EOS AT THE VIENNA T2

EOS Workshop 2022

Erich BirngruberericH.birngruber@gmi.oeaw.ac.at
Vienna Biocenter

https://www.clip.science/

CLIP

VIENNA T2: COMPUTE

- Project CLIP (CLoud Infrastructure Project), OpenStack based
- VBC Compute center for HTC workloads (Biology, physics, various others)
- Job scheduler is SLURM (+ HT-Condor CE)
- 200 compute nodes in pools, 250TB scratch filesystem, 100GBit/s Ethernet interconnect
  - CPU / compute nodes (total ca 8000c)
  - High-memory nodes (2TB)
  - GPU nodes 112 GPUs total
- Shared “public” partitions for generic workload
  - Separate partition for grid workload: 40 nodes, ca 1500 cores
  - Multiple experiments workload: Alice, Belle2, CMS
  - QoS + Fairshare for balancing between experiments

CLIP

VIENNA T2: EOS STORAGE

- Deployed early 2020, production mid-2020 (EOS v4.7.x)
- Today: EOS 4.8.78, QuarkDB 0.4.2
- RAIN 5+2 layout
- Converged instance for 3 experiments
  - Alice
  - Belle2
  - CMS
- Extra services
  - FUSEx mount on SLURM computes
  - For other physics data, and easy access

CLIP

VIENNA T2: EOS HARDWARE

- On baremetal commodity servers, 3 PB total raw capacity
- 3x JBOD with 84 HDDs @ 12TB as building blocks
- EOS installation
  - 3 mgm with quarkdb (on NVMe SSD)
  - 9 FSTs
    - 28 x 12TB HDDs attached (dual SAS-3 12 Gbit/s)
    - 3 FSTs per JBOD
    - 2x 40Gbit/s Ethernet
- Campus 40 Gbit/s uplink (recent upgrade from 10 Gbit/s)
- Current usage: ca 60+% used (in 2021: 40%), quarkdb ca 56GB (23mio files, 1.38 mio dirs)

CLIP

VIENNA T2: EOS MONITORING

- EOS builtin monitoring output is machine readable
- Prometheus exporter based on: `eos <cmd> -m`
  - Written to text file, shipped by node_exporter
  - Combined with other host metrics
- Metrics
  - IO, fsck, quota
- Alerting
  - General Linux OS state: CPU, Mem, disk
  - EOS Node status
  - EOS FS status
- Will look into upstream EOS Prometheus exporter

DISK FAILURE MYSTERY

- Starting in 2021: unexpectedly high disk failures rates were observed
- 252 total disks: 33 failures total (31 in 2021 alone): **>12% failure rate**
- Individual disk failures were tracked, disk serial
  - enclosure slot
  - FST server
  - Linux device

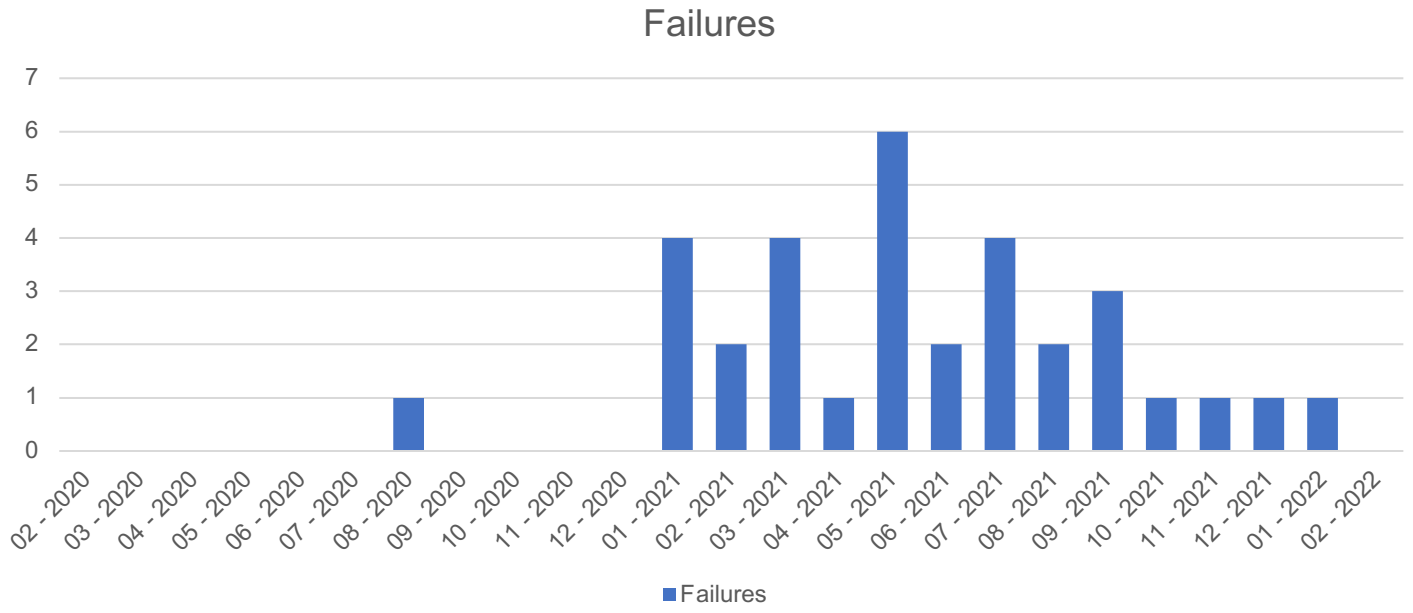

CLIP

DISK FAILURE MYSTERY: SOLVED

- failed disks all over (all servers, enclosures, slots)
- A bad model? Not according to Backblaze stats: (ST12000NM0038) https://www.backblaze.com/blog/backblaze-drive-stats-for-2021/
- Keep on looking

Some high capacity Seagate hard drives may fail prematurely or report PFA - Lenovo ThinkSystem

- https://support.lenovo.com/in/en/solutions/ht511631-some-high-capacity-seagate-hard-drives-may-fail-prematurely-or-report-pfa-lenovo-thinksystem (tl;dr bad firmware LCA6)
- Firmware update on all disks, drastically reduced failures since.

CLIP

VIENNA T2: OPERATIONAL TASKS 2021

- ✅ fsck improvements

  initialy high number of fsck report, drastically reduced after improvements to reporting

  some fsck still reported on deleted files (cosmetic) or when replacing disks (FS)
- ✅ reported accounting report discrepancy (used for experiment accounting)
- ✅ aided implementation of fs status on Linux multipath devices (eos fs ls)
- ✅ davs:// HTTP protocol access enabled for all experiments

CLIP

VIENNA T2: SUMMARY / OUTLOOK

- No plans yet for EOS5, who upgraded already (upgrade path)?
- Are there disk (type, failure) stats at CERN (from the 52k disks)?
- Fuse(x) is important our users: groups without VO, local cluster access only

- Whishlist
  - EOS client auth vs Alice token auth for (fuse) compute nodes (quota, find, etc)
  - Built-in Prometheus exporter

- Excited about next data taking period, Run 3 (our first one!)
- Thanks to the active community: Elvin, Andreas, EOS forums!

VIENNA T2: CODE

- https://www.clip.science
- https://github.com/CLIP-HPC/clip-grid-eos (private repo)
  - Ansible playbooks
  - Prometheus exporter
  - Other config details
  - (ask us for collab invitation Github @ebirn, @timeu)