

Diamond Light Source Site Report

James Thorne
james.thorne@diamond.ac.uk

Additional material from Dai Jenkins and Chris Reynolds

Diamond Light Source



Roof replacement



Diamond II

NEWS

Home | War in Ukraine | Climate | Video | World | US & Canada | UK | Business | Tech | Science

England | Regions | Oxford

Oxfordshire research centre gets £519m government investment

6 September



GETTY IMAGES

Compute

- 15,816 cores
- 404 GPUs (K80, P100, V100, A100)
- RHEL7+UGE+CFEngine
- RHEL8+Slurm+Ansible
- Infiniband

Storage

- NetApp: ~900 TB (science)
- NetApp: ~200 TB (controls/accelerator)
- GPFS02: 7.2 PB – end of life
- GPFS03: 16 PB
- GPFS04: 9 PB

Archive

- Run by RAL
- >54.7 PB on tape (+14.7 PB in last year)
- Powered by CTA


Total data registered in StorageD

Data collected (Tb) Files in system

54780.9

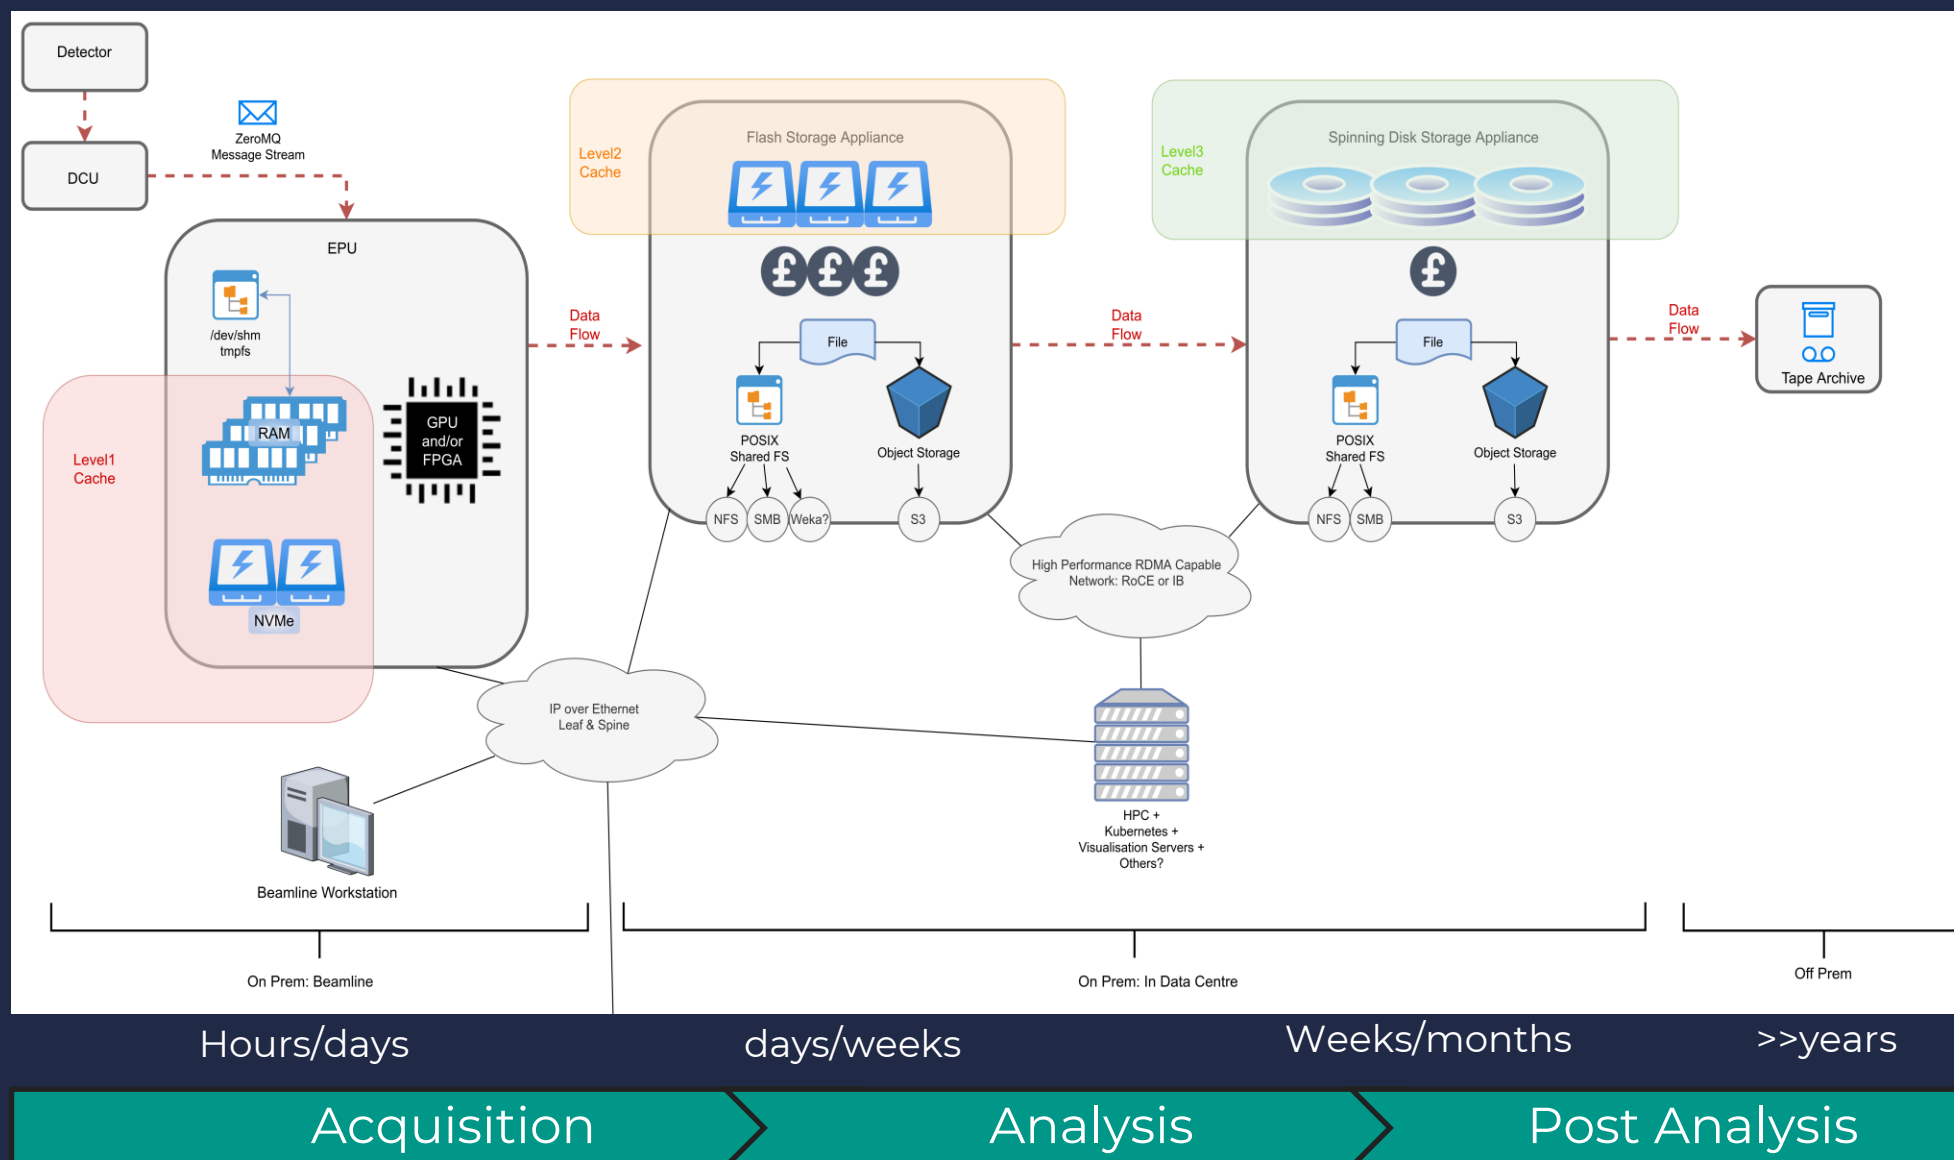
4365774559

Data Centres: CSCR4



- CSCR4 / CS06R: First water cooled racks for Diamond
- Main role to host compute and storage infrastructure for Initial Data Acquisition and Analysis;
- Designed to accommodate 1MW IT Load
- 40 Racks:
 - 20 Low Power (25kW) – Compute/Storage
 - 15 High Power (60kW) - Compute
 - 5 Networking (5kW)
 - ~ 15mins UPS
- Rack cooling
 - Low power – Air-cooled with 25kW Rear Door Heat Exchangers (RDHX)
 - High power - 60kW Direct Liquid Cooled (DLC)
- Will provide complementary capabilities and capacity to CSCR2 and CSCR3
- Projected handover - Calendar quarter 2/3 2025

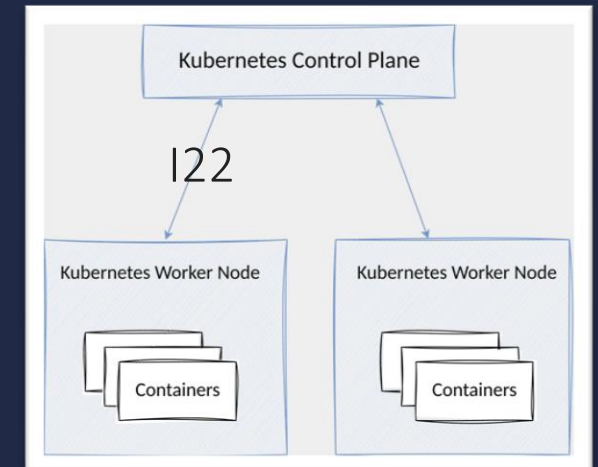
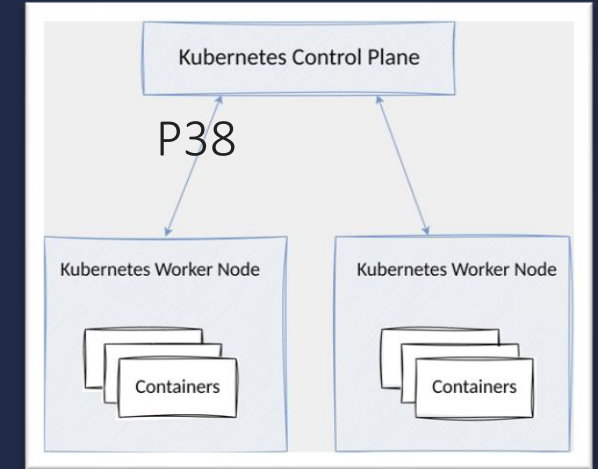
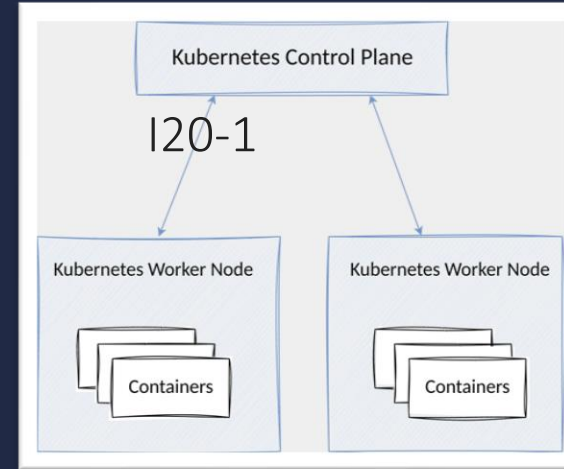
Storage: fast cache and tiers



Potential Compute Platform Alignment???

Kubernetes Cloud Platform

- Hypervisor Hardware purchased for "Kubernetes Cluster per Beamline"
- Handful of early clusters deployed: P38, I22, I20-1
- Multi cluster management tooling worked out.
- Moved Kubernetes storage to use new Netapp.



Cloud Bursting

IRIS Storage

- Working on a data mover/copy tool to rapidly transfer data from /dls to IRIS.
- Data copied will reside on IRIS CephFS based fast shared storage.
- 200TB of CephFS available.

IRIS Scheduler:

- Early prototyping with Slurm on IRIS.
- Job submission will be via Wilson/DLS Slurm REST API.
- Accounting will be via Wilson/DLS Slurm database.
- The Wilson Slurm REST API version needs an update to support this.

Miscellaneous

- Linux Distribution
 - Updating to RHEL 8
 - Lots of RHEL 7, especially on beamlines
 - **Still** some RHEL 6 😬
 - Thinking about the future – still RHEL?
- CFEngine and Ansible
- CVMFS
- Mishaps – ask me over coffee

Other Diamond talks this week

- [Graylog-as-a-service – this morning](#)
- [Grid Engine to Slurm – 15:35 today](#)
- [MFA for SSH – 14:00 tomorrow](#)

Questions?