Scotgrid Glasgow Site Report

David Crooks
for Scotgrid Glasgow

HEPSYSMAN, June 2017

Team

- Gareth (systems manager)
- David (systems manager)
- Sam (data)
- Gordon (local group/school/grid)

Current status

- "Headline details"
- SL6/CentOS 6/CentOS 7 (new services)
- ARC/HTCondor
- DPM
  - New storage still CentOS 6

Storage

- Deploying new storage, 3 × (host+shelf)
  - Each shelf has 60 10TB disks.
  - DELL R630/MD3120
- Toolchain to support HBAs vs RAID controller
- ZFS

Storage

Worker nodes

- Most recent worker node tranche deployed as VAC
  - HPE ProLiant DL60 / E5-2630v4 / 40 HT cores / 4GB/HT core
- Gareth has spent some time with deployment
  - full report to follow
- Non-trivial amount of work to deploy with our requirements (CentOS 7 / non-VIAB)
- Ansible roles for CentOS 7
  - Talk to Gareth!

VAC

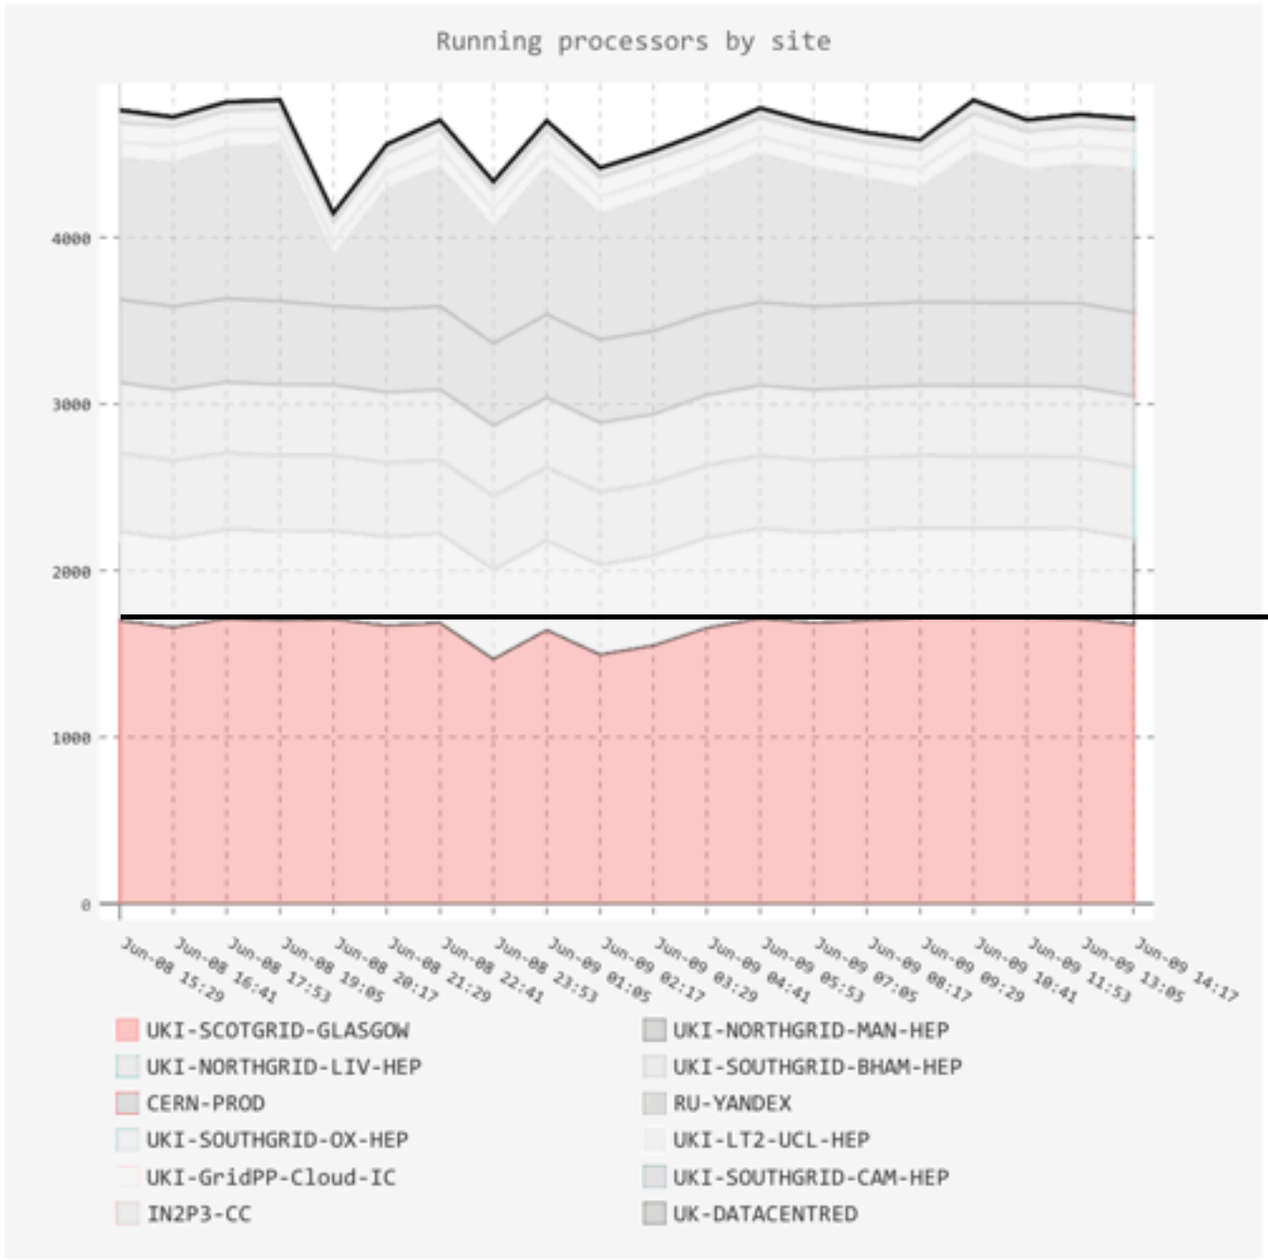
Running processors by site
4000
3000
2000
1000
0
1760
Jun-08 15:29
Jun-08 16:41
Jun-08 17:53
Jun-08 19:05
Jun-08 20:17
Jun-08 21:29
Jun-08 22:41
Jun-08 23:53
Jun-09 01:05
Jun-09 02:17
Jun-09 03:29
Jun-09 04:41
Jun-09 05:53
Jun-09 07:05
Jun-09 08:17
Jun-09 09:29
Jun-09 10:41
Jun-09 11:53
Jun-09 13:05
Jun-09 14:17
UKI-SCOTGRID-GLASGOW
UKI-NORTHGRID-LIV-HEP
CERN-PROD
UKI-SOUTHGRID-OX-HEP
UKI-GridPP-Cloud-IC
IN2P3-CC
UKI-NORTHGRID-MAN-HEP
UKI-SOUTHGRID-BHAM-HEP
RU-YANDEX
UKI-LT2-UCL-HEP
UKI-SOUTHGRID-CAM-HEP
UK-DATACENTRED

VAC

LHCb

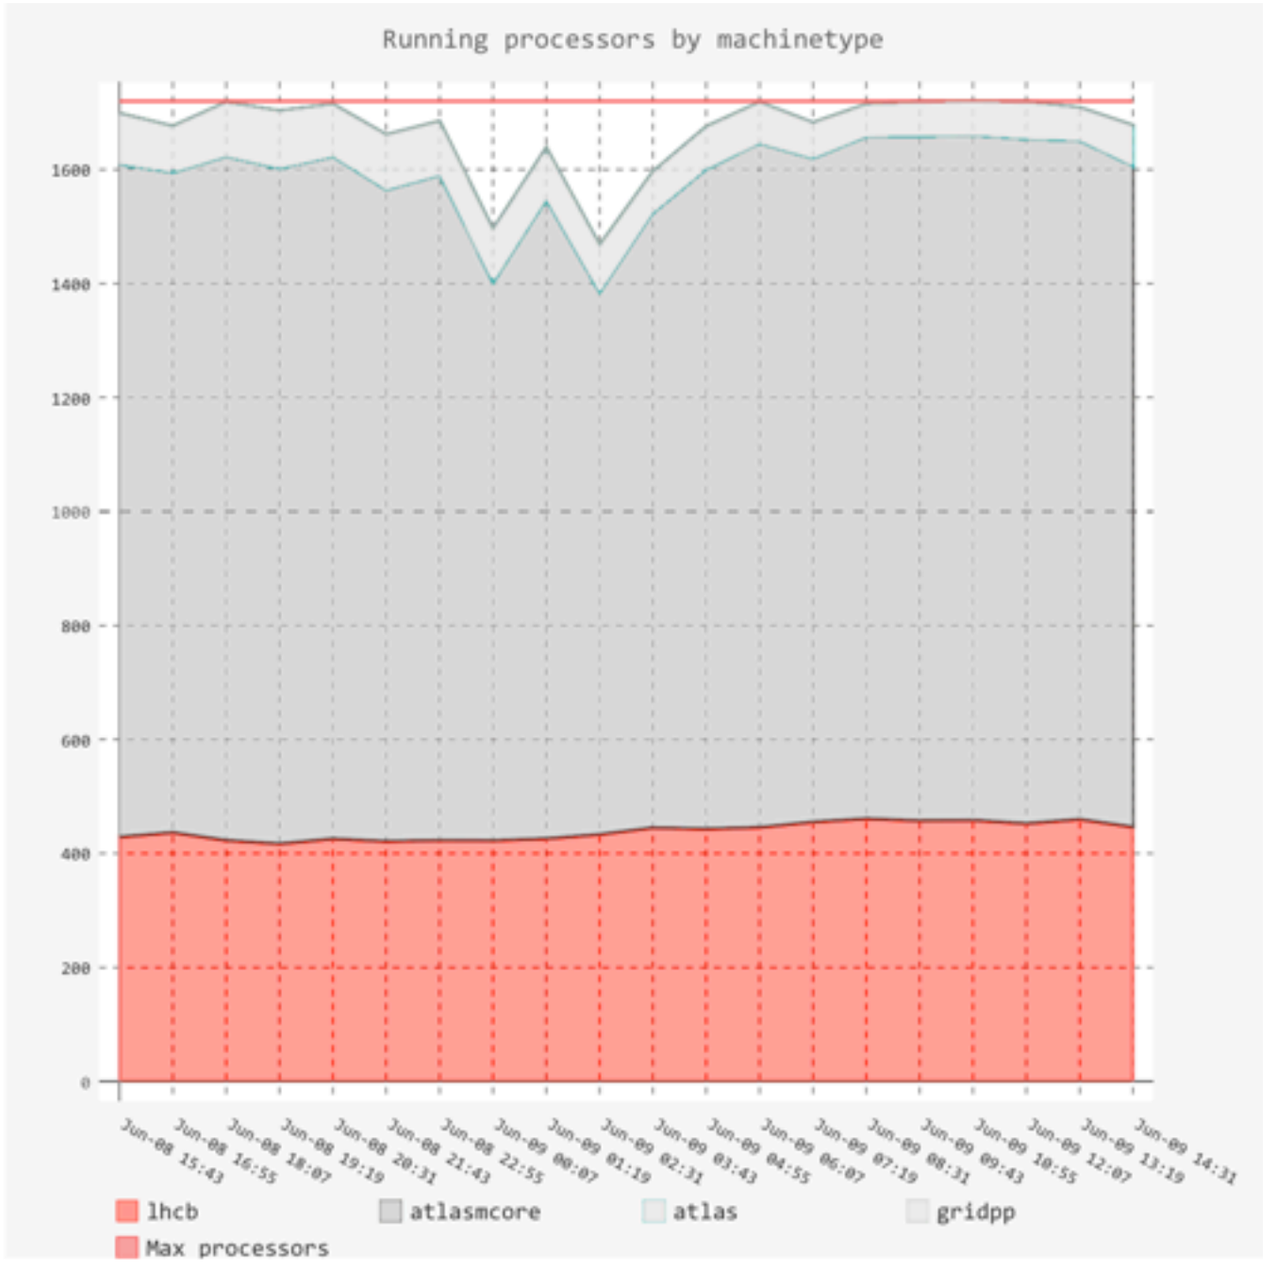
Running processors by machinetype
1600
1400
1200
1000
800
600
400
200
0
Jun-08 15:43
Jun-08 16:55
Jun-08 18:07
Jun-08 19:19
Jun-08 20:31
Jun-08 21:43
Jun-08 22:55
Jun-09 00:07
Jun-09 01:19
Jun-09 02:31
Jun-09 03:43
Jun-09 04:55
Jun-09 06:07
Jun-09 07:19
Jun-09 08:31
Jun-09 09:43
Jun-09 10:55
Jun-09 12:07
Jun-09 13:19
Jun-09 14:31
lhcb
atlasmcore
atlas
gridpp
Max processors

VAC

ATLAS Multicore

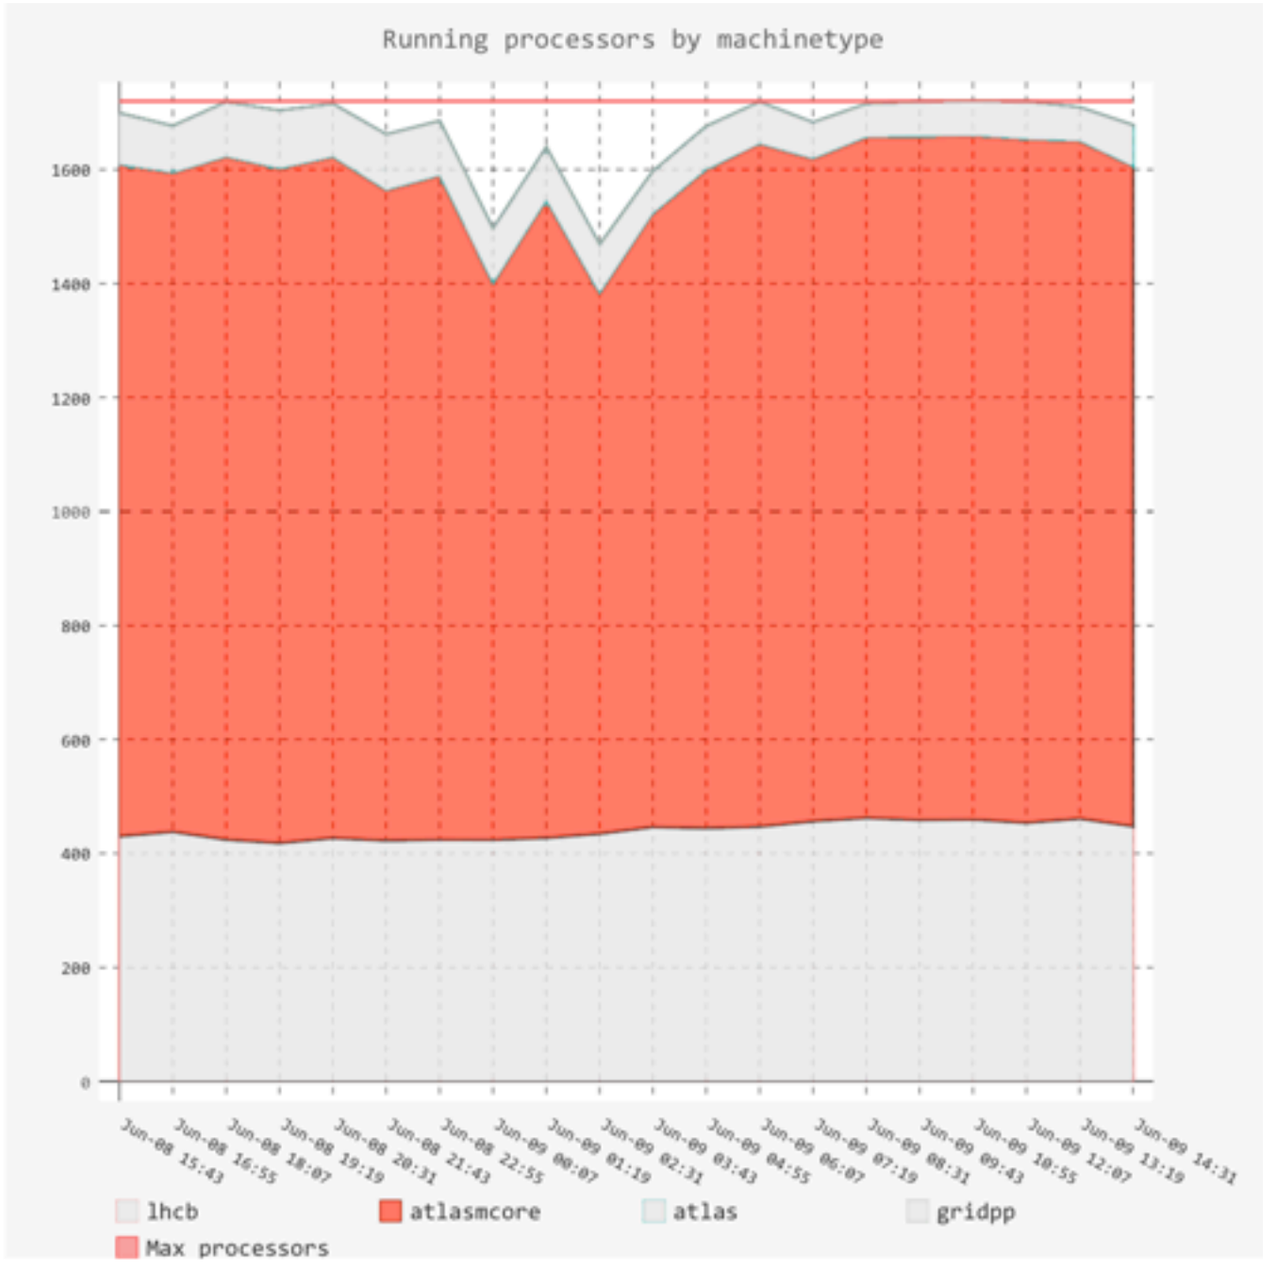

VAC

ATLAS Singlecore

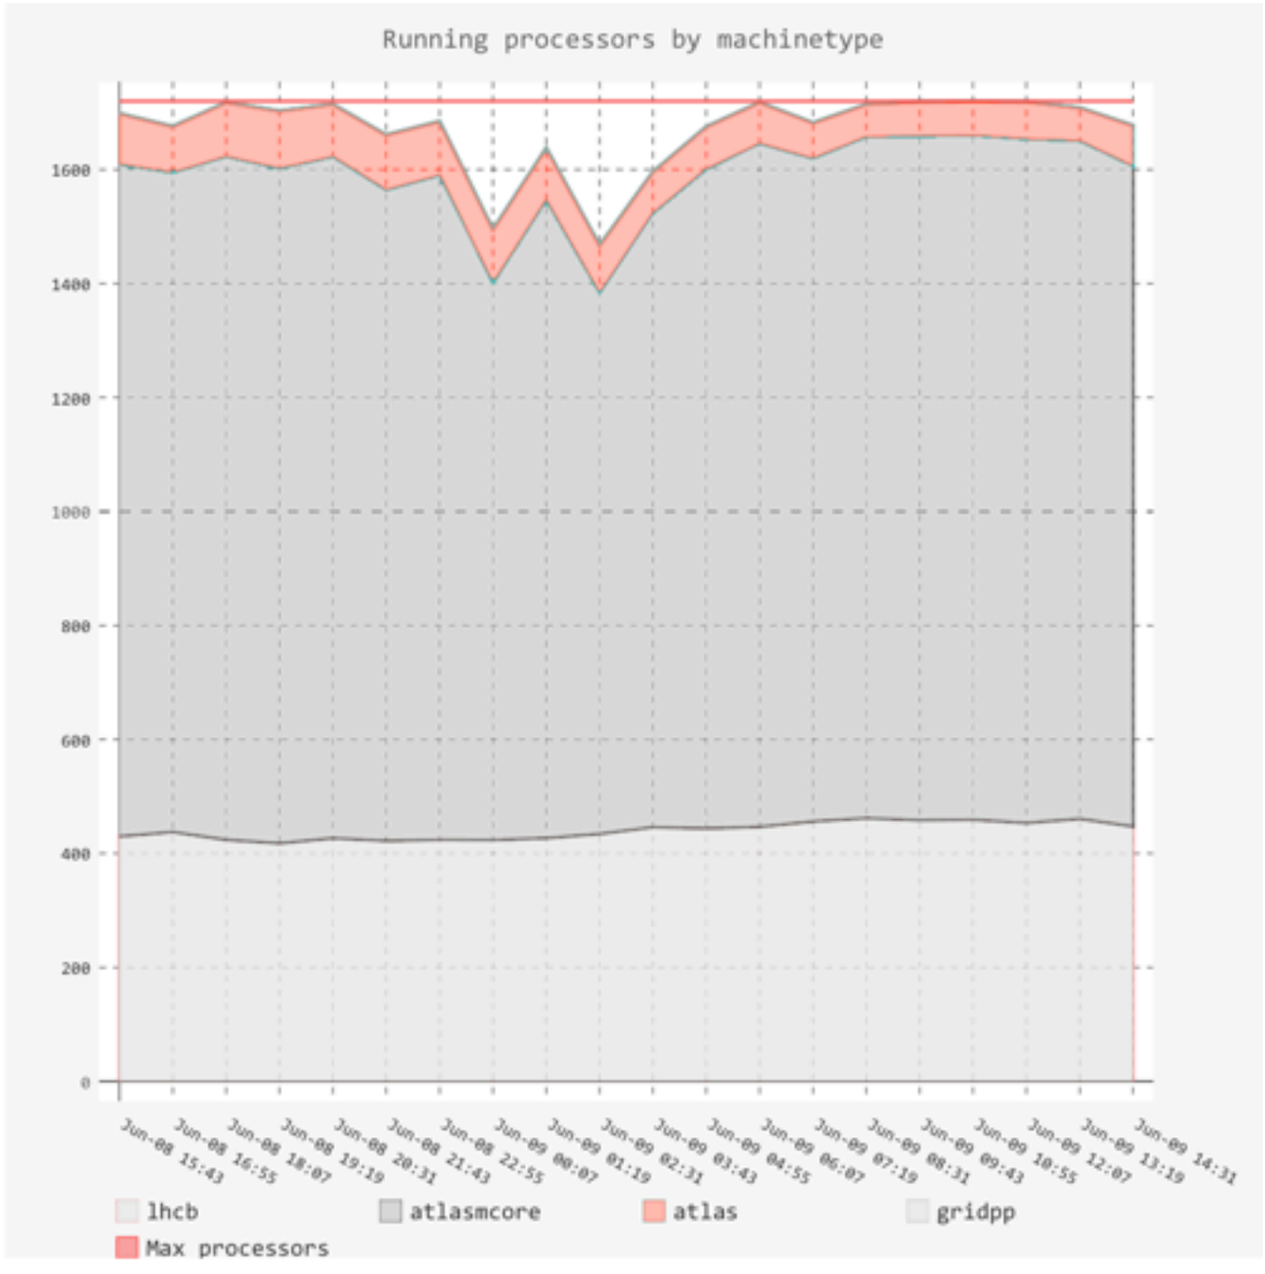

VAC (Current Issues)

- Scheduling
  - problems with scheduling single and multicore on a single VAC node
  - fairshare issues, need to optimise fizzle settings (when a new VM is spawned)
- Node failure
  - 1 to 2 node randomly reboot, may be related to KVM issues (under investigation)
- Visibility
  - usual problem with VM's not transparent to local monitoring, relying on vacmon.gridpp.ac.uk

VAC (Future Work)

- Connecting VAC to local condor pool
  - Scotgrid and Local Group
- Containers when available

Infrastructure

- Power work in 243d
- New distribution board and thermal cutout
- Proved effective after power surge last week
- Example of extremely successful estates project
  - Good communication
  - Clear timelines
  - Well realised goals

Infrastructure

- Water, water everywhere
  - *or,* The continuing adventures of The Jolly 141
- 141 uses drip trays to protect against A/C overflows (seen in early days)
- Following heavy rain and a blocked drain these flooded

Infrastructure

Infrastructure

Infrastructure

Infrastructure

Infrastructure

Infrastructure

Infrastructure

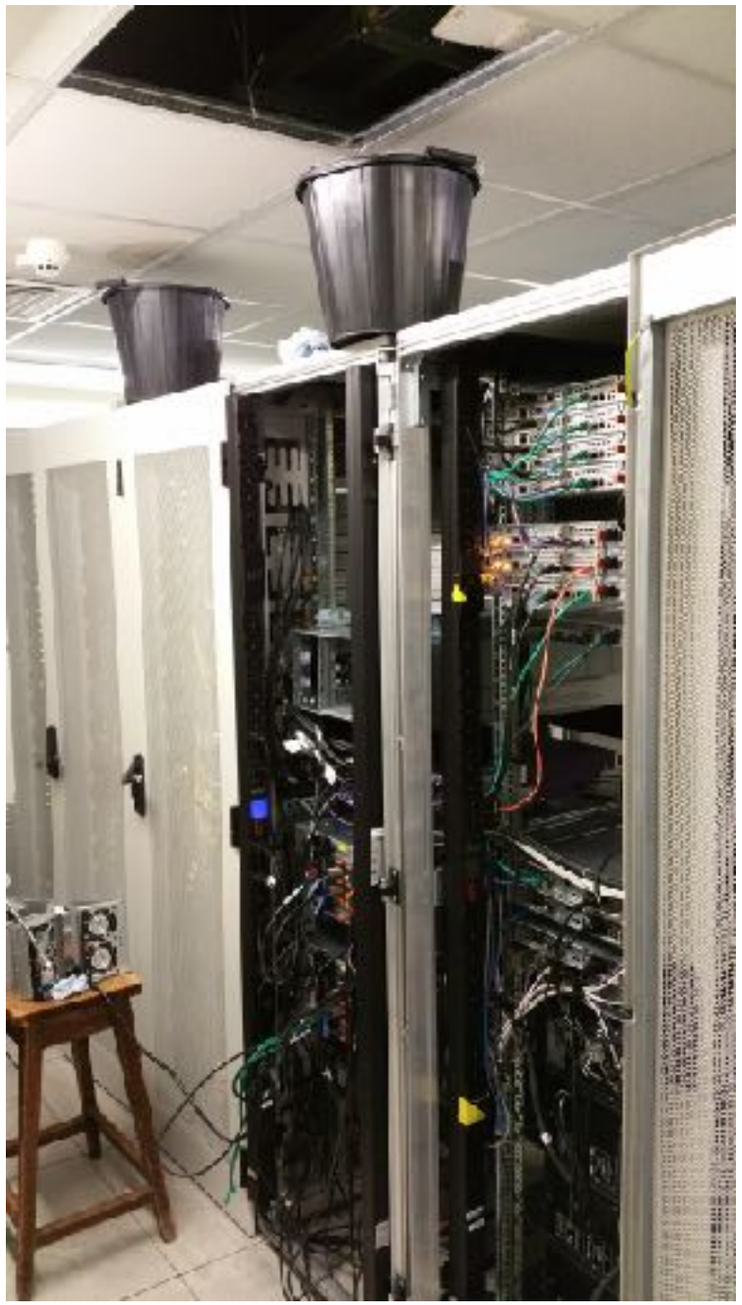

Data Centre update

- Kelvin Data Annex
- Planning permission in progress
- CapEx meeting soon
- Currently November/December timescale
- The answer to most questions is either
  - "Will be done for the DC" *or*
  - "Not before the DC"

Data Centre planning

- Procurement
- Networking
- Physical transition
- Logical transition (including SE design)
- Opportunity to revisit practices and policies

Data Centre planning

Networking

- Core network for KDA procured
  - Lenovo 40 Gb/s
- Continue to work with IT services on IPv6 provision
  - … “Not before the DC”

Networking

New networking

New lift trolley

The way forward (Principles)

- Distributed configuration
- Reduce single points of failure
- Policy and process based configuration
- Simplification and automation
- Version control and documentation

Provisioning

- Puppet → Ansible: orchestration of services
- VM hosting (KVM+scripts → *oVirt*)
- CI
- Workflow

Research

- Security Platforms
- Containers (Docker + Singularity)
- Storage caching
- Local group HTCondor
- Integration of Grid/Local group

Lightning topics

- Netbox
  - https://github.com/digitalocean/netbox
- Monitoring scripts
  - https://bitbucket.org/account/user/ukiscotgridglasgow/projects/MON
- ARGUS Ansible role
  - To be posted following sanitisation

Lightning topics

University of Glasgow

Rack L4

Created April 9, 2017 · Updated 2 months ago

Rack	
Site	Kelvin Building > 141
Group	None
Facility ID	N/A
Tenant	scotgrid
Role	None
Type	None
Width	19 inches
Height	42U (descending)
Devices	18

Non-Racked Devices

None

+ Add a non-racked device

Comments

None

Images

None

+ Attach an image

Reservations

None

+ Add a reservation

‹ Previous Rack | › Next Rack | ✎ Edit this rack | 🗑 Delete this rack

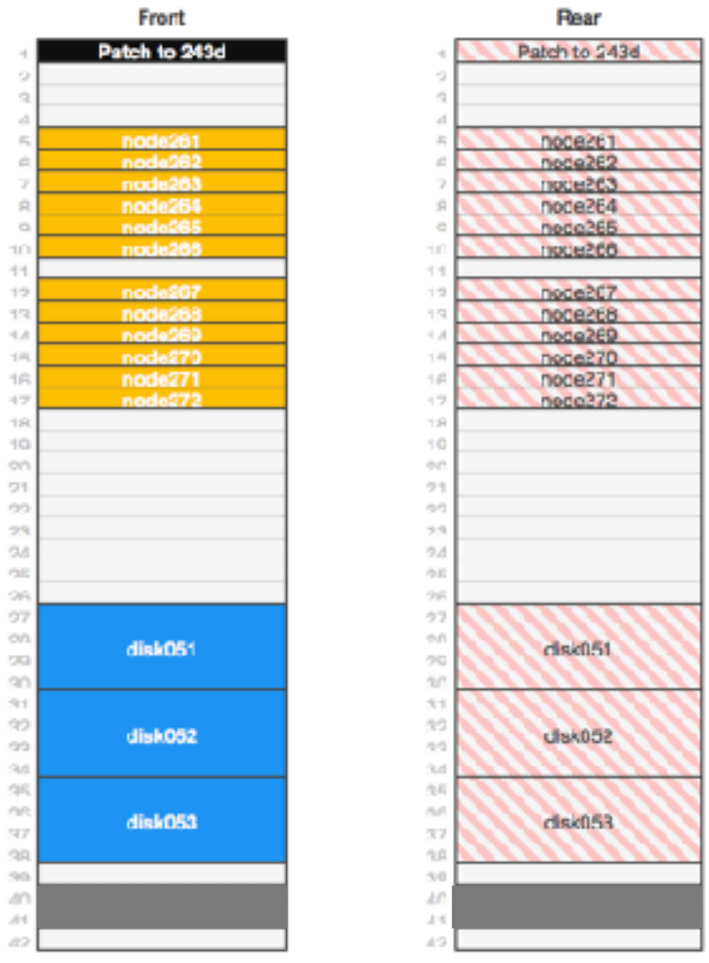

GridPP UK Computing for Particle Physics