



ROYAL
HOLLOWAY
UNIVERSITY
OF LONDON

RHUL Site Report

Antonio Perez,
Simon George, Barry Green, Tom Crane
HEP Sysman meeting @ RAL, May 2019



ROYAL
HOLLOWAY
UNIVERSITY
OF LONDON

Group Activities (update)



ROYAL
HOLLOWAY
UNIVERSITY
OF LONDON

ATLAS

- Benefit from strong collaboration in software support.
- Large Tier3 batch compute and storage resources for data analysis.
- DAQ test systems.

Dark Matter & Neutrino

- Detector development: lab DAQ systems and **significant quantities of test beam data to store and analyse.**
- Growing need for compute and storage resources to simulate and analyse data.
- Need help with things like installation, data movement.
- **Have overtaken ATLAS as the largest user of Tier3 storage.**

Accelerator

- Small DAQ systems.
- Many small activities around the world which generates unique and valuable data sets.
- Simulation: both embarrassingly parallel and multi-process (MPI) computing.
- Software development infrastructure (e.g. cdash server).

Theory

- Occasional use of Tier3 **and HPC clusters.**

Tier-2 9 Racks in off-site machine room

Still using 9 racks, no room to expand




Tier-2




- 9 racks in modern machine room provided by central IT service.
- 46 kHS06 in 120 WN (SL6)
- Cream/Torque CE x 2, moving to HTCondorCE & HTC
- 1.4 PB DPM SE in 47 servers
- 8 misc servers including 3 VM hosts running standard network services, provisioning, Grid services.
- New kit from GridPP6 & RHUL funding arriving soon:
 - 12x Dell PowerEdge C6420, 2x Xeon Gold 6148, 80 cores per node, total 12 kHS06
 - 9x Dell R540 HBA 120TB raw, total 1 PB raw
- Lack of space meant we had to remove old machines used to test VAC

Storage

Colours show distinct procurements




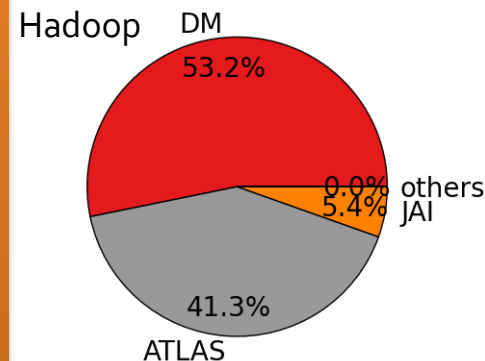
Compute



Tier 3



- 8 racks packed into home-made machine room
- Now connected to Tier2 by 10Gb/s fibre
- 6 kHS06 in ~100 old WN, mostly hand downs from Tier2, some upgraded
- Torque/Maui batch system with scaling problems
- Storage:
 - 100 TB Hadoop (3 replicas) using 90 WNs
 - 226 TB NFS scratch over 3 servers
 - 11 TB NFS Home
- 3 VM host machines running multiple redundant squids & NATs, interactive user machines, network services, development services, odds & ends.





Tier2

- Migrate batch system to HTC on CC7 WNs
- Replace CREAM CEs with HTCondorCE
- DPM upgrade? Too soon to replace with cache?
- Retire some older DPM pool nodes
- Add redundant network connections to storage and between switches
- Automate shutdown in case of power cut and backup generator fail
- IPV6: still waiting for local networks team to do the paperwork for the JANET DNS service.

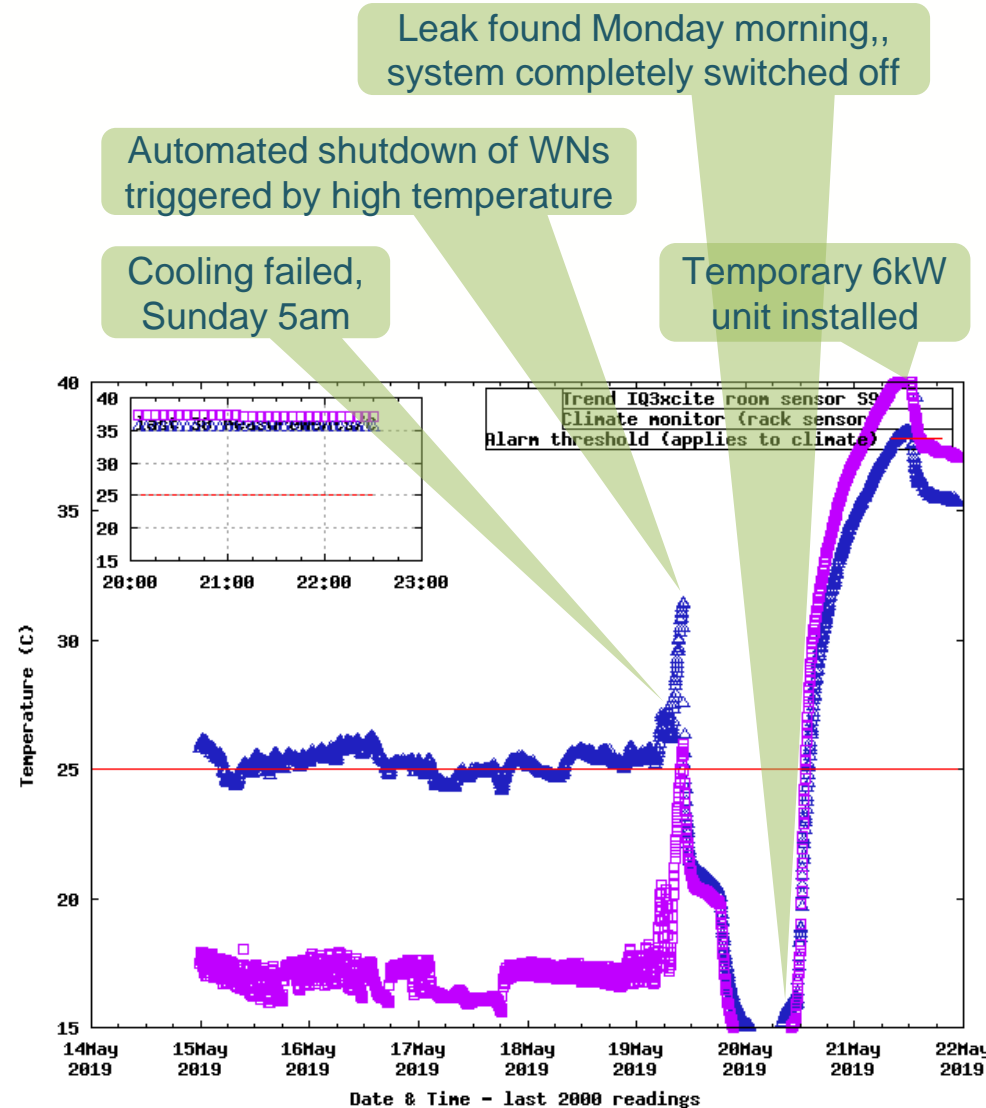
Tier3

- Consolidate batch function onto Tier2 batch system and retire old worker nodes
- Deploy new HDFS storage co-located with Tier2
- Upgrade Tier2-Tier3 connection to 100Gb/s
- Migrate home service to new hardware

Problems



- Tier3 hit by major aircon failure (again, last time was November 2017)
- Good test of our automated shutdown procedures ☐
- Leak in condenser system will take at least 2 weeks to fix 🙄
- Struggling to run bare minimum of systems in indoor room



Thank You