Echo Operations

Rob Appleyard

Current Ceph team: Tom Byrne, Vijay Rajabathar, Joshua Kitenge, Aidan Mccomb

Echo – Brief Introduction

- Echo is the Tier 1's disk storage system
  - Ceph
  - Currently 42 usable PB (58 raw)
  - 8+3 Erasure Coding
- External access via dedicated XrootD/GridFTP/S3 gateway nodes
- Internal access via an XrootD gateway on every worker node

Image from www.publicdomainpictures.net, available under https://creativecommons.org/publicdomain/zero/1.0/

Science and Technology Facilities Council

Storage Team Changes Since Last Summer

- Left
  - Marcello Armand Pilon (another post in SCD)
  - Matt Heath (left STFC)
  - Morgan Robinson (placement student)
  - Kyle Pidgeon (Grad rotation)
- Arrived
  - Vijay Rajabathar (Ceph team)
  - Joshua Kitenge (Grad rotation)
  - Maha Agilandamurthy (Antares team)
  - Jyothish Thomas (XrootD development)
- Still here
  - Rob Appleyard (Team lead)
  - Tom Byrne (Service development, wizard)
  - Aidan McComb (Disk operations, CASTOR)
  - George Patargias (Tape operations and development, ALICE)

Hardware changeover

May 2022:

- Storage generations from 2015-2018 in production, plus one rack (6.1PB) of 2020 on the legacy network.
- Cluster 73% full, fullest OSD 83% full
- Deployment of 2019 and 2020 hardware delayed in order to ensure that they were placed on the new network.
- Disks in XMA15 hosts increasingly flaky
- Several Dell16 hosts supported by replacement switches placed on top of the racks following a problematic network intervention on the 31st March.

Storage Hardware Changeover Sequence

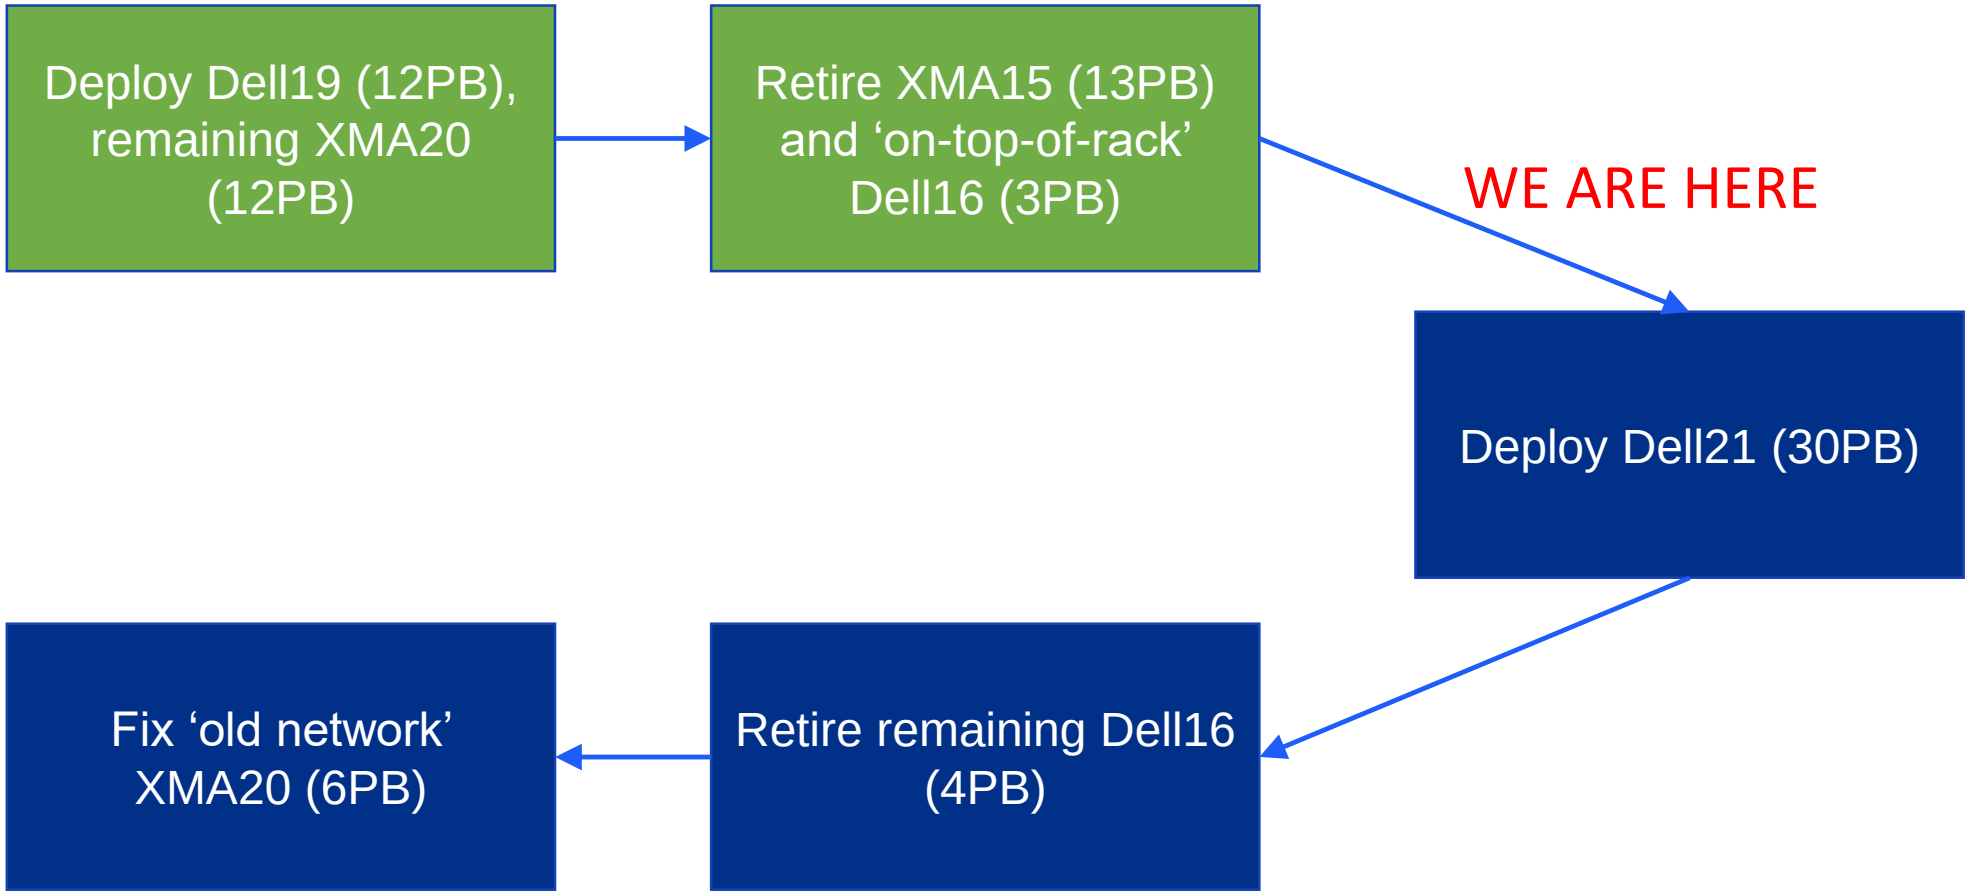

Science and Technology Facilities Council

Hardware Changeover (1)

- New hardware deployment was constrained by the network link between the Legacy Tier 1 core and the Superspine (4*100GB)
  - Data migrating from old network to new
  - Agreed max safe usage 65% of network capacity
- Weighted up each generation by 10% on alternate days
  - Limited point load on each host
- Took 29 working days

Hardware Changeover (2)

- Two constraints on hardware removal
  - Once again, network link between the Legacy Tier 1 core and the Superspine
  - CPU and disk I/O capacity of old hosts
    - Heavy load on our oldest hardware
    - Tried weighting down by 10% at a time in mid-July, result was slow requests seen by ATLAS and LHCb
    - Underlying cause was this happening to a 100Gb link (rack-level switch to core)

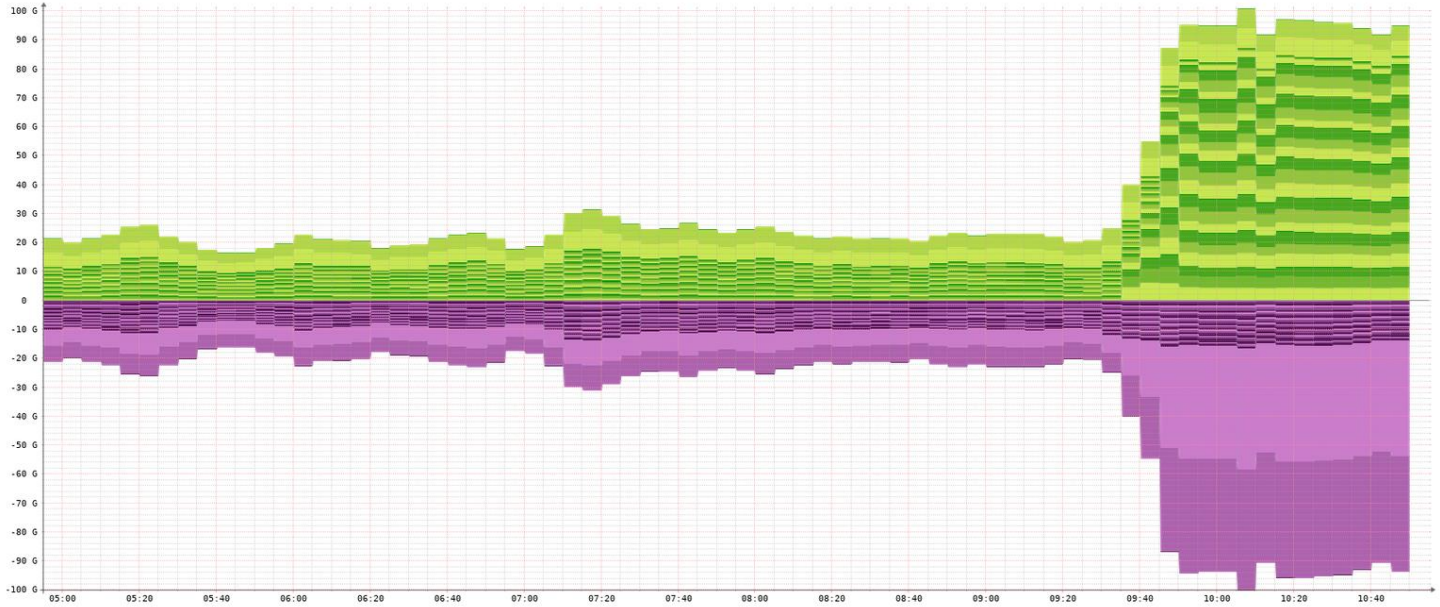

Science and Technology Facilities Council

Hardware Changeover (3)

- New plan: Weight down all hosts in generation by 5%/day
  - No more performance issues observed
  - Probably due to our host change but may just be different user load
- Status now: All XMA15s and ‘OTOR’ Dell 16s removed from cluster, Dell21 hardware to start shortly.

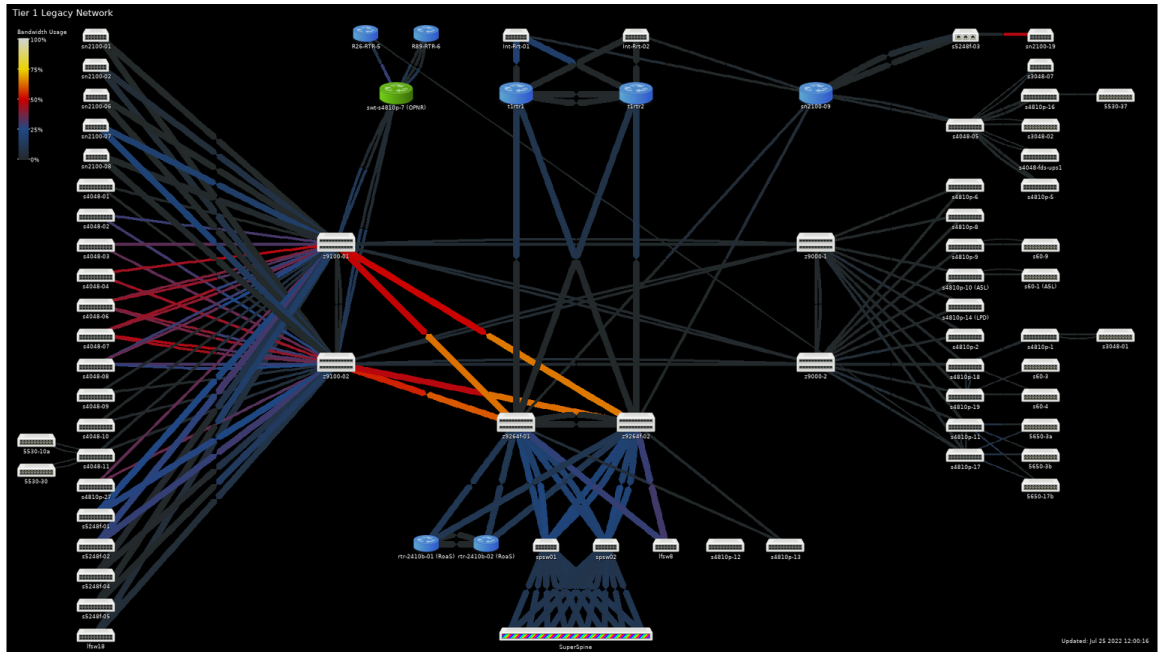

Science and Technology Facilities Council

Data Distribution

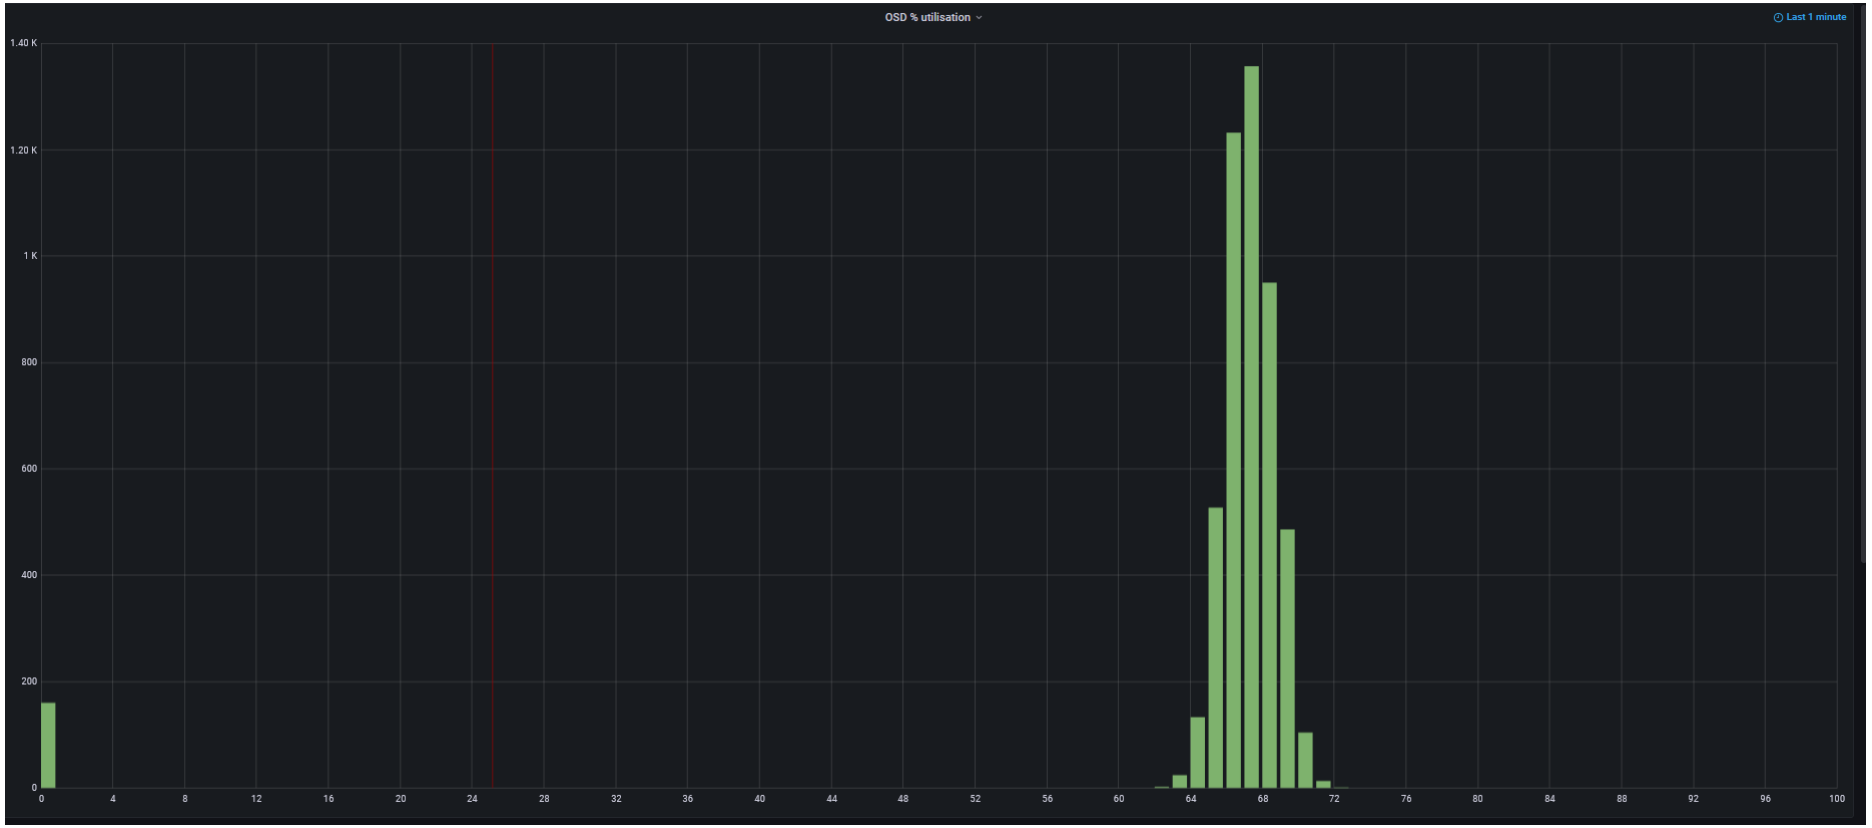

Science and Technology Facilities Council

Gateways and Monitors

- Several ongoing tickets for XrootD gateway issues
  - Jyothish and James Walder have been working very hard on these, and Jyothish has a slot on Friday
- We think the issues seen this weekend were due to a hardware issue on a gateway host (ceph-svc02)
- Changeover of gateway hardware in progress
  - 2015 hardware out of production, 2020 hardware in
- Looking to replace our round-robin redirection with a site XrootD redirector
- Planning to replace three very old (2015) monitor nodes with 1*2019 and 2*2020

VOMS Authentication

- We are behind where we want to be on rolling out VOMS-based authentication to all of Echo
  - Matt Heath completed this for our external gateways in May
  - …and then left ☹
- Vijay has picked up the job of migrating our WN gateways
  - Ongoing!

Agenda for the Next 6-9 Months

- In no particular order…
  - Complete the hardware transition plan
  - Upgrade Ceph from Nautilus to Octopus
  - Upgrade all hosts to a more-modern kernel
  - Rollout VOMS authentication to all worker nodes
  - Transition to rack-level failure domains

A Mimic Octopus.

Ceph Failure Domains

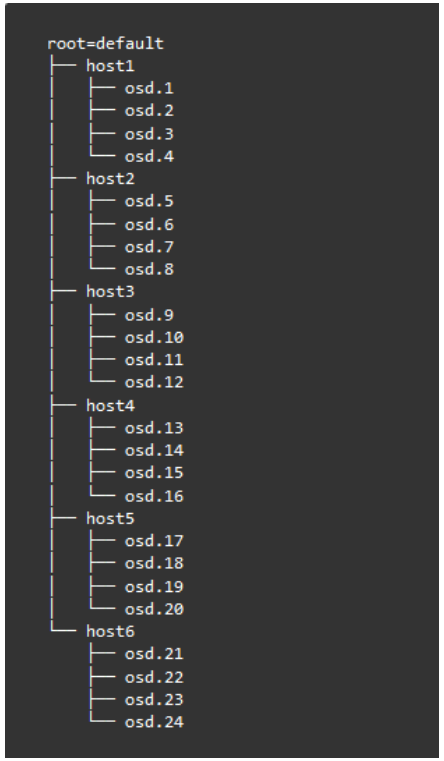

A CRUSH map with only host information

- Ceph's failure domain represents the smallest hierarchical branch in the architecture that cannot appear twice in a placement group
  - If the failure domain is set at host-level, every shard of an object will be placed on a different host.
  - If the failure domain is set at rack-level every shard of an object will be placed on a different host.

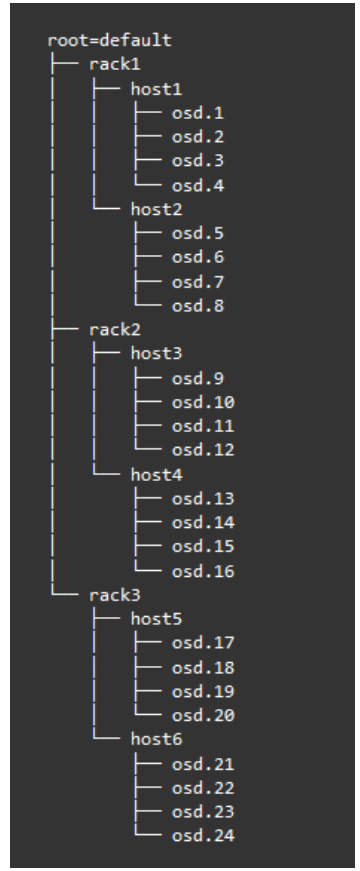

A CRUSH map with host and rack level information

Science and Technology Facilities Council

Move to Rack-Level Failure Domains

- Echo's failure domain is currently set to host-level
  - Choice made during cluster creation
    - 8+3 EC implies 11 racks as a bare minimum
    - Added complexity during a challenging period
- We think shifting to rack-level will bring several advantages
  - Much faster and simpler full-cluster reboots
  - Seamless TOR switch interventions
  - Simpler architecture
- How to get there?
  - Work in progress
  - Will likely require two full internal data reshuffles

Non-WLCG Ceph Services

- ‘Deneb’
  - 5.2PB HDD providing 8+3 EC CephFS
  - More hardware on the way
- ‘Sirius’
  - 700TB NVMe block storage for OpenStack cloud
  - Tight on space due to proliferation user files, prompting the creation of…
- ‘Arided’ (to be commissioned shortly)
  - Josh Kitenge’s grad project
  - SSD cluster providing 460PB (usable) 3* replicated CephFS storage mountable from the SCD Cloud

Some block storage

“Drawers and drawers of LEGO, CC 2.0 licence, via Roo Reynolds’s FlickR

Science and Technology Facilities Council

Final note

- I’m leading a project to replace SCD’s current system backup solution
- Very keen to hear from other sites…
  - What are your use cases?
  - What technologies do you use?

Science and Technology Facilities Council

Conclusions

- Big hardware/architecture changeover is going well
- Long sequence of improvements to come
- Ceph running is very stable
  - Gateway issues under heavy development, considerable progress made

Science and Technology Facilities Council

Thank you

Facebook: Science and Technology Facilities Council

Twitter:@STFC_matters

YouTube: Science and Technology Facilities Council