

# Scientific Computing Operations for ENP

August 2024

Bryan Hess

Thursday, August 1, 2024

 **Jefferson Lab**

 U.S. DEPARTMENT OF  
**ENERGY** | Office of  
Science



## Lustre Transition to new filesystem August 19, 2024

- The farm, ifarm, SWIF, /cache, /volatile, and tape storage services will be offline August 19 9:00am – August 21 9:00am for the Lustre upgrade
- Farm nodes will be reserved starting August 19 at 9:00am. Jobs that would run past that time will be queued until the upgrade is completed.
- Files in /cache and /volatile will be moved to the new servers
- We are synchronizing the filesystems in the background and are aiming for a small final sync.
- If the final synchronization runs long, we will stop it at 9:00am on August 21. Any "missing files" will be those written in the few days before the upgrade. We will maintain the old filesystem and can copy them. We will maintain /lustre19 as a backup for at least one month.
- As a reminder, we are making this "hard cutover" because of a limitation in older hardware that prevents us from running both Lustre filesystems simultaneously.

```
filesystem_summary:      bytes      Used      Available Use% Mounted on
                        4.7P      3.7P      1008.3T  79% /lustre19
```

```
filesystem_summary:      bytes      Used      Available Use% Mounted on
                        11.2P     1.9P       9.4P   17% /lustre24
```

## Asking for Your Help with The Lustre Transition


---

- To facilitate the Lustre transition on August 19<sup>th</sup> :
  - Consolidate or delete small files as possible (smaller than, say, ~100MB)
  - limit directly writing large file sets to /cache or /volatile (if possible) the week prior to the upgrade. This will help to keep the final file system synchronization from running long
- Important Note: Using jcache is not a concern. We presently write all jcache requests to both lustre filesystems.

## Open Science Grid Progress and Notes

- Work in recent weeks has been on EIC workflow debugging.
  - taking an in-depth look at submissions, job matching, efficiency, and accounting
  - we will take lessons learned here and apply them to other projects to improve job efficiency and accounting
  - OSG Debugging requires moderate coordination effort.
- Standing up a new sumbitter host for SoLID
  - This brings us to 5 projects: CLAS12, GlueX, MOLLER, EIC, and SoLID.
- Usage remains steady, with ePIC, GlueX, and CLAS12 showing highest demand.

### OSG Usage, Last 90 Days



## Open Science Grid Next Steps

---

- Next Up after submitter and job workflow is a look at data movement
- Pelican/OSDF installation is in progress on a new Data Transfer Node
- Following that, CVMFS Stratum 0 for JLab.
- Speaking of CVMFS
  - The JLab synchronization delays with OASIS are due to bottleneck that limits oasis updates to one active VO update at a time (across all VOs).
  - With this new understanding, we are adjusting the sync time to 4 hours.
  - This is a limitation not in the rsync itself, but in back-end processing on the OSG side.

## Reminders and Updates

---

- Farm Usage projections due August 16
  - email to Brad following format outlined in email
- Backlogged projects still in the queue
  - “long run” feature for jobs with preemption
  - fairshare modifications
  - Gitlab CI/CD for code.jlab.org (OpenShift complexity; case open with RedHat; working through it)
  - podman on the farm
  - Write-through cache transition to read-only