

## Hardware/Sysadmin/Puppet

### January 23, 2024

- Procurement:
  - **Update:** CMS and FermiGrid worker nodes (79 CMS, 21 FermiGrid) ready for production, other than a few problematic nodes, in progress.
  - CMS tape dCache servers (NVMe SSD storage) getting installed.
  - CMS and IF disk storage: All servers at FCC and getting installed. (This will take a while; some are waiting for network parts.)
  - GPU servers are here, getting installed into racks, but waiting for network parts.
  - **Update:** About half of general CMS servers are here, getting installed into racks.
- Preparing to contact experiments and other stakeholders with proposed dates for upgrades from SL7 to EL9 (or EL8 in a few cases).
- HPC:
  - Wilson Cluster downtime nearly complete (some cleanup still proceeding)
- Upgrading from SL7
  - Status:

OS version	SL7	EL8	EL9
# nodes	2012	674	511

## Other Activities

- Proxmox/RHV/GPCF
  - rsyslog has been configured on all nodes to send logs to clogger.
  - ocsinventory installed on all servers.
    - Everything appears to be working from the client standpoint, but SLAM says they are not receiving the data, so we are still trying to make this work.
  - Updated DNS entries:
    - proxmox.fnal.gov -> ssi-proxmox.fnal.gov
    - proxmox-report.fnal.gov -> ssi-proxmox-report.fnal.gov
- Miscellaneous
  - Two people out last week, so mostly just trying to keep up.
  - Finishing Wilson downtime. Our work is mostly done, but we see some possible InfiniBand issues on a small number of systems. Investigating.
  - dCache kernel updates/reboots on Wednesday performed with no issues.
  - NAS to Ceph migrations on Wednesday, gm2 and Nova completed (but not finished until later in the evening due to the rsync completion time). (Which is mostly because experiments put **millions** of files in a directory, for inexplicable reasons.)
  - Did some work with mu2egpvm06 to help the experiment do more testing of the Ceph mounts.
  - Interesting problem with novadaq-near-mgr-04 spawning a DAQ process with 100,000+ threads... appears to be a user software issue.
  - Working through some IPMI issues on new worker nodes.
  - Assisted SSA with some credential issues related to a couple of our Puppet git repositories. (A side effect of issues with their Cobbler server.)
- Installs/Moves/Retirements
  - All Wilson nodes have been reinstalled as part of the downtime.
  - csitest01/02 reinstalled as esworker130/131 and handed over.
  - Reinstalled new worker nodes cmswn6001-6079 and fnpc23101-23121 to fix a partitioning issue.
- Normal operations: 58 logged interventions last 3 weeks—replacing disks, rebooting hung systems, fixing corrupt RPM databases, retiring systems, fixing xinetd issues, clearing caches, fixing hung NFS mounts, fixing Puppet errors, etc.


## HPC (High Performance & Parallel Computing)

- Wilson Cluster upgrade to EL8 nearly complete.




## Incidents

SCF - Weekly Incident Volume by Assignment Group



## Requests



# Ticket Trend

