

Scientific Computing Operations for ENP

December 2024

Bryan Hess

Thursday, December 5, 2024

 **Jefferson Lab**

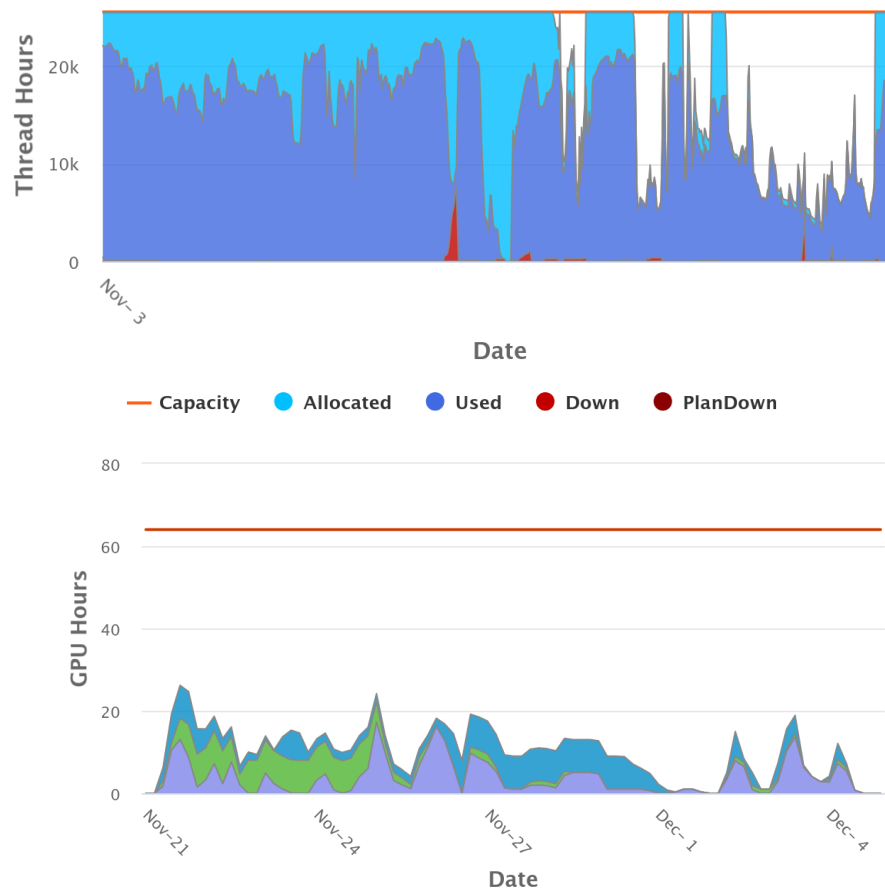

 U.S. DEPARTMENT OF
ENERGY | Office of
Science



Farm Operations Overview for November

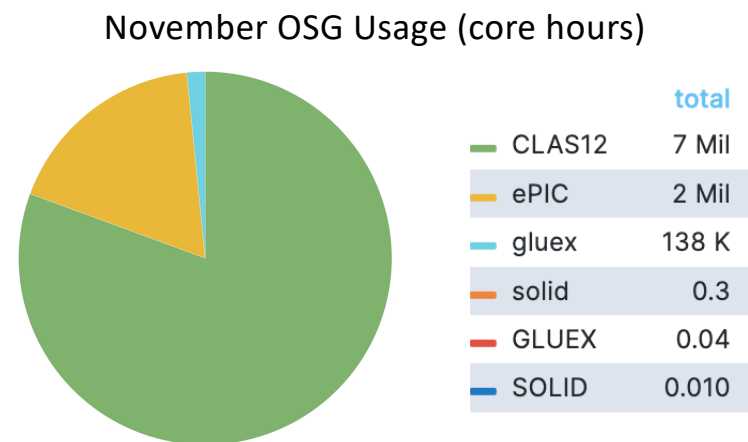
- Excellent farm node availability: >99%
- Strong demand early in the month, light use later in the month.
- No significant issues.
- There was a brief, planned outage for for read-only /cache transition.
- GPU demand remains steady, but below capacity.

Slurm Accounts Usage (CPU Hours)



OSG Usage for November

- November OSG Usage remains strong with five JLab-affiliated projects running jobs.
- We are working with OSG to increase the backfill on the farm for empty job slots, aiming to increase utilization and reduce idle CPU in periods of low demand
- Reminder: Token based writes back to JLab Lustre are available using Pelican. This requires some manual steps to map tokens to users and writeable areas on the Data Transfer Nodes. Please submit a ticket for new requests.



Status and final steps for read-only /cache transition


- Successful completion of this change on November 19th
 - /cache is now read-only from farm, ifarm, and data transfer nodes
 - We received no problem reports after this change.
- Clean up of loose ends:
 - Directory reconciliation /mss vs. /cache underway to clear up
 - Small files will be pushed to tape to make (and maintain) /cache a subset of /mss
 - All files on cache will be treated equally with regard to deletion policy, regardless of size.
- Have there been any significant challenges from your perspective?

Gitlab - Code.jlab.org update

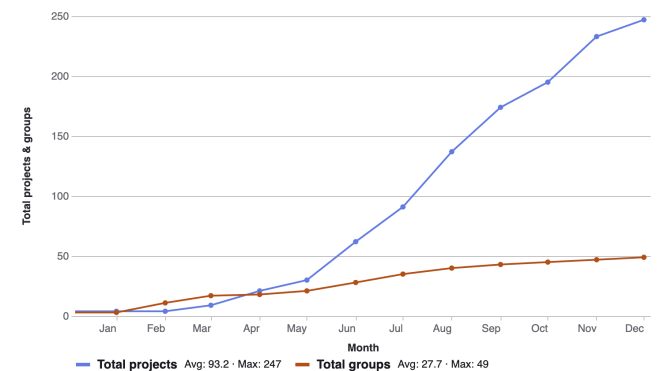
- Usage continues to increase.
- Regular security and stability updates
- Open Issues
 - Increased object storage capacity for containers
 - CVMFS integration for CI/CD (this has OpenShift/Kubernetes interactions that are being worked through)
 - Changes to the user onboarding code to improve the experience
 - Changes to pages.jlab.org to support logins as an option for documentation.

Projects	Groups	Users	Issues	Merge requests	Pipelines
249	50	124	96	113	966

Total users



Total projects & groups



Upcoming Work: December

- Install new ssd-based storage for data caching from the halls
 - lifecycle replacement
 - This was awaiting networking and is now read
- Retire end-of-life QDR InfiniBand switches in the farm and server racks
- Starting with December maintenance, dormant SWIF workflows will be archived to history tables
 - Dormant: any workflow that has been inactive for 6 months or more
 - The first sweep to history tables will be done on maintenance day in case of long table locks
 - Subsequent sweeps to history will be done as a regularly scheduled jobs
 - Commands and web portal interfaces to examine archived workflow history will be provided.
- A check will be added to Jasmine to confirm that files written to cache or other disk pools have the expected file size (to address a rare file truncation issue)
- Planning for stable farm operations over the break. Minimal staffing during the shutdown, support for significant outages but not for changes or day-to-day support

Other Topics

- Container Security
 - Patching and update strategies for containers remains a challenge
 - DockerHub security grading shows many security vulnerabilities
- Container Support for Farm Jobs
 - Apptainer support continues as the standard solution
 - Podman is available after last maintenance day change; some loose ends, if you use it, please submit ServiceNow tickets for problems
- Work after the New Year
 - Farm25 procurement
 - Database Server upgrades
 - preparation for running experiments
 - CVMFS
 - stratum zero servers
 - support on CI/CD

Hall Data Management Plan

- (Placeholder for Discussion)