting on the web page
- Data flows from Halls to data center not easily viewable
 - Bryan will look into adding a page/tab for this
 - Add LQCD data flows as well?