



ALICE

ALICE-USA

T2 and AF Computing @ LBNL

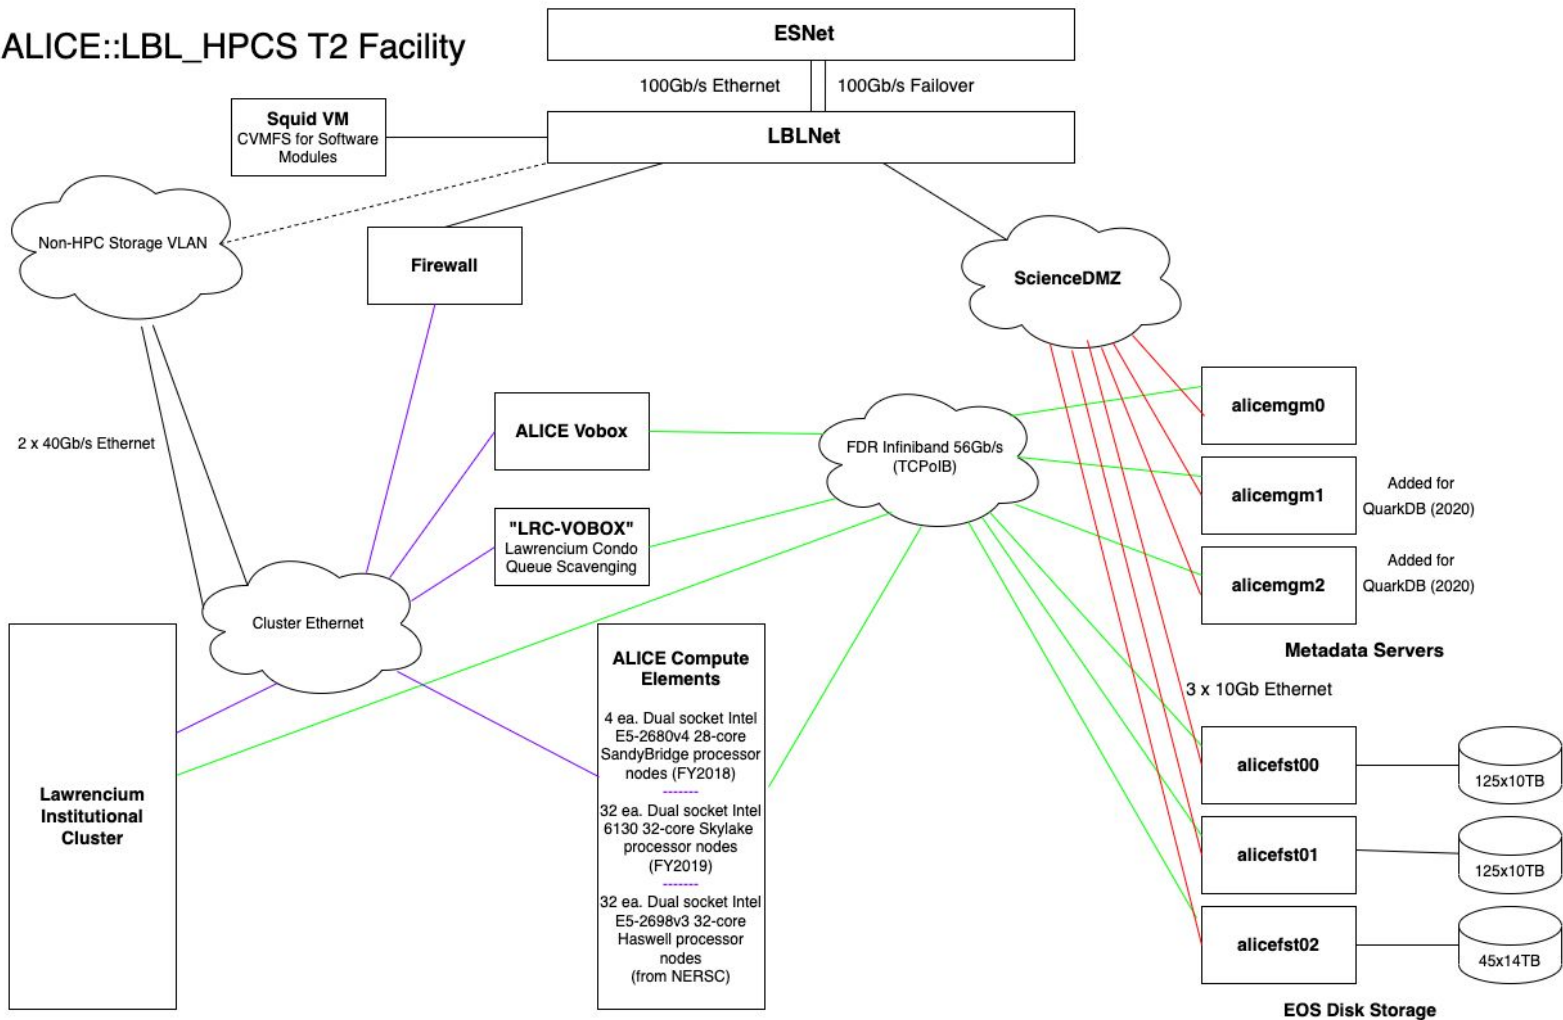
ALICE-USA Annual Meeting
May 15-17, 2023

John White LBNL

Lawrencium Overview

- **Lawrencium:** Institutionally-supported HPC Linux Cluster
 - ~1,300 nodes/ ~25,000 processor cores
 - Intel processor architecture
 - 4.3PB Lustre parallel file system for Scratch supporting condo-style buy-in
 - Scientific Linux 7 OS, SLURM job scheduler, Singularity containers, Open OnDemand
 - Nvidia GPUs for machine learning
- Three models of access
 - PAYG \$0.01/core-hr
 - Condo Cluster
 - Dedicated cluster support
- PCA - PI Computing Allowance
 - Free - 300K SUs per year
 - Can be shared with your staff or pooled together with other PIs

ALICE::LBL_HPCS T2 Facility



ALICE Compute Elements

4 ea. Dual socket Intel E5-2680v4 28-core SandyBridge processor nodes (FY2018)

32 ea. Dual socket Intel 6130 32-core Skylake processor nodes (FY2019)

32 ea. Dual socket Intel E5-2698v3 32-core Haswell processor nodes (from NERSC)

EOS Disk Storage

Storage Futures

- New 2x40Gb routed link from LBLNet to HPC network
 - Currently providing a 6PB CEPH platform
 - S3/RADOS
 - Cephfs
 - NFS
 - New VAST offering
 - Home, group, software
 - Formulating condo pricing
- Scratch File system
 - DDN SFA 18K
 - 4 controllers, Virtualized servers
 - First 64MB of every file on NVME Flash (Progressive File Layout)
 - Rest on Spindle
 - 440TB Buy-in available (non-purged space)

Cluster Futures

- Upgrades in the (not near) future
 - Warewulf 4
 - Not terribly important to this audience BUT
 - There will be a few fits and starts
 - Rocky 8
 - Mellanox OFED LTS 4.9 requirement to support old hardware
 - SPACK
 - Bare-bones images already running on ~110 nodes
 - Just let us know when for alice0
- LHCOne done
 - Migrating to new flexible routing layout on Saturday
- Migrating from flat files to LDAP

EOS (ALICE and former STAR)

- 4 MGM nodes
 - HTTP interface recently re-enabled
 - Need a downtime for migration to new Datacenter
- 4 FST nodes
 - 6 JBODs split among all 4
 - Populating all remaining slots, Splitting for AF usage (order showing up on Friday)
 - Near-infinite expandability (1024 drives per server)
 - Need a downtime for migration to new Datacenter (maybe)

Purchases

- 2 Chassis (8 nodes) for ALICE
 - Air Cooled
 - Installed recently, being configured.
- 92 Drives (16TB) purchased for Analysis Facility Shuffle