

Experimental Computing Laboratory

July 27, 2023

Advanced Computing Systems
Research Section

Jeffrey Vetter, Section Head

Steve Moulton, Systems Engineer

ORNL is managed by UT-Battelle LLC
for the US Department of Energy



U.S. DEPARTMENT OF
ENERGY

Systems Status

- All new heterogeneous-support servers deployed
 - Milan0: 2 * Nvidia A100 GPUs. Available.
 - Milan1, Milan2: Groq AI Accelerators
 - These cards can be paired – (re)deployment driven by researcher needs
 - No current usage
 - Milan 3 difficulties resolved
 - All Milans available as “Milan” slurm partition
 - At least one project using slurm to dispatch apptainer jobs

System Outage August 15, 2023

- Deploy ZFS file server as primary file service
 - User-accessible snapshots performed hourly for user files
 - Hourly kept 48 hours
 - Daily kept 14 days
 - Weekly kept 4 weeks.
 - /noback will no longer be necessary. I will happily help users fold their /noback files into their home file tree.
 - Still please be reasonable about disk usage – it isn't infinite. Clean up CI artifacts, etc.

System Outage August 15, 2023 (2)

- LDAP proxy being upgraded
 - Longer session authentication keys required in latest Redhat releases
 - Should be transparent
- Moving hypervisor services to Rocky9
 - Will update/move infrastructure VMs
 - No change to user VMs
- Lots of OS updates need to be applied.

This is likely to take a full day.

Pending installation

- Xilinx U280 FPGA server-class card
- Next Silicon Maverick card
- Nvidia Jetson AGX Orins
- Nvidia Jetson Nanos – have several; might make nice minicluster?
- Achronix VectorPath S7t-VG6
- Nvidia Xavier board reflash
- Nvidia Bluefield 2 DPU
- Snapdragons – not installed yet.

Accounts and Other Management Activity

- All accounts requested in place
- Siemens project in place, awaiting account requests
- Ft.ornl.gov shut down and deleted.
 - Internal documentation on devdocs.ornl.gov
- Excl.ornl.gov is becoming in-house gitlab-served resource (but still available externally).
 - AWS has been problematic with small (1/10 server) management. Looking forward to move.

FY 2023 purchases *Now is the time*

- MI300 – based system
- Intel-Intel system (stampede 3-like)?
- Multiple GPU heterogeneous system for AI and performance research
 - MI250, H100 (Narasinga) (both for AI and Siemens work)
 - Heavily using Cades Cloud system – could disappear at any time
- Nvidia Hopper
- Ideas, requests, dreams?
- Crusher-like resource (leconte-like) -- Pedro

*There is a tide in the affairs of men. Which,
taken at the flood, leads on to fortune*

Questions/Discussion?

- TRC is coming
 - But so is Christmas
 - Appears that construction activity safety issues continue
- Query Users on outage timing
- Siemens discussion?