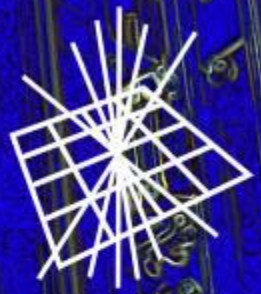

GridPP

UK Computing for Particle Physics

RAL Tier1 Status / Operations

GridPP37- Ambleside

Gareth Smith, STFC-RAL

Some Figures

For year April 2015 -> March 2016.

Availabilities:

- OPS: 99.5%
- Alice: 100%
- Atlas: 98.6%
- CMS: 97.5%
- LHCb: 98.8%

- Farm Occupancy: 91%
- Job Efficiency: 84%

Tier1 Hardware

- CPU: ~140k HS06 (~14.8k cores)
  - FY 15/16: additional ~106k HS06 in test
  - E5-2640v3 and E5-2630v3 (Supermicro (XMA), HPE)
    - 4 sleds in a 2U chassis
    - Back to 1G NICs
    - 4GB per virtual core; 2 hard disks (RAID0) per server.
  - XMA batch in use. HPE batch expected in service in a couple of weeks.
- Storage: ~16.5 PB disk
  - FY 15/16: additional ~13.3PB raw
  - CEPH spec
    - 2 x CPU, 64GB RAM, 36 x 6TB HDD, SAS HBA, 2 x 2 port 10G NICs
    - 4 year warranty.
- Tape: 10k slot SL8500 (one of two in system)
  - 44PB (T10KC/D)
  - We have 15 'D' drives. 6 more being ordered.
  - Migrations to D-only started.
    - Atlas (~6PB) recently completed. LHCb (~3PB) to follow.

Disk Servers

- Castor Disk Servers – RAID 6 data array:
  - Recent batches with the LSI Megaraid cards
  - In past (some years ago) we have generally not updated firmware unless problem. However now:
    - Make sure we have new version tested in preprod.
    - Update opportunistically – but not many "opportunities"
  - …. older servers seem to require firmware updates.
  - Not all vendors disk drives seem to show same rate of failures with age.
- Newest disk servers (for CEPH) - without RAID cards:
  - Need to re-work Nagios monitoring as we currently monitor the RAID cards to get disk failure information.
  - New test found using SMART (smartmon) – work has been needed to tailor for what we want.

Tier1 Network

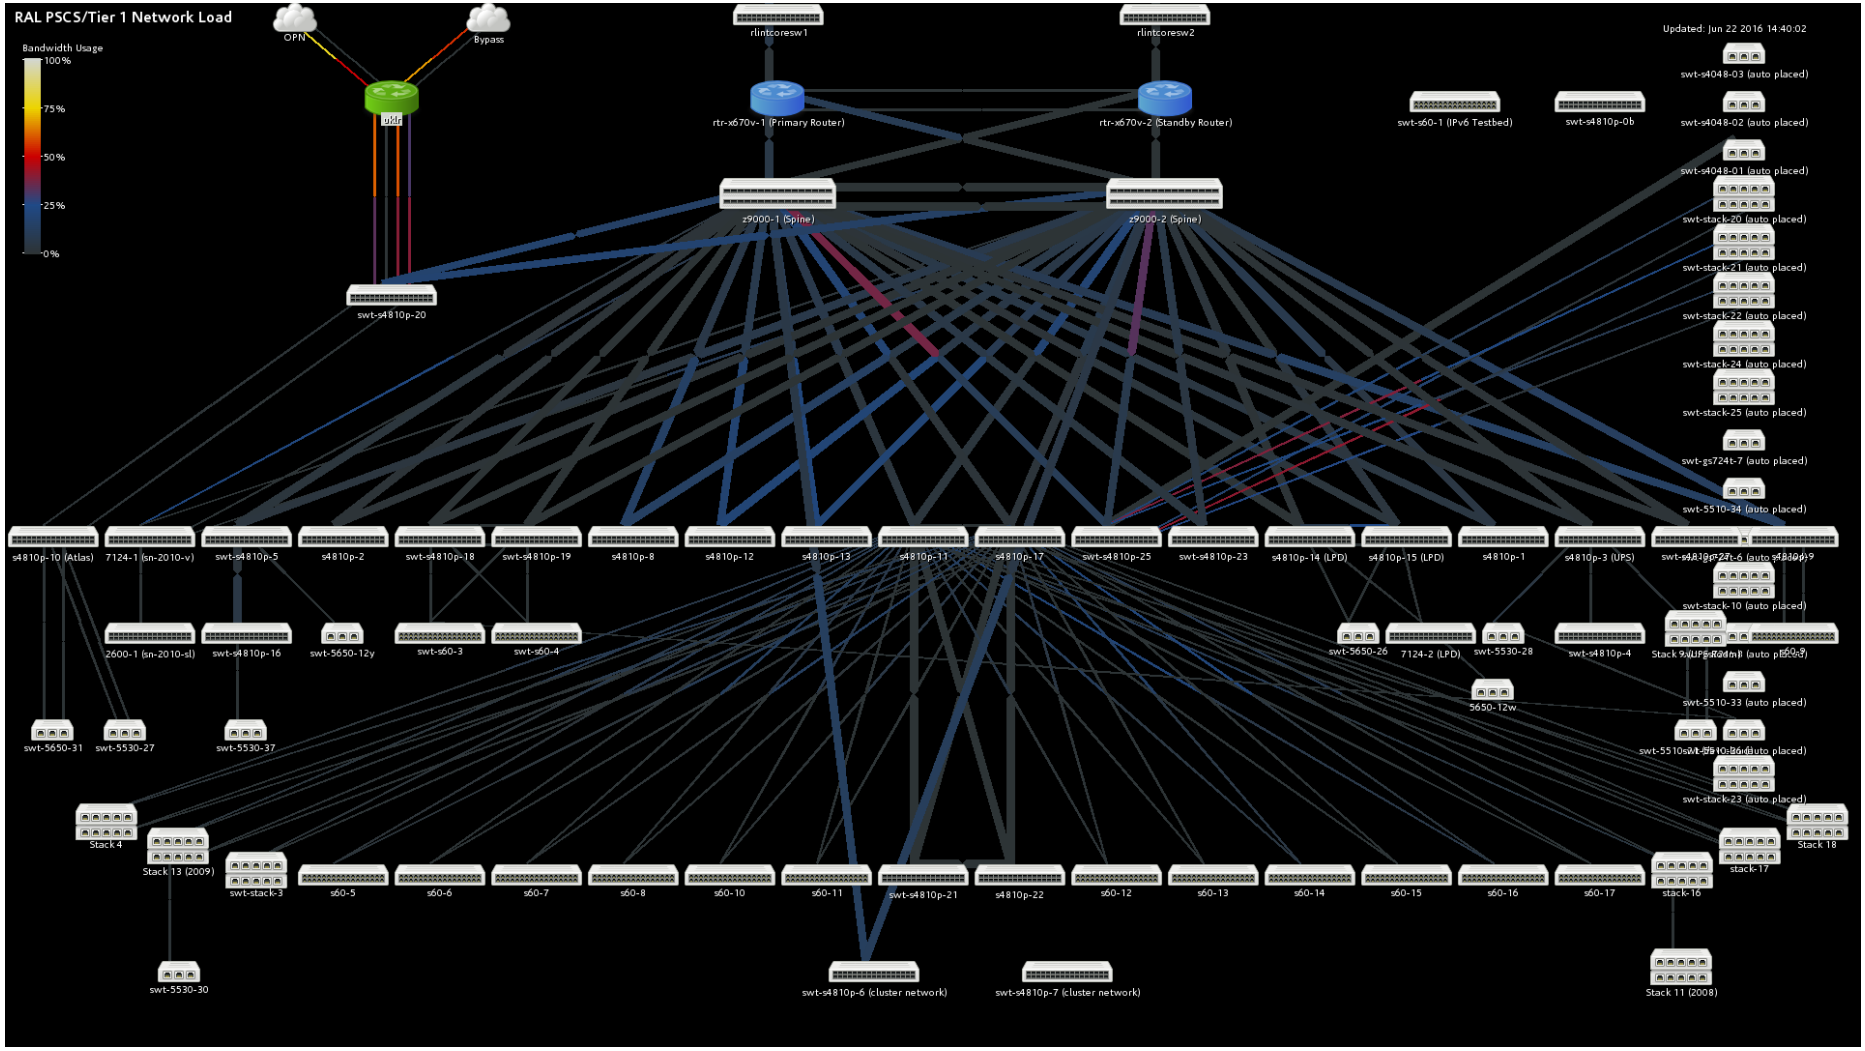

Tier1 Network

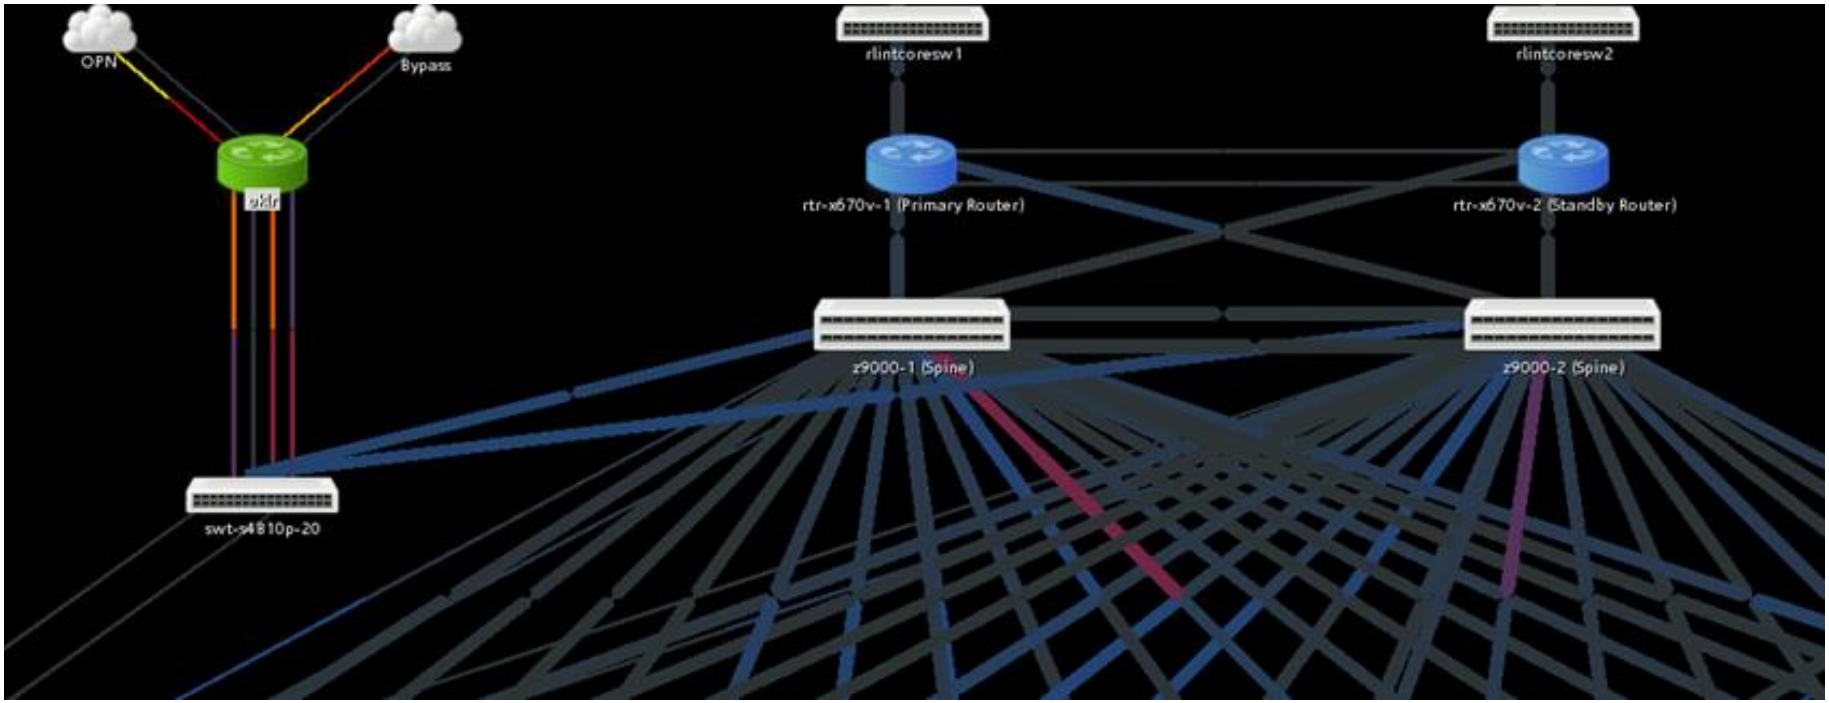

OPN Link - Last year & month

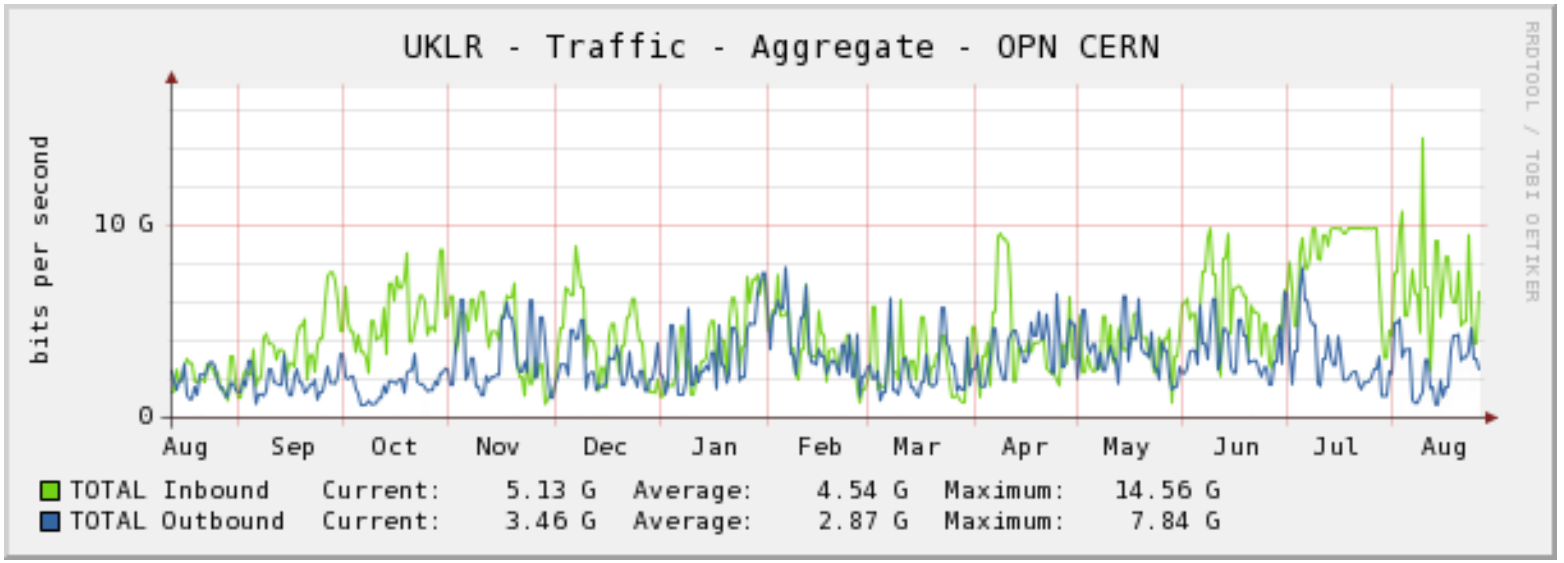

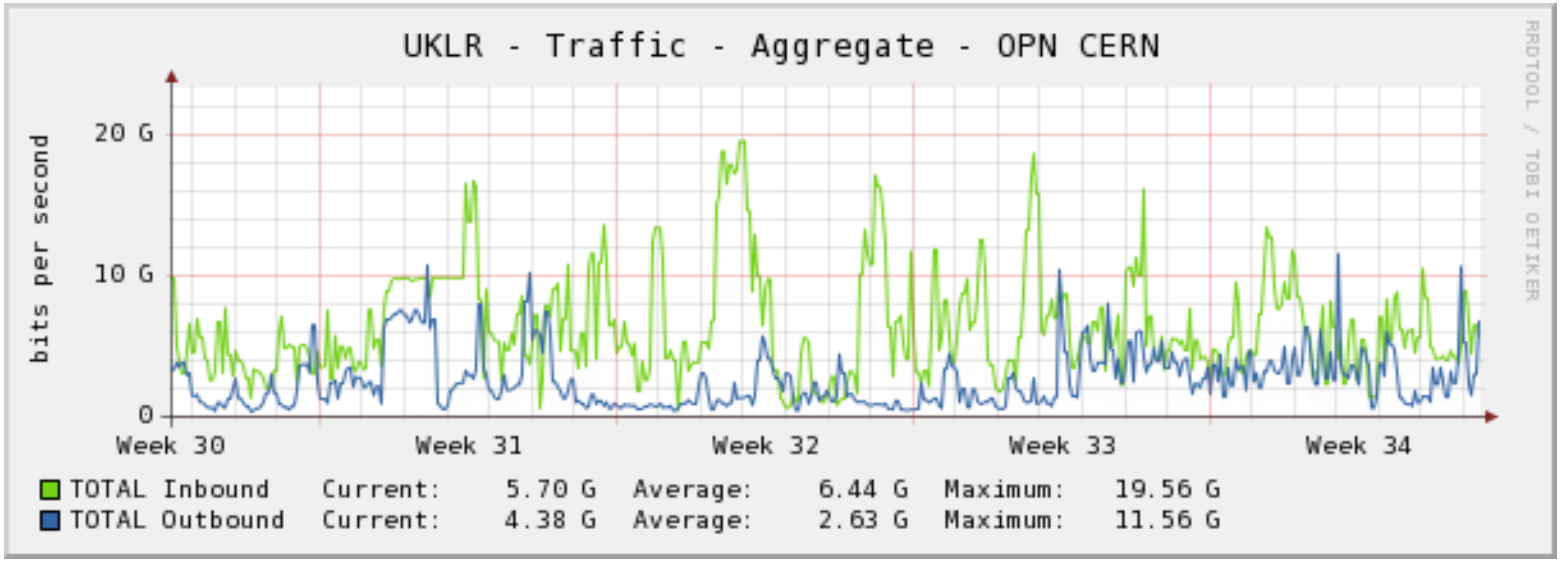

Networking

- Replace ancient UKLR router
  - Stacked pair of Dell Force10 S4810P
  - Will provide 40Gb/s pipe to border
  - Direct connection to Tier1 core
  - Landing point for OPN
  - Enables IPv6 traffic through this route.
- Expansion capability of Tier1 network topology limited – need to look at moving from L2 to L3 mesh or other architectures
- *Both 'primary' and 'backup' OPN links now brought into use to increase bandwidth to 20Gbit.*

Castor

Currently running v2.1.14

- Stable running
- Improvements in data throughput from disk thanks to optimization of IO scheduling.
- Starting to see limitations of database throughput. Database refresh in pipeline. A period with:
  - High transaction rate on Atlas stager database
  - Delays applying updates to the standby databases

• v2.1.15 in test

- Tests ongoing. Have had series of problems (database performance, problems with test database itself, configuration issues). Now looking encouraging – almost ready to schedule update.
- Updating (SL5) SRMs delayed until this is done.

Summary Plan:

- Move to ECHO (CEPH) for disk.
- Continue with Castor for tape.

Tape Library Problems

During May/June we had appreciable problems with the tape system.

- Two SL8500 Libraries each with 10,000 tape slots. Next to each other with common control system which provides the interface between the libraries themselves and the tape servers.

Three overlapping problems:

- 1) Significant hardware problem – fix confused by an "EBF" (Early Board Failure).
- 2) Faulty replacement handbot led to safety concern.
- 3) Problem with library control software hanging.

Neil Gaiman. Neverwhere:

"He had noticed that events were cowards: they didn't occur singly, but instead they would run in packs and leap out at him all at once."

Virtualisation etc.

- Moving VMs from a Windows Hyper-V 2008 to Hyper-V 2012.
  We have a high availability cluster as detailed 5 hosts, 130 logical cores, 1280GB RAM total
  - Tiered EqualLogic arrays (82TB total)
    - Eql-varray2: 62TB, 7.2K SATA
    - Eql-varray3: 22TB, 10k SAS
- The re-organisation of teams within groups in SCD links us with others with VMware experience. Looking at this.
- There has been a project to investigate containerisation and management of services using Mesos. Looking at its wider role in the Tier1.

Services

- Batch system (HTCondor)
  - Works well
  - At end of last year rolled out a new method for draining worker nodes for multi-core jobs, enabling us to run pre-emptable jobs on the cores which would otherwise be idle.
  - Machine Job Features rolled out.
- ARC CEs
  - Four in production.
  - Being migrated (to Hyper-V 2012) with more disk space for logs.
- FTS
  - Two instances Production and “Test” (used by Atlas).
  - Service fronted by HAProxy load balancers.

Services – Continued

- LFC
  - In use for non-LHC VOs.
  - Recently had refresh of underlying (Oracle) database system.
- WMS:
  - Two production WMS & L&B in service.
  - We did have three WMS – rationalised to two a couple of months ago.
- CVMFS
  - Stratum 0: For EGI and UK regional VOs. Working on a new instance with more space.
  - Stratum 1: (Replicas from Stratum 0s at CERN, RAL, OSG, NIKHEF) Recently upgraded. Now a HA two-node cluster.
  - We are part of test procedure for clients. Recently installed 2.2.3-1.

Cloud

- Cloud
  - Service using OpenNebula, storage based on CEPH ('Sirius').
  - Department and wider use in STFC (e.g. ISIS), including access by LOFAR
  - Backfill to run Worker Nodes. These have 8 cores, with typically 30 – 50 running at any time. Have been doing so for a year or so now.
  - Development of OpenStack Cloud underway.