

EOS usage and upgrade in Cape Town

CERN, CHPC

Sean Murray,

February 4, 2020

CERN/CHPC/UCT

Cape Town



Cape Town



Old/Current storage

We currently have a very badly designed storage configuration.

- We lost 100TB of data due to raid5 with 30 disks.
- SAS card limits each instance to c. 4Gbps.
- Reworked to 11 raid6 arrays after the above loss.

New Storage

Some History

- Been wanting to change this storage since 2016.
- Money approved in 2018, but storage to sit in Pretoria.
- Alocated as part of a ceph cluster
- Cape Town to Pretoria, is Geneva to Sofia.
- Lesson : If one irritates people enough they just give you money to get you to go away!

2 designs

- RAIN on 24 disk servers.
- ZFS 60 disk machines.

replace core networking

- currently have 10G Dell Force 10 in HA 48 port
- singular direct 10G uplink to NREN backbone
- 2 new Mellanox MSD2410-CB2R 48 port 25Bge/8/100GbE
- MGM will host disks as well for cost reasons.
- ATLAS and ALICE "completely" separate.

24 disk servers

- 2x 16 Core
- 64 GB RAM
- 2x480GB STA SSD
- 2x 1TB NVME
- dual 25GbE
- 24x16TB drives
- other stuff required for a supermicro box

60 disk servers

- 2x 16 Core
- 64 GB RAM
- 2x480GB STA SSD
- 2x 1TB NVME
- dual 25GbE
- 60x16TB drives
- other stuff required for a supermicro box.

The above 2 specs were more less done 18 months ago.

- 24 disk boxes allow us RAIN
- 60 disk boxes give us 4 zfs zones 15 disks each.

Acknowledgements

My colleague Lepeke Phukungoane from CHPC in Cape Town.

Opinions, comment welcome.