SciOps + ENP March 2023

Status Updates and Upcoming Work

Bryan Hess

Thursday, March 2, 2023

Jefferson Lab

U.S. DEPARTMENT OF ENERGY | Office of Science

JSA

Farm Changes for Feb 2023

- Issues in the memory tracking code resolved this month
  - Slurm job arrays were under-counting memory use. This has been resolved.
  - Nuisance reporting (very small memory jobs) were removed from notification
  - Repeats with large numbers of memory-inefficient jobs will be temporarily limited (MaxJobs)
  - This mechanism is helping to keep CPUs busy
- A total of Eight NVIDIA 80GB A100 GPUs have been added to the farm.
  - These are not part of Jupyter by design to keep people from parking on them
  - Slurm Feature name is gpu:A100
- Apptainer OOM Kernel Bug in CentOS 7
  - There is a kernel bug in CentOS 7 that can allow OOM in Apptainer to take down a node
  - We saw this on Feb 6 where a batch of user jobs crashed many farm nodes
  - We have a kernel module in place to mitigate it, but it is not sufficient in every case. This bug goes away when farm upgrades beyond CentOS 7, meanwhile beware of jobs that severely underestimate memory use and run in a container.
  - https://github.com/apptainer/singularity-userdocs/issues/417

Jasmine Evolution for Efficiency and Community Tool Use

- We have generalized the use of two internal XRootD disk buffer pools for Jasmine to cover all cases:
  - one for writing (including hall data)
  - one for reading to reduce tape reads, especially for small or recent files
- All network traffic to/from Jasmine is now XRootD, including the jget client.
- Since Nov. 2022, the Jasmine internal disk cache system has prevented ~1PB of data extraction from tape.
  - 152,000 hits for 71,000 unique files.
- This has improved farm use, decreased latency for data reads, and reduced the number of drives needed.

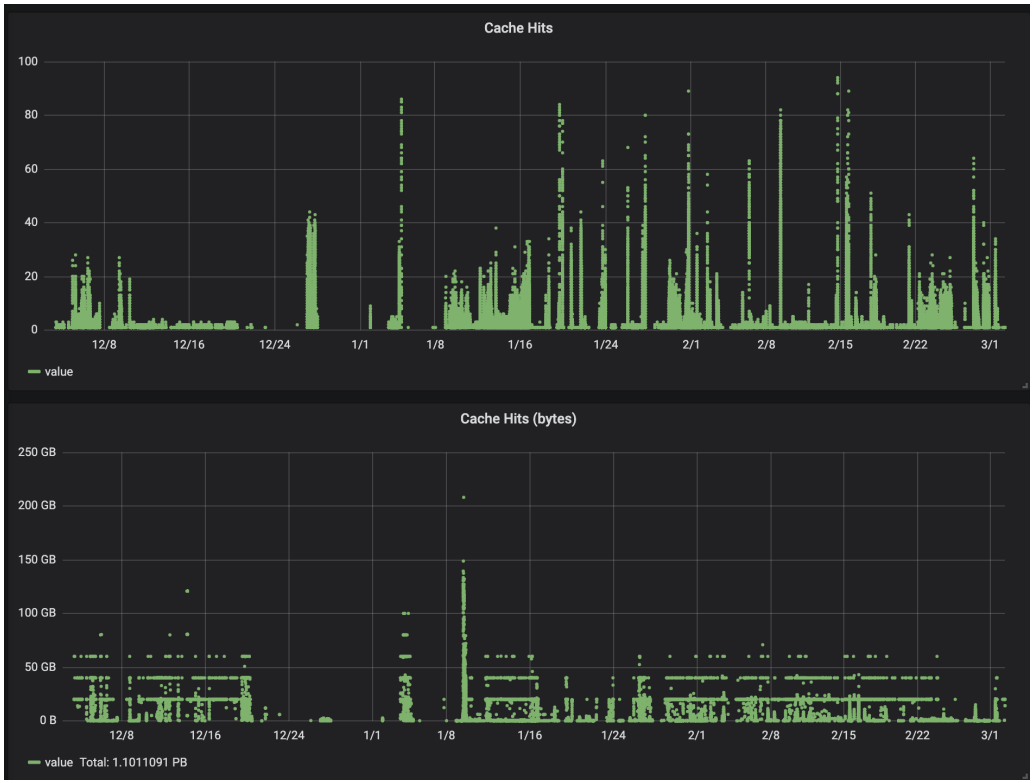

One Week Raw Data Ingest: from halls (+) / to tape (-)

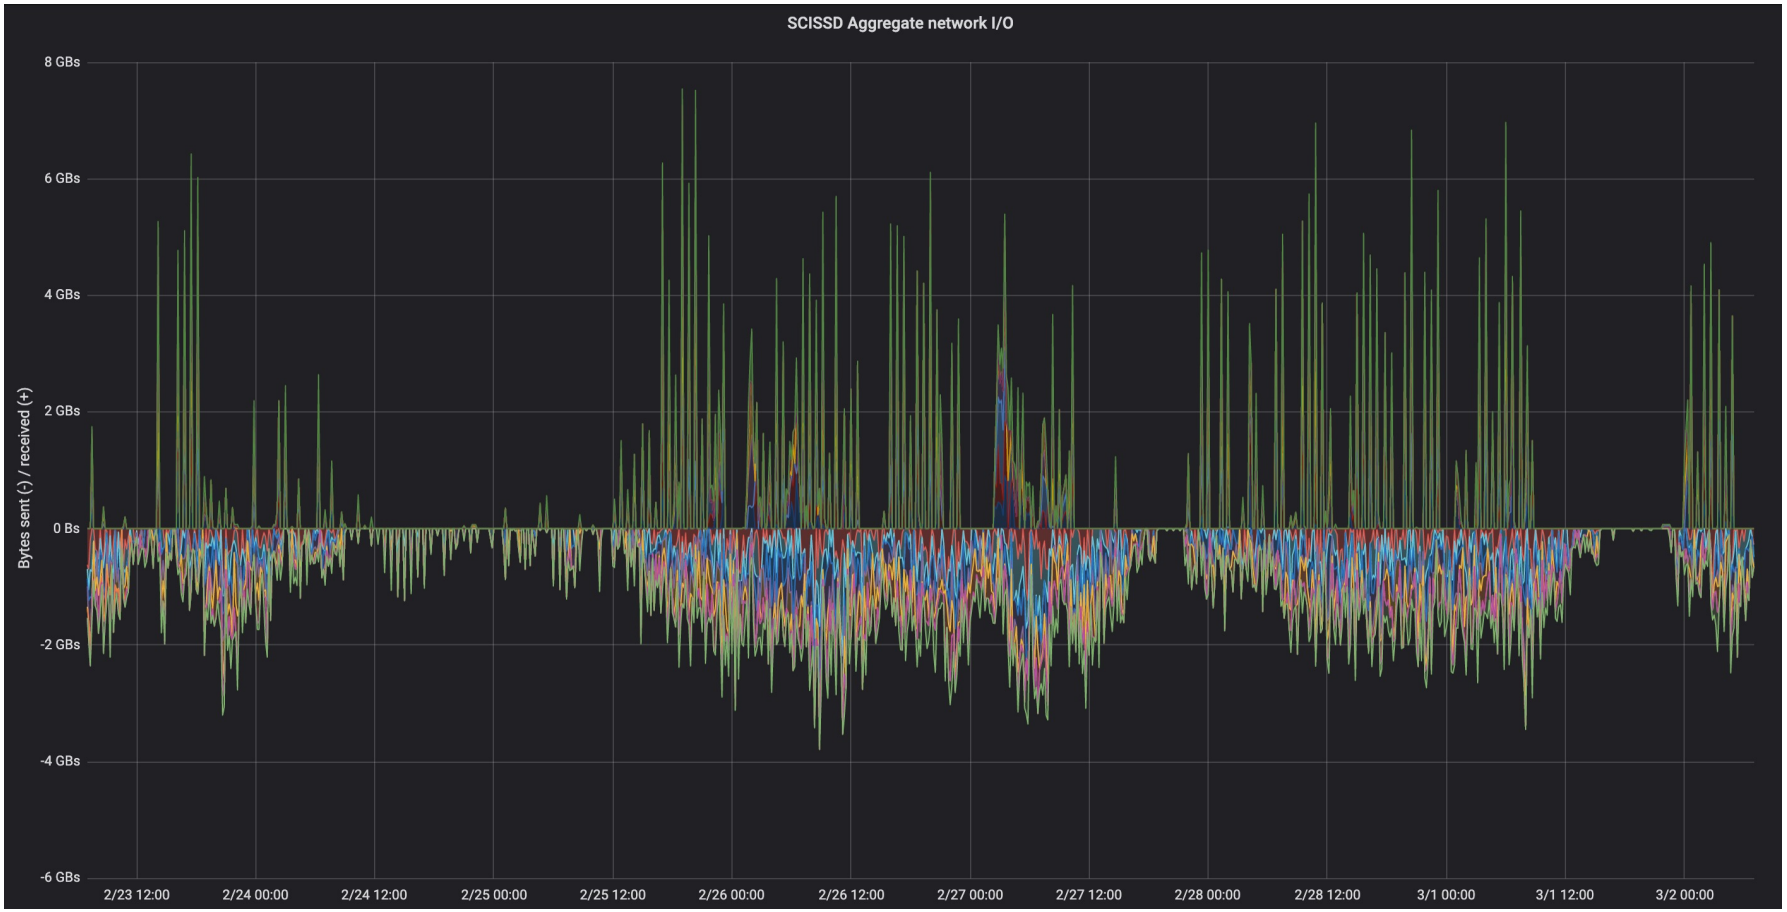

SWIF Input Cache Management Last Month

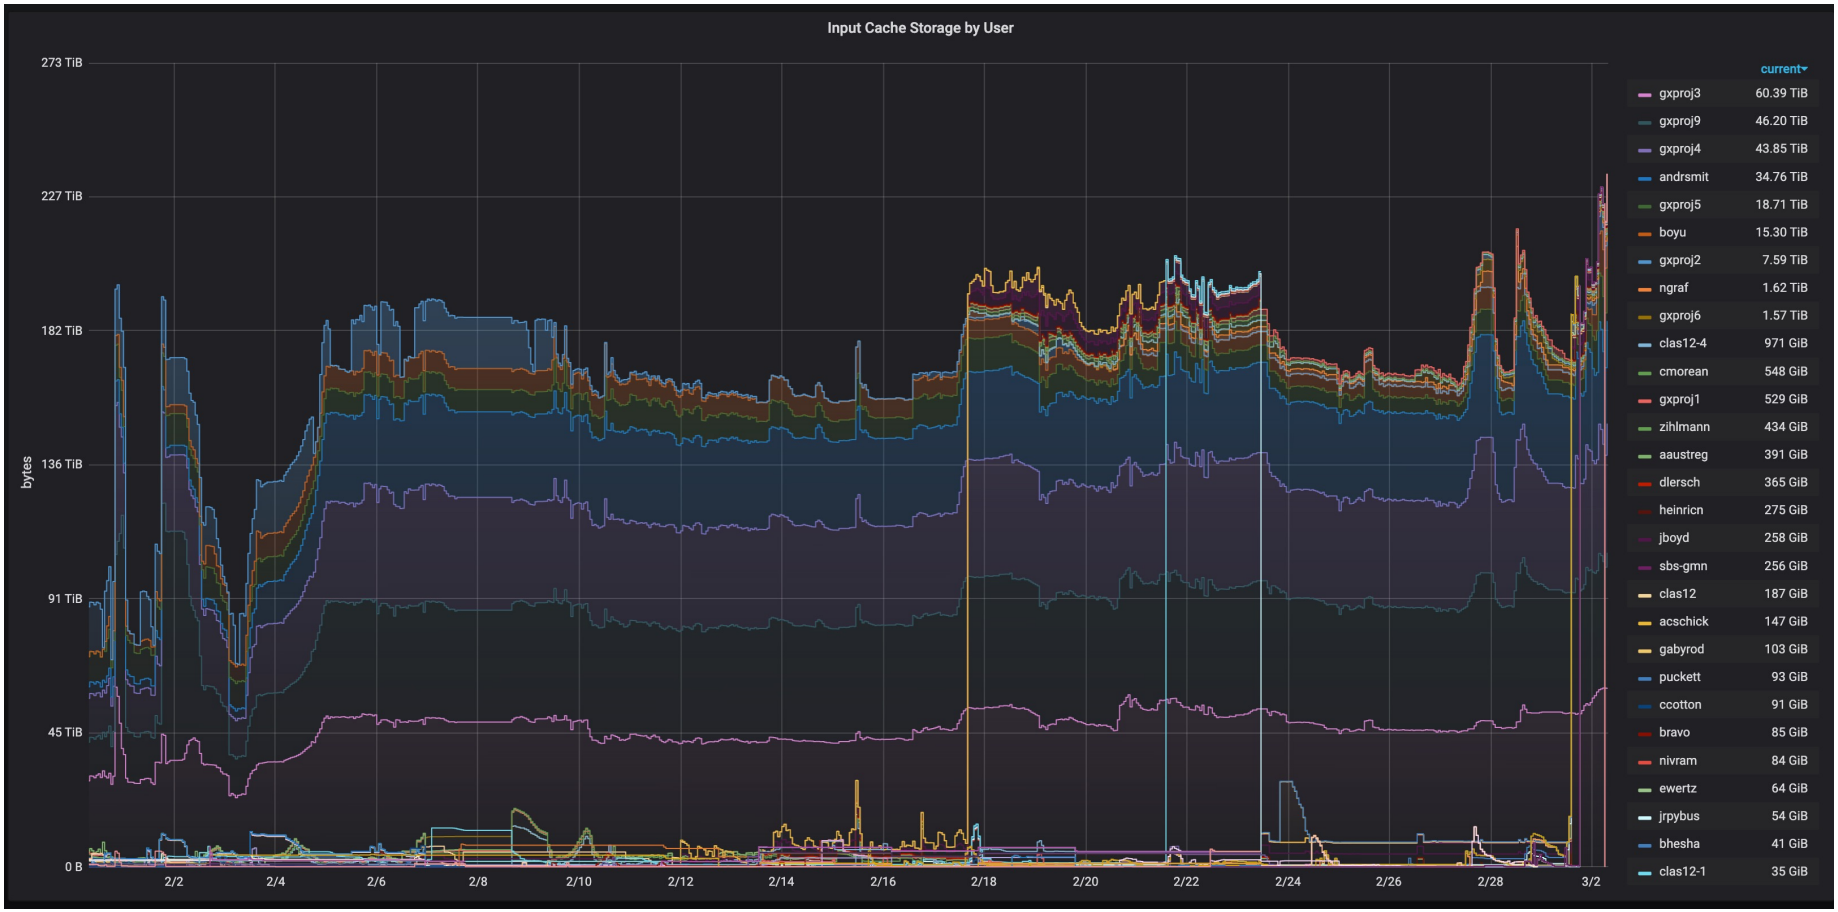

Reminder: MFA access to ifarm – March 21, 2023 cutover

- All Token enrollments have been issued
- Anyone who is having problems with enrollment or token use should contact the helpdesk for assistance
- A final email will go out prior to the change with the timing of the firewall change.
- There are no system changes needed; Maintenance day changes are all network layer.

Testbed Construction

- We are building a testbed for system and software evaluation
- A place to vet risky changes
  - slurm versions
  - new operating systems
  - puppet configuration changes
  - patches
- The farm13 nodes are the worker nodes
- firewalled from production systems
- No shared components
- For system development by multiple tenants.
- Aim is to improve overall quality and availability of production systems

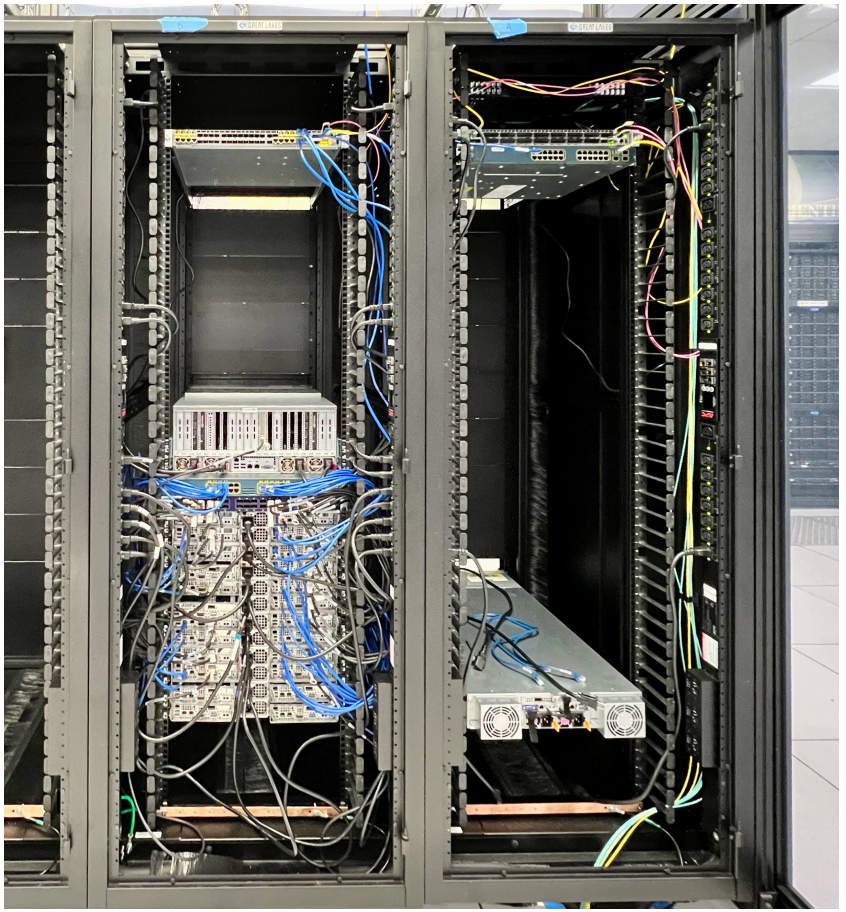

Maintenance Planning

- March
  - First Maintenance after the SAD begins
  - Database Upgrade for Jasmine, SWIF, and related services
    - commands and web pages will be interrupted briefly for patching and reboot
    - Slurm will continue to run
  - Routine patching and reboots of systems
    - most changes are OSG related

- April/May
  - Slurm Upgrade
  - Further tasks to be determined

- On The Horizon
  - Next farm Operating system and worker node configuration
  - CentOS 7 is showing its age; time to move on

AlmaLinux

- After Last Month's discussion we evaluated Alma and confirmed that we can target it as the next Farm Operating System
  - Our base work with Rocky and RHEL translate easily.
  - OSG is building support for Alma 9 now
  - Alma 9 can be the user facing "next OS target"
  - For non-user-facing systems we can remain on RHEL 8 if needed for stability
- As with the transition from 6 --> 7, we imagine that containerization will be the means to support legacy code
- Suggest containerization in general as a strategy for running on multiple sites and reducing JLab-specific paths, or map them to container paths that make path/configuration changes avoid the need to change code.
- We will stand up test systems in coming months. First on the testbed, then in a test queue on the farm for early access.

Rucio Exploration: Update from Ying

- A containerized version of the Rucio (1.30.0) **development/test environment** is set up on a VM rucio2.
- It has 14 individual docker components (rucio server, database servers, 4 xrootd servers, monitor pieces, and more).
- So far, following rucio functions are tested:
  - Create/list/manage a rucio account (user and group account).
  - Create/list/manage DID (data identify), such as file, dataset and container.
  - Use pre-configured res (Rucio Storage Element) for file transfer, includes upload and download files.
  - Manage res limits, file replica rules.