

# T2 Hardware Refresh

Pete Eby  
[ebypi@ornl.gov](mailto:ebypi@ornl.gov)

Steve Moulton  
[moultonsa@ornl.gov](mailto:moultonsa@ornl.gov)

Slides:  
Dec 7th 2022 @ ORNL

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

# Worker Nodes – New Hardware

- 4 Silicon Mechanics Rackform R4420.v8 quad chasses.  
Each node:
  - 2 Intel Gold 5318Y (24 Core, HT, 2.1 GHz)
  - 256 GB Memory
  - Broadcom NetXtreme 25 Gb adapter – 2x SFP28
  - 240G & 960G SSDs
- hepspec benchmarks and projections [here](#)
- Racking orders in progress

# New Production EOS

- FSTs: Two Storform A506
  - AMD EPYC 7543P Processor 32 Core 2.8 Ghz
  - 256 GB Memory
  - 2 \* 240 GB SATA SSD
  - Broadcom HBA 9500-8i 8 port SAS3/SATA/NVMe
  - Broadcom NetXtreme 25 GB (2x SFP28 ports)
- JBODs: Two Storform DS5-84 (Seagate Exos E 5U84)
  - 1.51 PB Raw storage – 84 X 18 TB Drives
- TBD: (n)fst.services per node with (t)fsids per service
- Option to upgrade to 40Gbit links to ORNL border router

# Single Disk vs Parity Options

Method	Cons	Pros
RAIN: Read rate	50% read rate reduction?	???
RAIN: w/o EC rebuild time?	Need clarified	???
RAIN 6	mixed bag (capacity, read/write rates, etc.)	Data integrity
Failure / availability domain	fsid size (huge with 18T disks)	disk size
Software/Hardware RAID ●	Perf envelope (xrootd threads / disk TB) / fsid.size fte overhead questionable	Just works.... Huge fsids, perf envelope
EC	Lacking full documentation	Presents pool of disks so no huge fsids?
Use of gateway (with RAIN, EC)		

# Replacement MGMs

- 2 Silicon Mechanics Rackform R335.v8
  - Xeon Silver 4310 2.1Ghz, 12 core (HT)
  - 256 GB Memory
  - 4 \* 480 GB SSDs
- QuarkDB members

# New Operational Storage – ZFS/NFS Server

- Silicon Mechanics Storform R5032.v8
  - Xeon 6342 2.8 GHz 24 Core (HT)
  - 384 GB Memory
  - 2 \* 480 GB SSDs
  - 100GB Optane SSD
  - 6 \* 10TB SAS 3.0I 12.0Gb/s
  - Broadcom NetXtreme 25 GB 2x SFP28
- Will also serve as new hypervisor
  - VMs: vobox, osg node, slurm, login node, monitoring, etc.