




Science and
Technology
Facilities Council

Tier-1 in GridPP7

Alastair Dewhurst




Introduction

- DC24 Overview
 - Details will be covered by many others
- Antares
- Echo
 - Ceph – I will present slides made by Rob.
 - XRootD – Jyothish Thomas
- Batch
- Network – James Adams
- Projects
 - VMWare
 - Victoria Metrics
 - Procurements

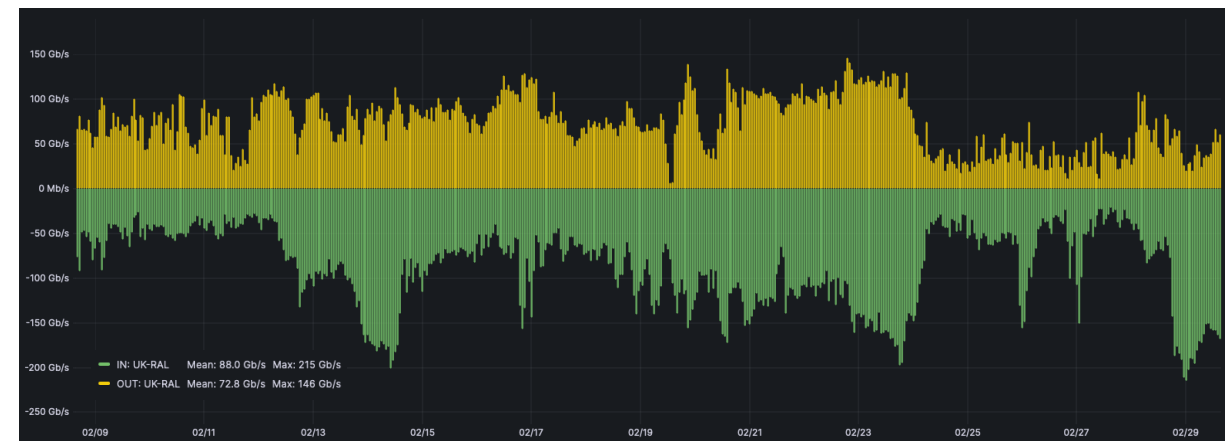
DC24 - Overview

- Overall DC24 was very successful for RAL.
 - Significant network problems at the start of challenge resulted in low initial throughput.
- At the end, throughput was 2 – 3 times higher than we ever saw before:
 - We found a new set of bottlenecks well above Run 3 normal load that can be worked on over the next few years.
 - Fantastic effort from the Tier-1 to debug, fix/optimize in real time.

WLCG DC24 Monitoring



Site Link Monitoring

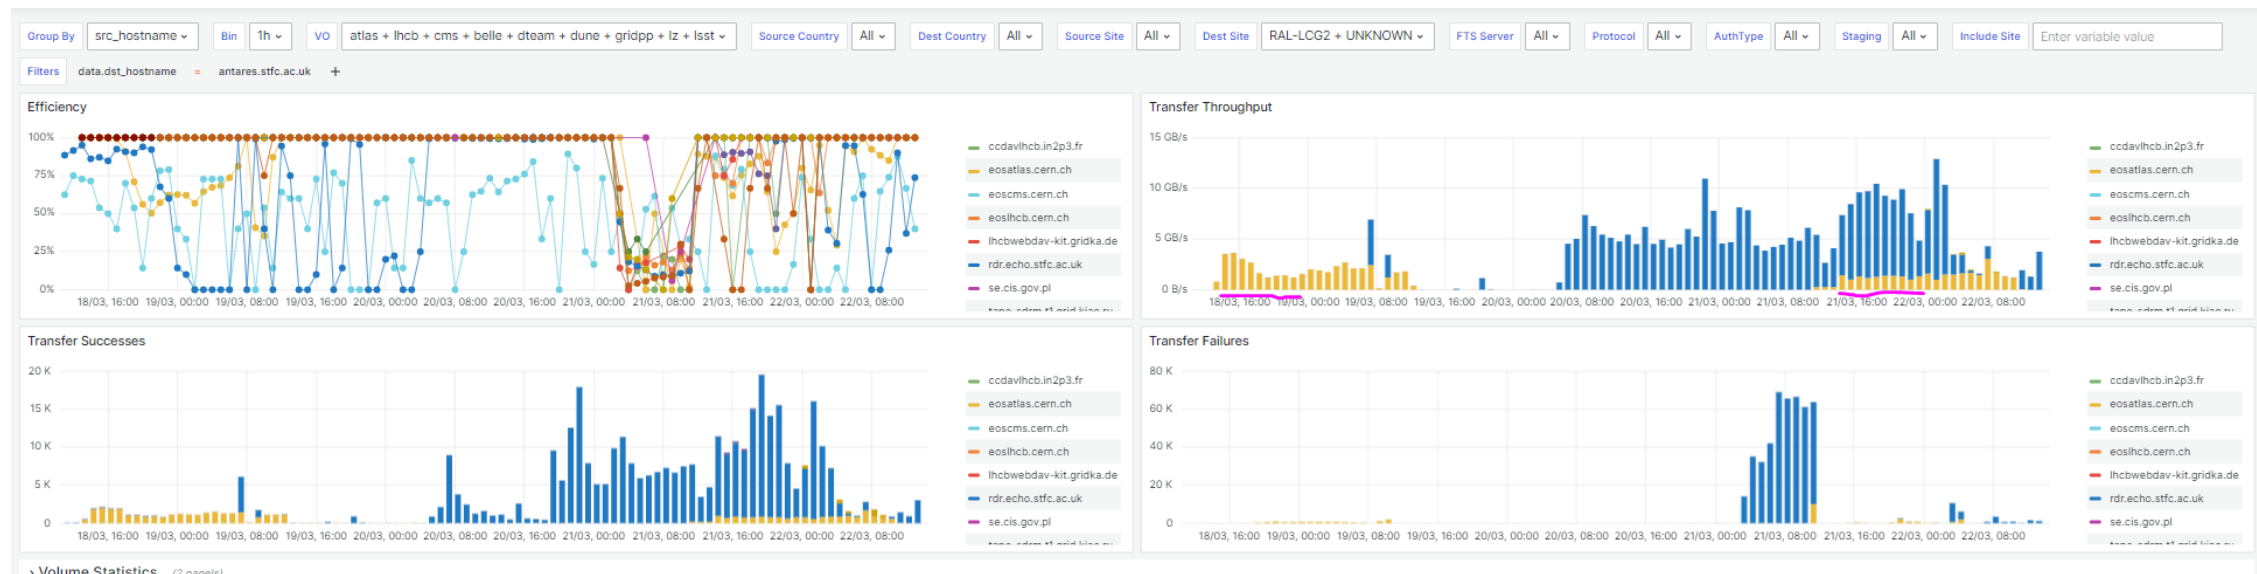


DC24 – Things to improve

- Deletes
 - Under high load deletes go very slowly. Deletes quickly caught up after DC24 finished. Described in Jyothish's Talk.
- Load Balancing
 - We got lots of experience trying to balance the load across the gateways. Described in detail in Jyothish's Talk.
- Check-summing
 - When a file is written via FTS, it is then recalled to the gateway to calculate the checksum. It is on Brian Bockelman's ToDo list to add check-summing while the data is being streamed.
- Gateway Hardware
 - We hit the 25Gb/s network limit on our gateways, we have purchased a 100Gb/s capable gateway. There is explicit effort in GridPP7 for improving throughput.
- Tokens
 - For various reasons these had to be disabled. Most of the improvements will come from the VOs, but there are some XRootD patches needed.

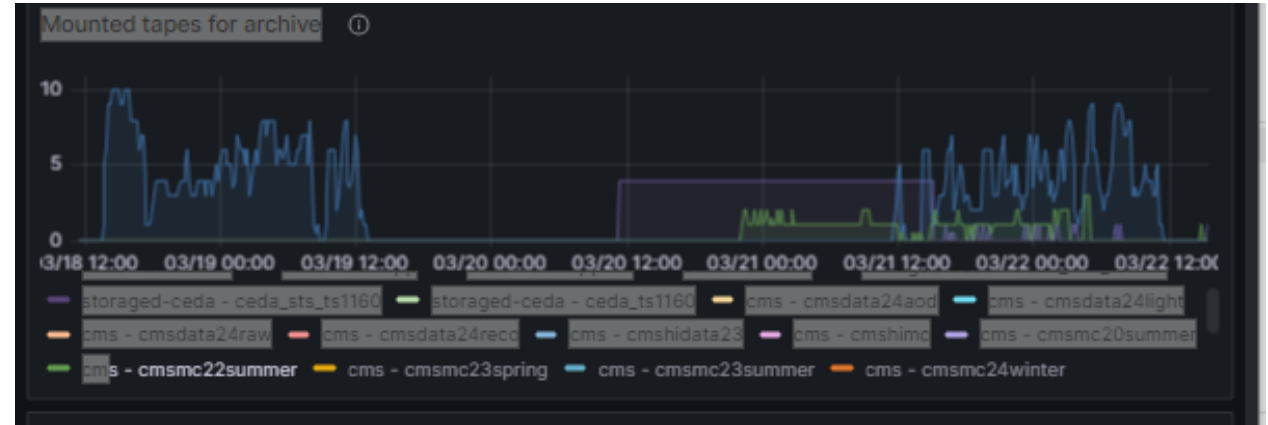
Antares – Data Export test

- Webdav and the HTTP REST Tape API mean that Antares can be accessed directly.
 - Raw Data export will be going directly from CERN to Antares.
- In the repeat DC24 test we asked for Antares to be tested.
 - Various teething problems (e.g. FTS config wasn't correct, deletes were slow).
 - Antares was ingesting data at 10GB/s but most of it wasn't DC data...

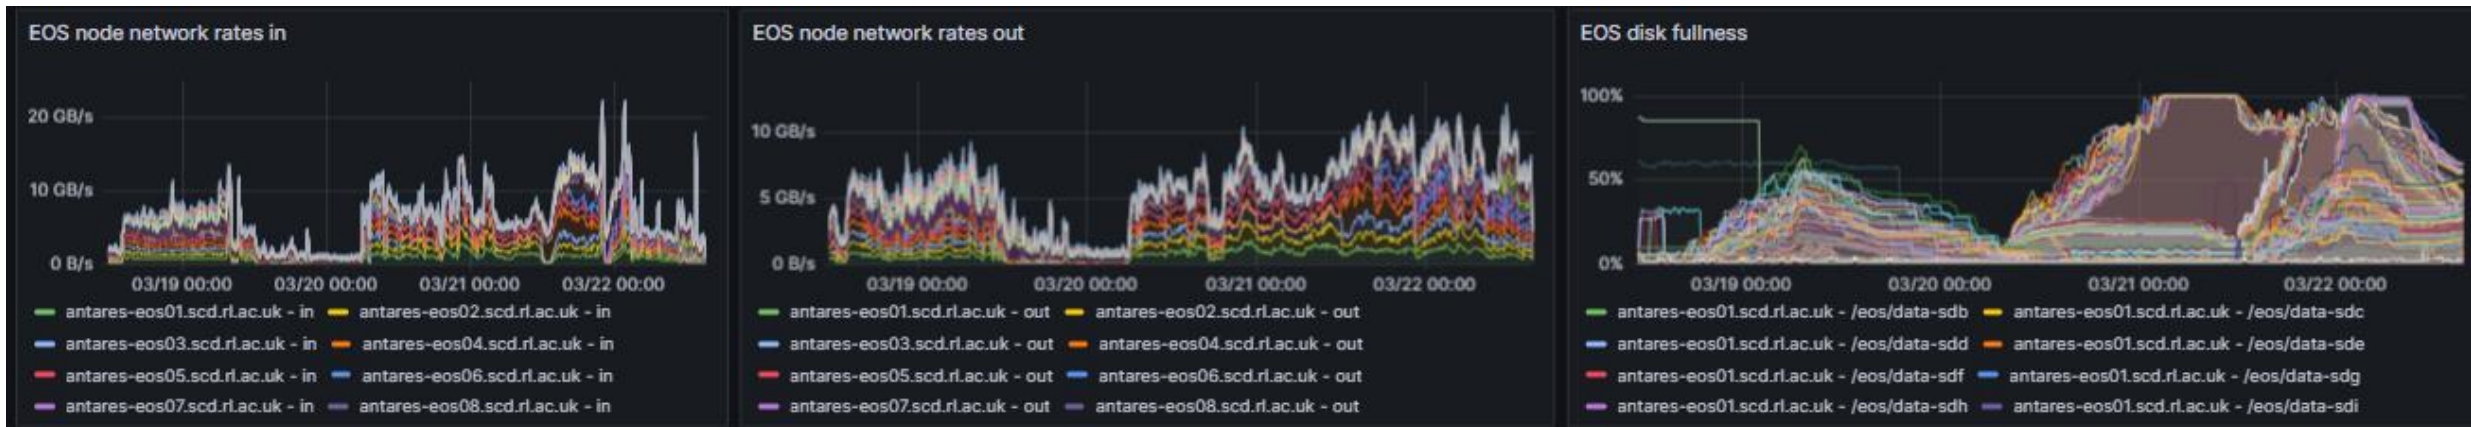


Antares – Tape backend

- Tape backend was definitely well tested!
- The EOS buffer did fill up during the test, which then reduced transfer speeds.
- SSD costs are coming down so buffer can be expanded.
- Plenty of room to optimize tape drive usage.



Blue in data challenge, Purple is atlas21mcother, Green is cmsmc23summer

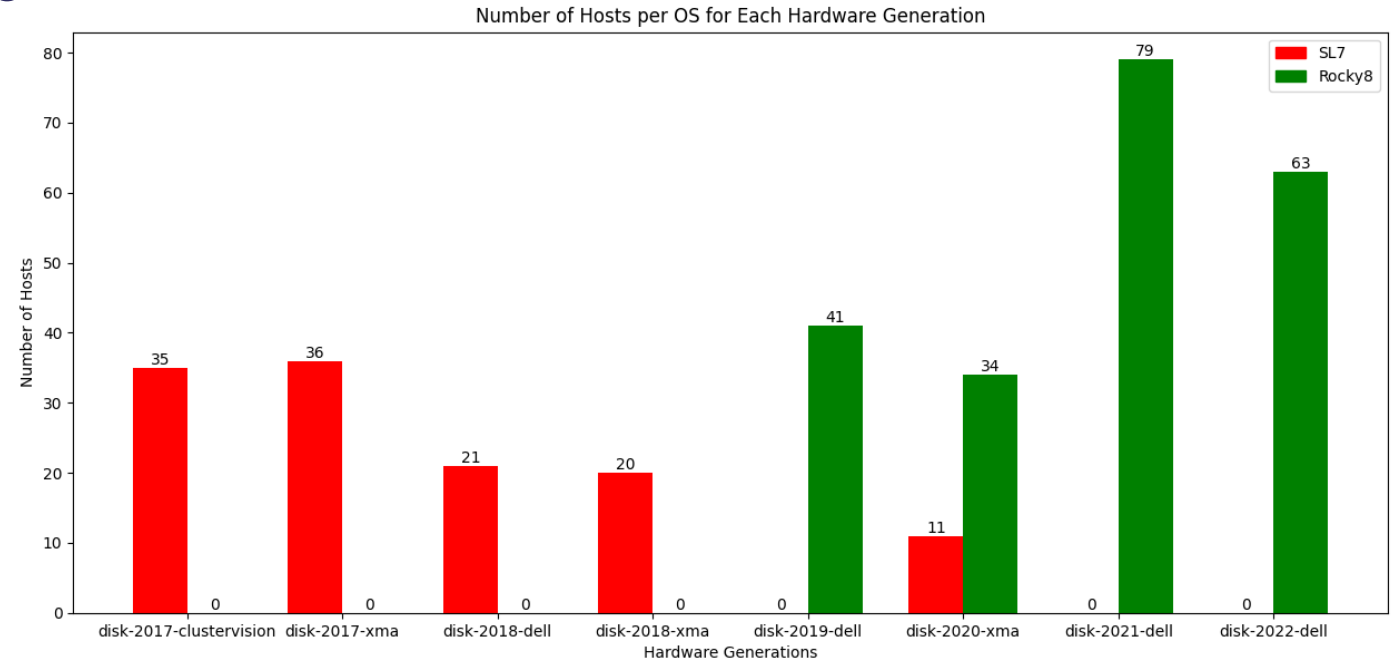


Echo status and plans

- Echo recently exceed 100PiB of raw capacity.
- Backend storage working very well.
 - Many operational changes can be done transparently, and the team is trying to automate them where possible.
- 2024 Q2:
 - SL7 → Rocky 8
 - Decommission 2017/18 generations of hardware
 - Continue deployment/weighting up of 2022 hardware
- Q3
 - Complete deployment and decommissioning
 - Upgrade to Pacific
 - Start deployment of 2023
- Longer term we still would like to implement rack-level failure domains.

Rocky 8 Upgrade Progress– Echo Storage


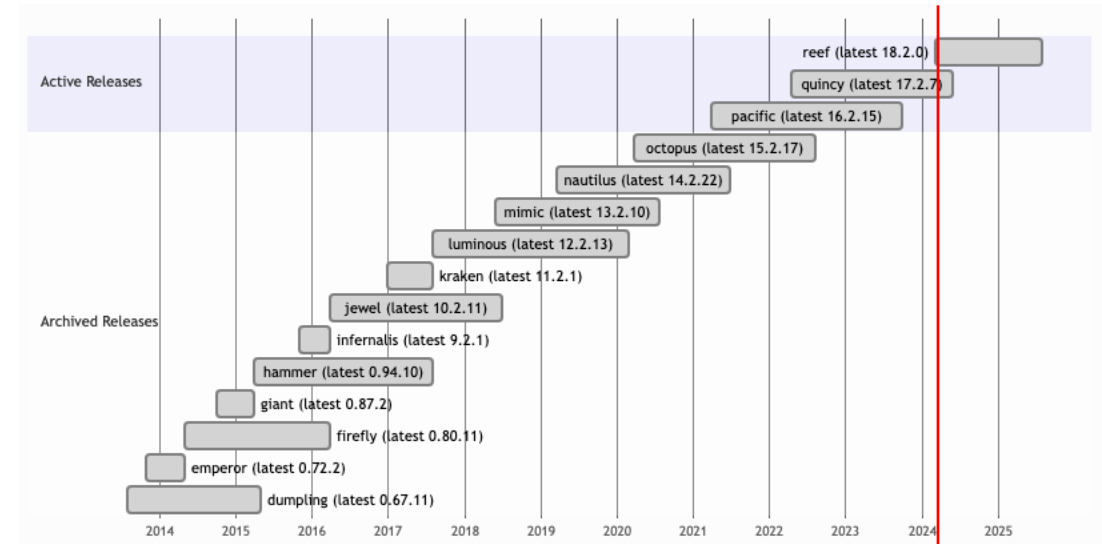
- OS upgrades to Echo storage nodes are automated
 - Working hours only
- Each storage generation has its own foibles
 - Development needed each time
 - We intend to decommission 2017 before the SL7 EOL
 - If we can decommission 2018 as well, we will.



Credit to Maksim for the graph and much of the work.

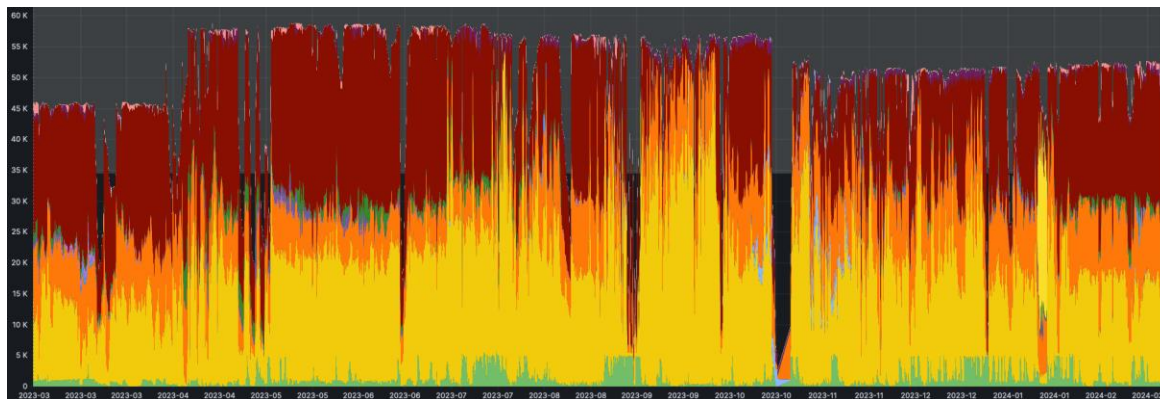
Ceph release schedule

- Our last major Ceph upgrade was in May 2020!
- Currently we are on Nautilus.
- Double upgrade to Pacific by September 2024.
- Double upgrade to Reef for April 2025.
- In GridPP7 want to aim to be looking to upgrade to latest release once it has been out for 6+ months.



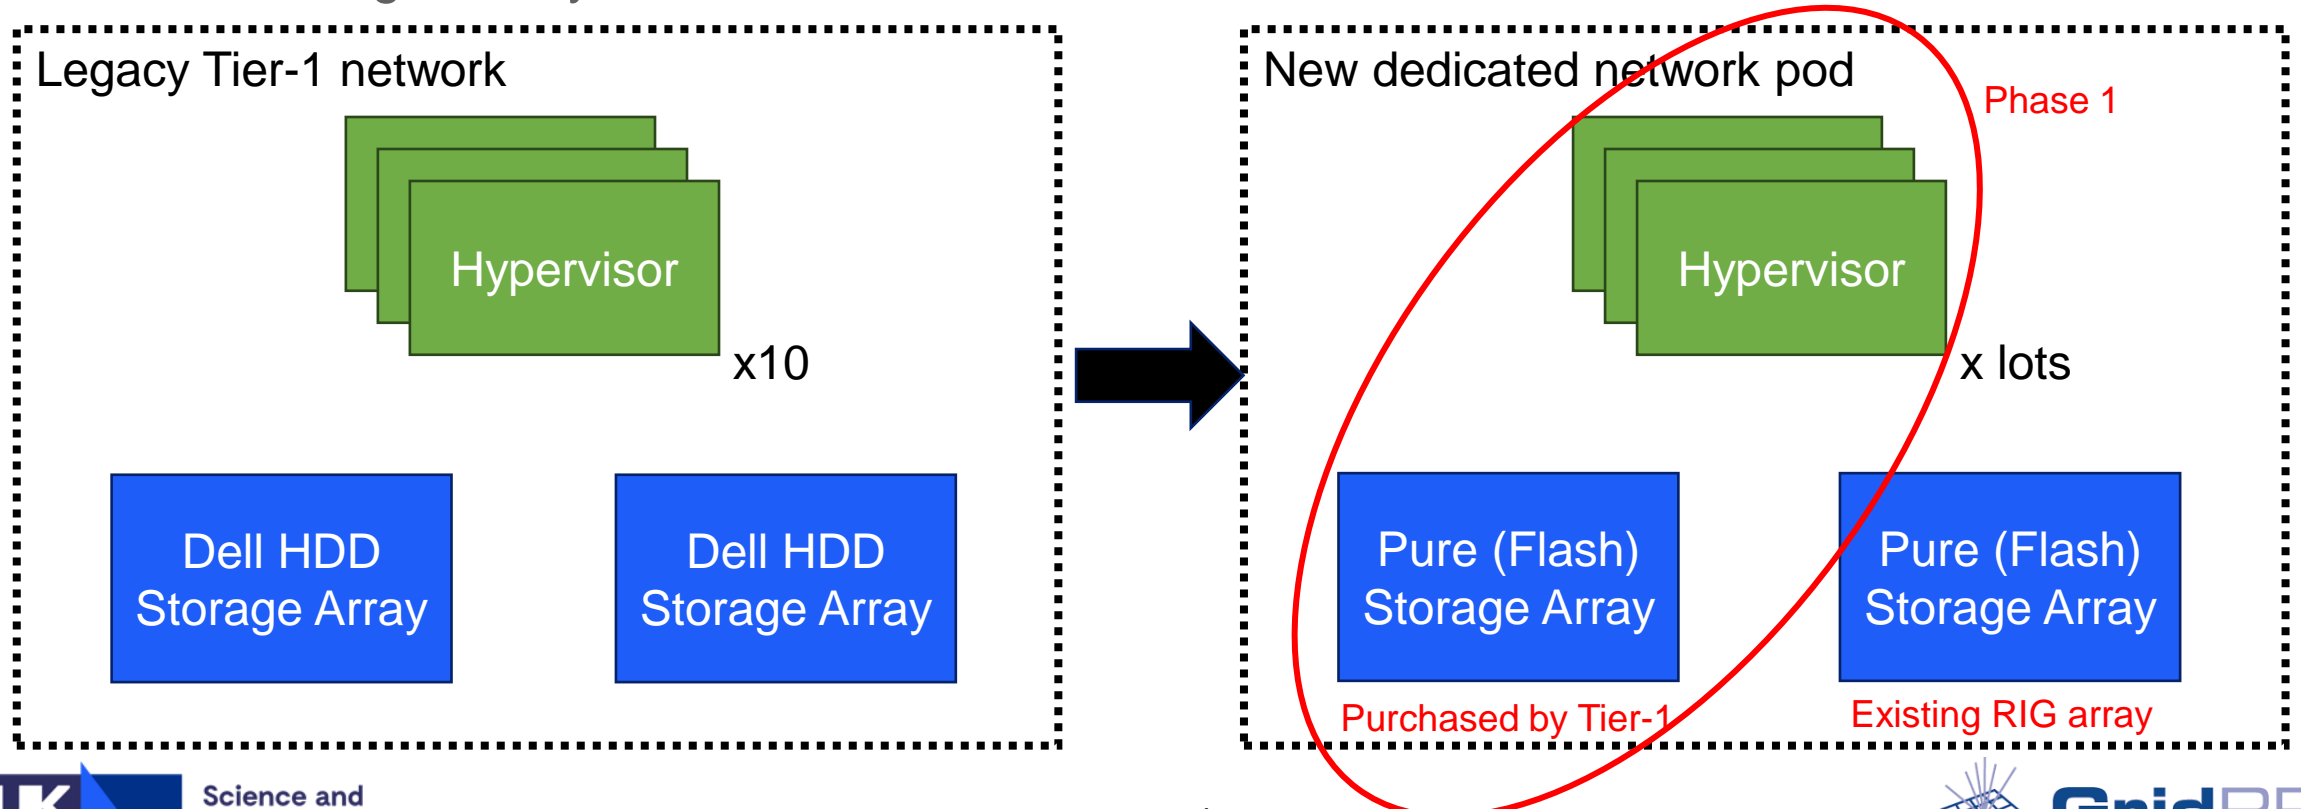
Batch Farm

- HTCondor and ARC CEs continue to perform very well for the Tier-1.
 - The utilization of the resources very impressive when compared to other science areas.
- The most frequent operational problem is with VOs not always receiving their desired fairshare.
 - This problem will become more complex as we add different architectures.
- Tom Birkett has setup a pre-prod cluster (which runs actual work and is accounted, and can be used to test the latest patches).
 - We have a single node dual stack which is being tested.



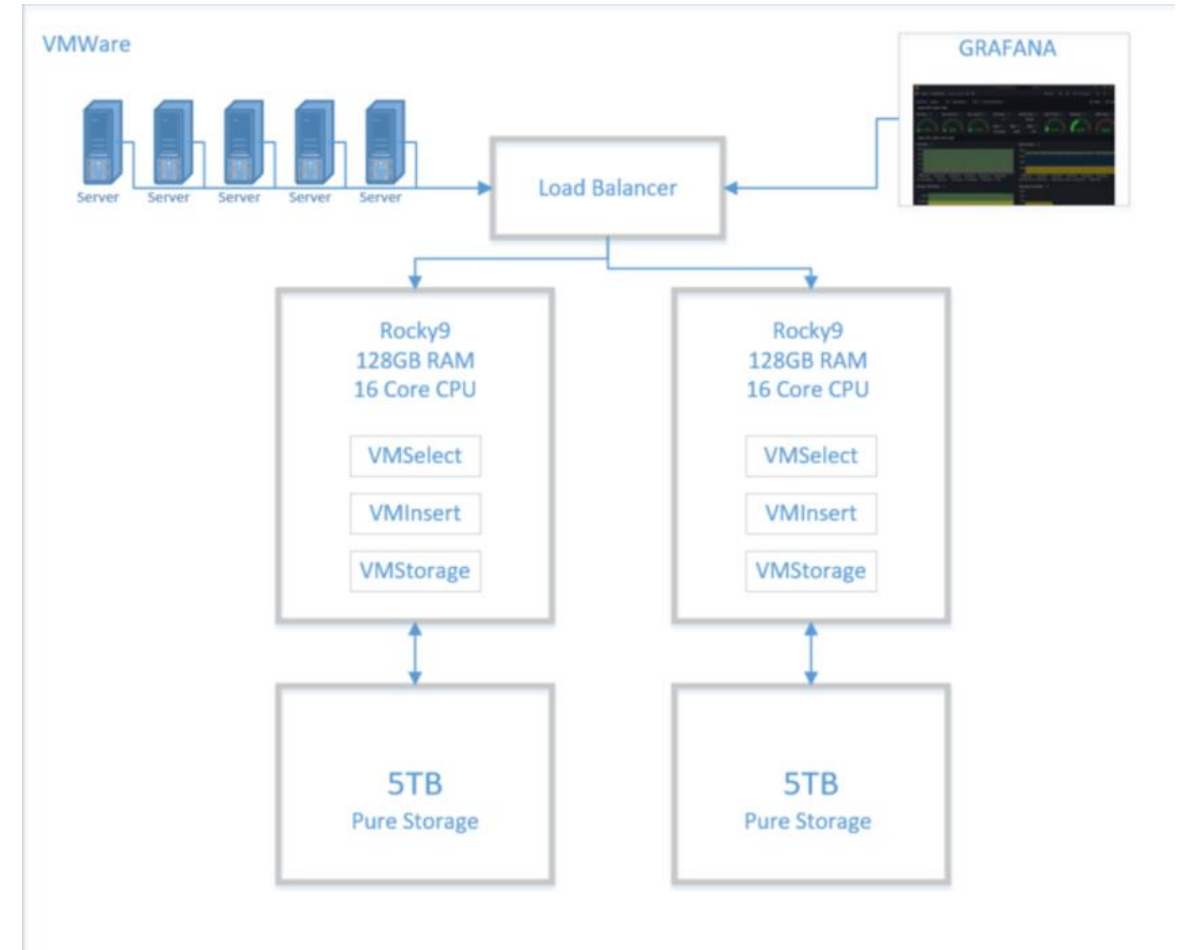
VMWare replacement

- SCD currently runs multiple VMWare instances.
- Tier-1 VMWare Cluster is 5 years old and we plan to replace it.
 - Joint SCD project underway to replace it.
- “Small” issue with VMWare license costs means we are migrating to something else which is being actively researched!



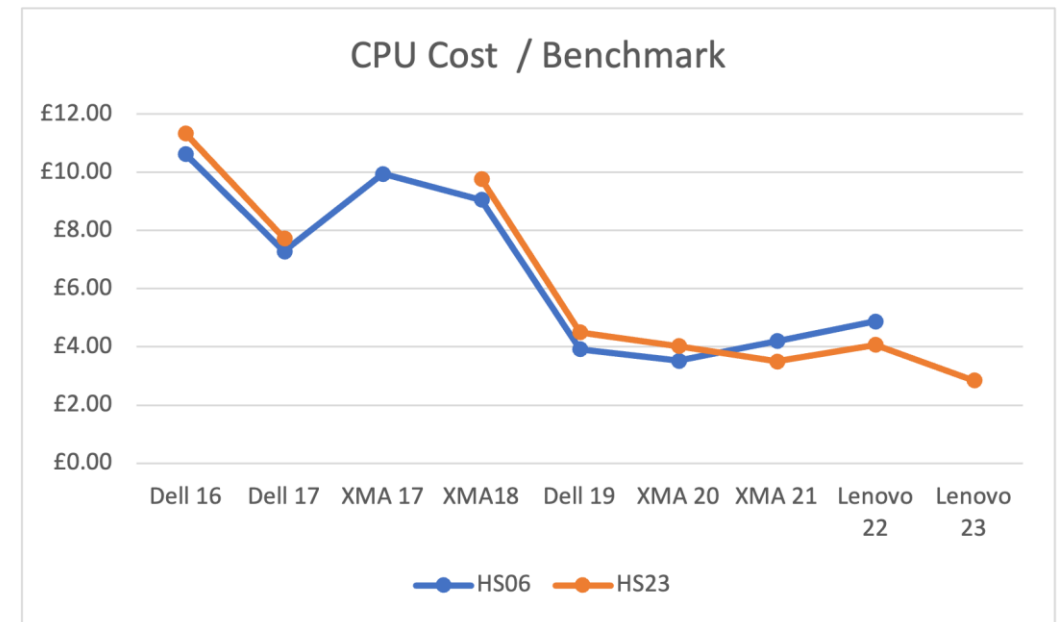
Victoria Metrics

- We currently use InfluxDB for our time series monitoring.
 - Very old version – Security is OK because not publicly accessible.
- We are running a (joint) project to replace it.
 - New service will be called Timon = Time MONitor
- We currently plan to move to Victoria Metrics.
 - I am negotiating to get the Enterprise version for a 3 year contract at ~£20k per year for 0.5 million data points per second ingest.



CPU procurement

- This year SCD did a joint CPU procurement for JASMIN, SCARF, Tier-1 and the STFC Cloud.
 - Compute, Memory and OS Disk were identical, differences in networking and local storage.
 - Total value of £7,772,995.52 (inc VAT) was awarded to Lenovo.
- 407 servers in total (36 for Tier-1) with:
 - Dual AMD EPYC 9654 (HS23= 6000.6 each).
 - 24 x 64GB memory (= 1.5TB memory).
 - JASMIN purchased some 6TB memory servers.
 - 480GB SSD OS drive.
 - Mellanox ConnectX-6 25Gb/s NICs.



Disk and other procurement

- We purchased 32 Dell Storage nodes
 - 28 x 22TB drives.
 - 2 x 16C/32T CPUs
 - 480GB SSD OS Disk
 - 256GB memory
 - 25Gb/s Mellanox NIC.
- Echo gateways
 - 4 with 25Gb/s NICs
 - 1 with 100Gb/s NICs
- We are buying individually metered power supplies.
- Earlier in the year we purchase a single ARM server which we integrated into the farm, we intend to buy more this year.



Science and
Technology
Facilities Council

Questions?