

Storage at DLS - One Size fits no one.

Frederik Ferner <frederik.ferner@diamond.ac.uk>

Diamond Light Source Ltd

April 23, 2020

Outline

Diamond

Storage and their use at DLS

Detectors — Data Acquisition

HPC — Data Analysis

Visualisation

Current Implementations

Diamond

Diamond Light Source (DLS) is the UK's national synchrotron facility. It is located at the Harwell Science and Innovation Campus in Oxfordshire, UK.

- ▶ Synchrotron
- ▶ many beamlines/detectors
- ▶ users/scientists need quick feedback on data quality for experiment planing



Storage uses at DLS

- ▶ Shared use: Data Acquisition, Processing and Visualisation in parallel
- ▶ some HPC for processing
- ▶ Scientists want to see data on Desktops
- ▶ NFS (and CIFS) access required
- ▶ Complex access control needs

Data ingest: Detectors

- ▶ many different detectors, different special issues
- ▶ two main classes:
 - ▶ many files per scan (.cbf, .tiff, .jpg)
 - ▶ small number of large files (hdf5)

Detector — Pilatus

Pilatus:

- ▶ relatively small files (typically 6MB files at $\sim 100\text{Hz}$, ~ 10000 per data collection)
- ▶ latency until data visible on cluster critical
- ▶ custom tool to copy from detector to file system (rsync to slow)

Detector — Eiger

Eiger

- ▶ example of a detector writing (mostly) hdf5 files
- ▶ relatively large files, 3-5GB for Eiger, generated at about 2GB/s sustained
- ▶ Single Writer Multiple Reader (SWMR), needs POSIX (not NFS)

Detector — PCO Edge

PCO Edge Camera

- ▶ Windows drivers only*
- ▶ 1GB/s
- ▶ tiff images
- ▶ ... or with DLS software hdf5 files
- ▶ tomography, experiments can write data at full rate for hours
- ▶ 100+ GB per collection frequent

Data Analysis: HPC or really HTC

- ▶ Jobs scheduled on SGE cluster
- ▶ many nodes reading data in parallel
- ▶ quick turn around... stat() and read rates import
- ▶ lots....



Visualisation/Interactive Access

- ▶ Scientists want to see their data
- ▶ Workstation type machines, interactive access
- ▶ NFS/CIFS



Impatient Scientist Scenario

- ▶ detector creates files
- ▶ user wants to confirm all working => runs 'ls' on workstation
- ▶ 'ls' 60+s
- ▶ expected NFS behaviour?
- ▶ eventually hotfix....

Data Filesystems at DLS

- ▶ GPFS — Initially chosen for single stream performance
- ▶ 3 separate file systems
- ▶ 2 separate Data Centres
- ▶ IB
- ▶ Ethernet

Questions

Questions?

