



CSAID operations report

Farrukh Khan on behalf of
Computer Services Infrastructure (CSI) group

January 23, 2024

Public Compute

- Not necessarily a public compute item but we observed SNOW issues on Monday.
 - Emails of tickets to the primary were delayed by a good 2 hours or so.
 - The first primary page also played the automated message with a delay of 15-20 seconds.
 - Were there any known issues with SNOW at the time?

Public Compute

- Intermittent failures while renewing tokens with `osg-token-renewer` continue.
 - Dave Dykstra is following up with CILogon in an email thread to troubleshoot the failures
 - Internally, we have a SOP for the on-call primary to renew tokens manually (which seems to go OK so far)
- GlideinWMS pilots failing to run DUNE and CMS global pool workloads on Alma9 worker nodes is still investigation. OSG pilots are unaffected.
 - Nothing major to report for FermiGrid otherwise. Things are running stable.
- Wilson cluster downtime (RITM1896578) was extended by a day to look at InfiniBand errors.
 - We are currently deliberating opening the cluster up at reduced capacity to the users and troubleshooting the errors in the coming days
 - From CSI perspective, major changes include upgrading Slurm to 23.11, migrating the monitoring over to landscape and restricting worker node SSH access for users. More detailed list of changes made is in the ticket

All things CMS

- GlideinWMS pilots failing to run DUNE and CMS global pool workloads on Alma9 worker nodes is still investigation. OSG pilots are unaffected.
 - Nothing major to report otherwise for CMS Tier1.
- My original plan was to upgraded lpcschedd3 and lpcschedd4 to Alma9 this week, but they didn't drain in time.
 - A user has high memory jobs (5 GB) in the queue. Since their priority takes a hit for taking up three slots, the jobs are running slowly
 - I will now look to upgrade the schedds next week
- Nothing major to report from CMS SI and factory side.