

SciOps+ENP March 2024 Meeting

Bryan Hess

Wednesday, March 6, 2024

 **Jefferson Lab**



Farm Alma9 Deployment Progress

- Knowledge Base Articles related to Alma9
 - [KB Version of March User Annoucement](#)
 - [Slurm - Requesting features or GPUs](#)
 - [Environment Modules and CVMFS](#)
- Next batch of Nodes being upgraded this week
 - Slurm features used to control node use within production partition
- Ability for larger, farm-local /home
 - starting in April on a user-by-user basis to start
- full podman container workflow
 - subuid allocation plan begin finalized
- Notes
 - No /apps (moving to CVMFS)
 - No /site (tools moving to /usr/local)
 - No jget/jput to on farm nodes
- OS indicator coming to web page. Meanwhile →

```
# As of March 7:  
  
$ sinfo -N -p production -o "%12N  
%50f" | grep e17 | wc -l  
188  
  
$ sinfo -N -p production -o "%12N  
%50f" | grep e19 | wc -l  
46
```

Facility Update

- 208V/50A temporary circuits were added to HACS 3 (Home of farm and LQCD clusters)
 - Energized Friday March 1
 - Needed to be able to fill HACS cabinets to capacity at higher power densities
 - In the coming weeks, certain clusters will be consolidated to make use of the new circuits and free up contiguous rack space. This will begin with farm16 and sciml
 - Making space for recently delivered GPU nodes with A800 GPUs.
- Data Center Expansion Design starting
 - We are near electrical capacity
 - We have the go-ahead to proceed to design
 - Lead time on electrical gear is long (>1 year for PDUs)
- PO awarded for Chilled Water piping fix
 - Awaiting schedule from FM&L

Hardware Deployment Updates

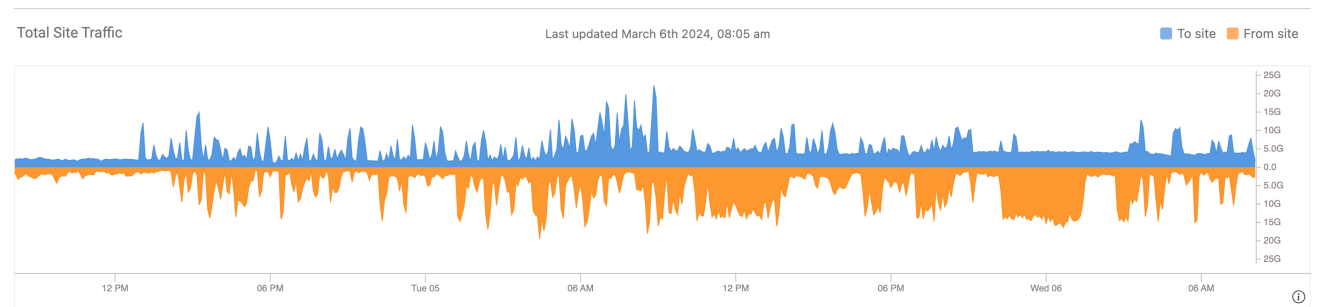
- Lustre 24 Installation
 - Awaiting delivery of metadata server replacement SAS disk shelf
 - Awaiting delivery of back-ordered transceivers
- NVMe storage evaluation
 - evaluation beginning
 - was also stalled for transceiver delivery delays
- Two new GPU nodes (with NVIDIA A800 GPUs) have arrived
 - They will be installed as soon as we make rack space for them.

Gitlab early access: code.jlab.org

- Early User testing phase for Jlab gitlab server, code.jlab.org
- Features
 - Support for container registry
 - Login/ access with home institution credentials, including authorization approval step for non-jlab accounts (open from off-site soon)
- Coming next
 - Cyber security reviewing now to open off site.
 - CI/CD support using Kuberbetes is in development – openshift installation now April.

Networking Changes

- Data Center Fabric: InfiniBand Core is now HDR (200Gbit)
 - Lustre switches being upgrade to HDR now
 - Gradually, leaf switches will be upgraded as part of lifecycle replacement
- The first 100Gbit Internet connection is being used regularly.
 - Data Transfer Nodes (DTN) and Science DMZ were already 100Gig, so increased bandwidth is already available
 - New Data Transfer Node for GlueX production data processing awaiting transceivers



Recent Operational Issues

- This week we had a recurrence of the system-crashing OOMs that we first saw in January
 - This kernel issue fixed in Alma9
- Tape Library Issues
 - several instances of dropped tapes. Replaced grippers and springs under contract.
 - Continue to discover the occasional tape from the 2020 humidity incident where the tape has degraded enough to need recovery by the IBM factory staff.