CMS Computing Resources

August 16, 2021

Jeremiah Mans

Where did it come from?

- Most of the resources were originally purchased using post-2008 “stimulus funds” (the ARRA act)
  - As a result, the nodes are not new
  - On the “positive” side, the rate of increase in CPU power over the last decade hasn’t been that high…
- Storage has been upgraded and replaced repeatedly through departmental and some limited DOE funding over the last decade

Nodes

- Most nodes are identical:
  - Two physical CPUs with four cores (and hyperthreading), resulting in 16 “CPU” per nodes. Intel Xeon L5520@2.27GHz
  - 24 GB of RAM
- Login nodes: zebra01.spa.umn.edu, zebra02, zebra03, zebra04
- Compute nodes: scorpion1.spa.umn.edu – scorpion48
  - Some currently non-operational: 7, 13, 18, 48
- Storage nodes:
  - whybee1 – JBOD server for /local/cms/user drives
  - whybee0 was historically used for a physical RAID array, which has been decommissioned for some time
  - whybee2 exists, with no apparent function that I’ve identified...

Storage

- Home directories from CSE IT over NFS
  - Migrated from physics home directories several years ago
  - Quota/size is quite limited
- /local/cms/user JBOD server
  - ‘Tray’ of drives, with some built-in redundency. Generally used as an extended home directory by users
- HADOOP/hdfs distributed filesystem
  - Drives located in all the scorpions (and other servers in the other HEP clusters)
  - Files are broken into 64 MB chunks and the chunks are distributed across the nodes of the cluster. Chunks are duplicated (typically three times) for redundancy and parallel I/O efficiency
    - Read loads during I/O bound jobs are then distributed across the cluster
  - Total drive count ~110, with drive sizes between 2 TB and 5 TB. Total storage capacity (with three replicas) is O(150T)
    - More total size, those drives are actually on nodes outside the CMS cluster...

The Future… maybe

- Historically, we had the ability to receive direct deliveries of full datasets through Phedex
  - This has been dead for some time, first due to software issues, then because Phedex is being shut down/replaced with rucio
- We have long desired to have the ability for CRAB jobs to deliver the output files from a job directly to Minnesota
- The CSE IT team has been working VERY SLOWLY on this for O(3 years), to get the newer OSG software installed.