

Diamond Light Source Site Report

Frederik Ferner
Scientific Computing Team



COVID-19 Impact

- Spring 2020: All experiments cancelled except those related to the pandemic.
- Summer 2020: Increase in experiments and remote beamline operation
- Autumn/Winter 2020: Remote operations again – mostly still the case
- Huge increase in use of remote access methods: Zoom, Slack, Teams, NoMachine.
- PCR and lateral flow testing on site.

Remote access

- NoMachine
 - Always had remote access for beamline control
 - Increase number of nodes in the service
 - Limit memory per user
 - Nodes with V100 GPU to provide remote visualisation
- SSH
 - Second bastion host
- New Fortinet SSL VPN

Storage

- GPFS02, 7 PB
- GPFS03, 16 PB
- NetApp
 - 200 TB replicated to Google Cloud using "NetApp On Command Cloud Manager"
 - Decommissioned smaller, on-site replication targets

Compute

- (Still) Univa Grid Engine
 - 13,000 cores
 - 420 GPUs (48 K80, 192 P100, 180 V100/V100S)
 - Dell, Lenovo, Supermicro
- Historic accounting information now all in Elasticsearch
- Kubernetes
 - Jupyter, DASK, HTCondor, ancillary services
- Prototype "cluster-as-a-service" portal

Challenges

- Further demands for remote access
- Partitioning of data acquisition, initial/auto processing, post-processing, "data visits"
- RHEL licensing?
 - Change to another distribution?
 - Support for GPFS
- Move to Slurm from Grid Engine?
- More data and on disk for longer

Recruiting for a sysadmin

- <https://vacancies.diamond.ac.uk/vacancy/storage-systems-administrator-senior-storage-systems-administrator-435678.html>
- Closing date 21 March 2021