


SciOps + ENP

Updates for November 2023

Bryan Hess

November 10, 2023

Jefferson Lab



AlmaLinux 9 Farm Rolling Upgrade

- **Expected Timeline**

- November Maintenance Day – small set of Alma9 nodes released for testing jobs (treated as production nodes with Slurm **features** to select, submit issues via ServiceNow).
- January Maintenance Day – begin transitioning the farm to Alma9, reducing availability of CentOS7.9 nodes.

- We welcome your feedback and suggestions!

- To minimize the number of breaking changes to partitions and slurm configuration we will retain the production partition.

- The production partition will contain BOTH CentOS7 and Alma9 hosts. Selection is controlled using Slurm features instead of partitions. (See next slide)

- Use of CVMFS instead of /apps

- More portable; works with OSG as well; can work on other systems; no nfs
- Use of modules to facilitate the change
- What software do we need to provide in CVMFS? (lots of gcc versions...)

- Transition documentation/HOWTOs under development.

Slurm Production partition feature changes for Alma9

- **CentOS 7 Only Jobs**

- No *current* change is required (see deprecation notice*).
- [If no constraints is specified, we will default you to centos79 during this transition.](#)

- **Alma9 Only Jobs**

- Use constraint to specify EL9 nodes (ex. --constraint=el9).
- Deprecated features will not be allowed (see deprecation notice*).
- After transition, no change required.

- **OS agnostic jobs (example: container-based workflows)**

- During the transition, to use all farm nodes set --constraint=centos79|el9
- After the transition, centos79 **MUST** be removed (see deprecation notice*).

- ***Deprecation Notice** - As nodes transition to Alma9, we will deprecate the following *features*:
 - general (same as not specifying any constraints)
 - centos79
 - amd, xeon (unused)
 - gpu, TitanRTX, T4, A100 (not used on production partition and handled by gres as needed)
 - Attempting to use deprecated features on Alma9 nodes will result in a job submission error.

Storage

- Lustre Upgrade from Lustre19 → Lustre24
 - After Supply Chain delays, all hardware has shipped from the vendor
 - We will rack up, burn in, and test all hardware components
 - Lustre installation and performance baseline to begin in January
 - Depending on RAID choices and performance trade-offs, we expect near 10PB of storage, roughly double what we have now.
 - This will ultimately contain /cache and /volatile.

Farm Utilization Improvements

- Farm Utilization and Uptime have been trending up. (90% utilized in October!)
 - We use Slurm job data to find efficiencies
- Tracking memory efficiency
 - Actual job memory use vs requested is tracked, looking for outliers only.
 - If there is a group of jobs with very low memory use the user is emailed
 - Criteria: More than 100 jobs in a day with <20% memory Efficiency
 - If the trend continues for several days, the user's max simultaneous jobs is limited
 - <https://scicomp.jlab.org/scicomp/slurmJob/memoryEff>
 - This is a success; we leave far fewer CPUs idle because of memory starvation now.
- Next Up: Tracking Slurm Job CPU efficiency
 - Looking for outliers only
 - Groups of jobs, not singletons
 - Example: jobs that request 16 cpus and use 8 over the life of a set of jobs.
 - Monitoring phase; no automated enforcement yet.
 - Checking our analysis with individual users to validate what we are seeing
 - This can recover unused CPU over time.
 - <https://scicomp.jlab.org/scicomp/slurmJob/cpuEff>

Work in progress/ Previews

- GitLab (code.jlab.org)
 - Integrated CI/CD
 - Accessible with federated identities from on and off site
- XRootD Storage for GlueX
 - Target is streaming processing of data on OSG
 - DTN hardware identified and being installed
 - XRootD deployment needs to be generalized/operationalized based on existing cases.
- ESNet/JLab Upgrade to 2 x100Gbit
 - We have light!
 - Anticipate commissioning and transition to production in December.
- Tape Library Outage
 - Tentatively planned for Jan 2 – 8 to align with Accelerator down time
 - Adding second robot, space for additional tape drives
 - Reconfiguring power
 - This may include a wider Data Center power outage; This is in discussion with FM&L now.
- Podman support for containers on ifarm and farm