RHUL Site Report

Govind Songara, Antonio Perez,
Simon George, Barry Green, Tom Crane
HEP Sysman meeting @ RAL , June 2018

Manpower

Antonio - Tier-3/Tier2 Sysadmin - 1FTE .
Govind - Tier-2 Grid Admin - 0.5 FTE.
Simon - Site Manager.
Barry - Hardware/Network Specialist.
Tom - All rounder.

Group Activities

ATLAS

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

Dark Matter

- Detector development (lab DAQ systems).
- Growing need for compute and storage resources to analyse data.
- Need help with things likes installation, data movement.

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 cluster.
- Interest in MPI.

Tier-2 9 Racks @ Huntersdale site

Upgraded from 8 to 9 racks.

Tier-2 Grid

- 9 racks in modern machine room provided by central IT service.
- New Kit: 10 x Dell PowerEdge C6420, Xeon(R) Gold 6148 CPU @ 2.40GHz, 40core, SL6-HS06-score: 966 / CC7-HS06-Score 1004 with HT
- 48 kHS06 in 150 WN; Cream/Torque CE x 2, moving to ARC/HTC
- 1.4 PB DPM SE in 47 servers
- 8 misc servers including 3 VM hosts running standard network services, provisioning, Grid services.
- Network restructured. Using stacking to solve some reliability problems. See next slides.
- 0.5 FTE
- Vac-in-a-box nodes updated to latest version and CC7.

Storage

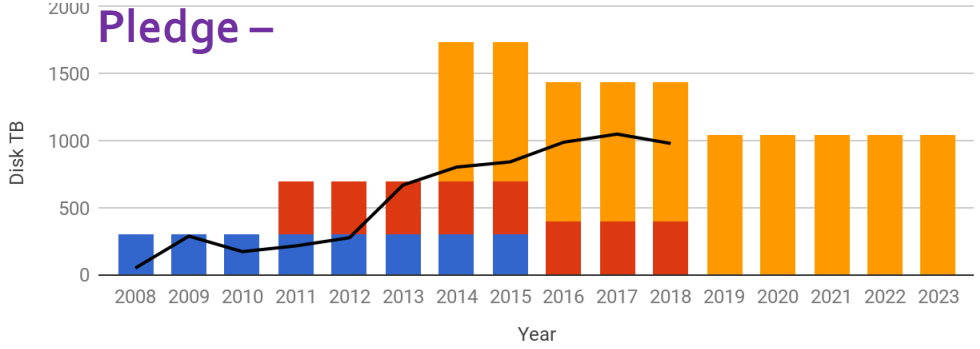

Compute

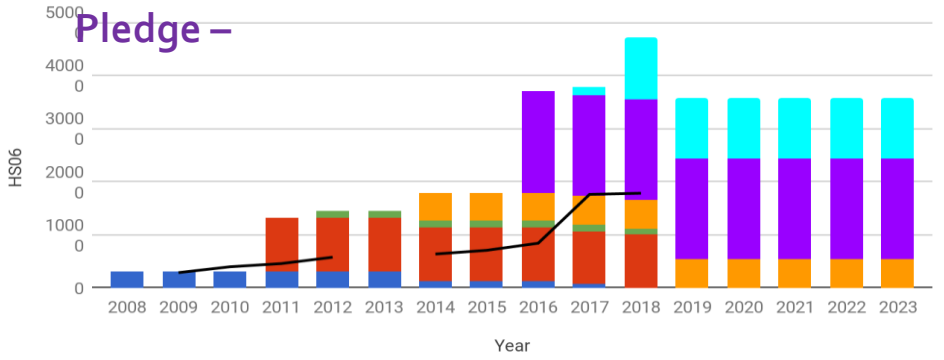

Network design 2016

Network design 2016

- Cacti monitoring shows no bottlenecks
- Running out of ports on pair of well-connected S4810s
- Outages caused by issue with individual PC6248
- Connection to Tier3 only via slow Campus LAN

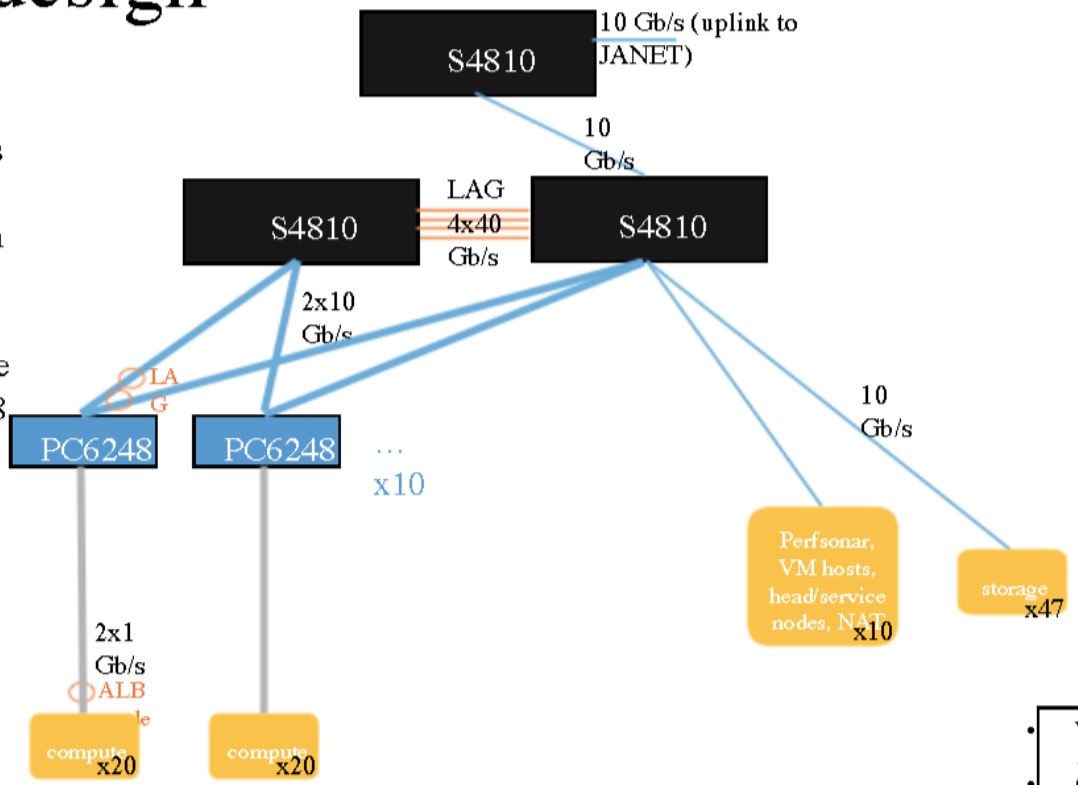

- VLANs and control network (IPMI, APC, switches) not shown.
- Connections to specific S4810 switches just indicative not literal.
- Multiplication factor xN includes nodes, switches and links.

Network design 2018

ROYAL HOLLOWAY UNIVERSITY OF LONDON

Network design 2018

Main changes:

- Stack S48 10s to have high bandwidth between all 3 switches
- Compute nodes connected to two switches for redundancy
- Fibre to Tier3 commissioned at 10 Gb/s, plans to upgrade to 100 Gb/s
- Harmonised VLANs with Tier3

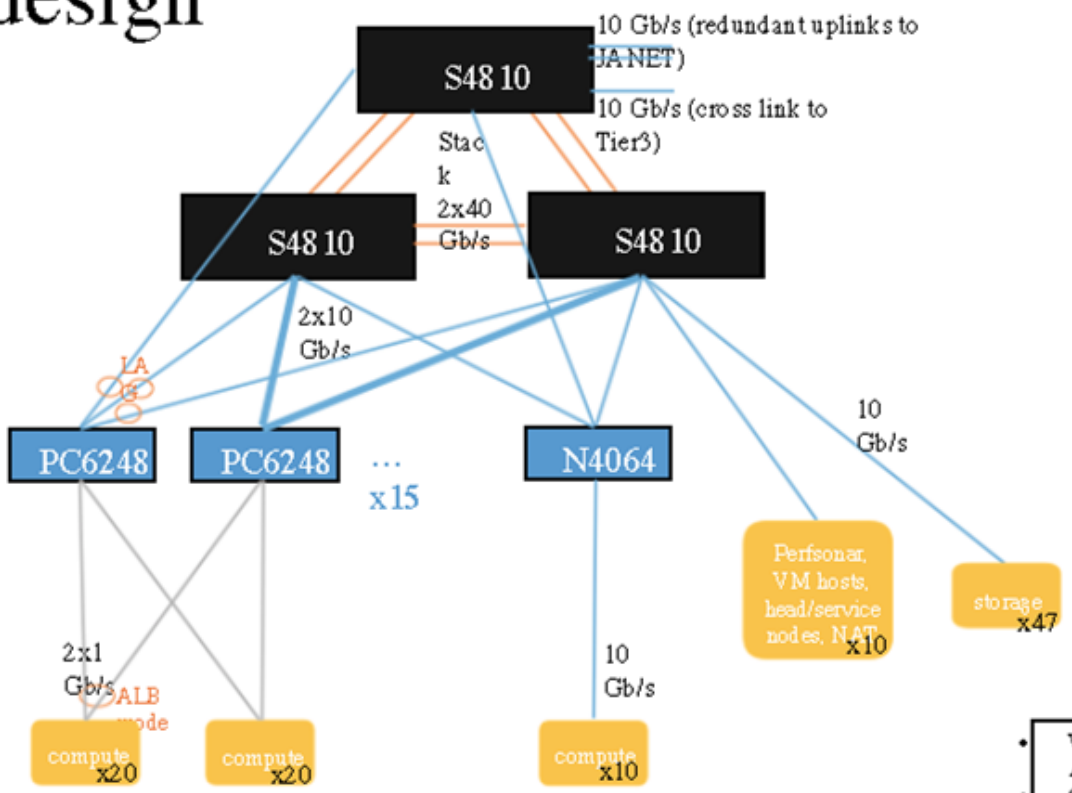

- VLANs and control network (IPMI, APC, switches) not shown.
- Connections to specific S4810 switches just indicative not literal.
- Multiplication factor xN includes nodes, switches and links.

Tier-2 Grid

ROYAL HOLLOWAY UNIVERSITY OF LONDON

Issues

- Some services impacted due to the network upgrade.
- **Long term expansion limited by rack space and cooling,** need to decommissioned old kit to make room.
- The need to have custom routing to put storage traffic on our private network currently prevents us from expanding VIAB.

Planning

- CC7 Deployment server, services, cluster, etc.
- CC7 DPM pool node upgrade
- IPV6 roll out still on going.

Tier -3

ROYAL HOLLOWAY UNIVERSITY OF LONDON

- 8 racks packed into home-made machine room, a long way from Tier2 on 1 Gb/s link
- 6 kHS06 in ~100 old WN, mostly hand downs from Tier2, some upgraded; Torque.
- Storage:
  - 304TB Hadoop using 90 WNs
  - 262 TB NFS scratch over 7 servers
  - 11 TB NFS Home
- 11 servers running standard network services, mainly as Vms.
- Hadoop updated from 2.4.1 to 2.7.3
- Network mapping – NetDisco, lldp on all hosts

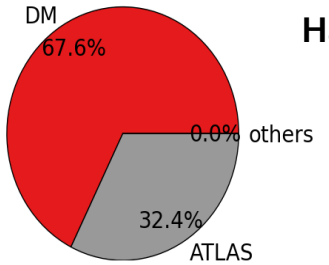

Hadoop

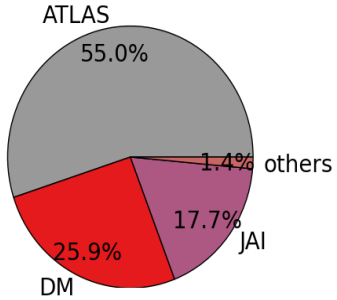

Scratch

Tier-3

ROYAL HOLLOWAY UNIVERSITY OF LONDON

Issues

- Aircon failures. Shut down script creation.
- Hadoop mismatch version between namenode and datanodes prevented load balancer to work.
- Issue found in HTCondor setup.
- Problem found with HTCondor and AMD cpus where a HTCondor was crashing only when running under AMD.

Planning

- Migrate batch system from torque to HTCondor.
- Hadoop test environment update to version 3.

HPC

New HPC located in Huntersdale.

- **1 xcat node + 11 nodes:** 20 x Intel(R) Xeon(R) CPU E5-2640 v4 @ 2.40GHz, 62GB RAM. 22Tb of storage.
- Running mainly, SLURM and easybuild under Centos 7.2.
- Using Salt as a configuration manager.
- Centos 7 AD authentication setup on going.


ROYAL
HOLLOWAY
UNIVERSITY
OF LONDON

Thank You


ROYAL
HOLLOWAY
UNIVERSITY
OF LONDON