

RAL PPD Site Report

Chris Brew

Ian Loader



The Department

- ~ 70 Staff
 - Joint Appointees, Visitors, Staff on LTA at CERN, UK Liaison Office at CERN
 - 3 Main LHC Experiments
 - Numerous “Small” Experiments
 - Hosting Department for MICE
 - Remit to provide support for UK Particle Physics
 - GridPP Tier 2 site



The Team

- Team of 5:
 - Chris Brew (Management, Linux)
 - Kevin Dunford (Windows)
 - Ian Loader (Linux, Web Servers)
 - Will Pilcher (Windows)
 - Deepak Jindal (YinI, Windows + Web Devel)
- But not:
 - Federico Melaccio – Left in April

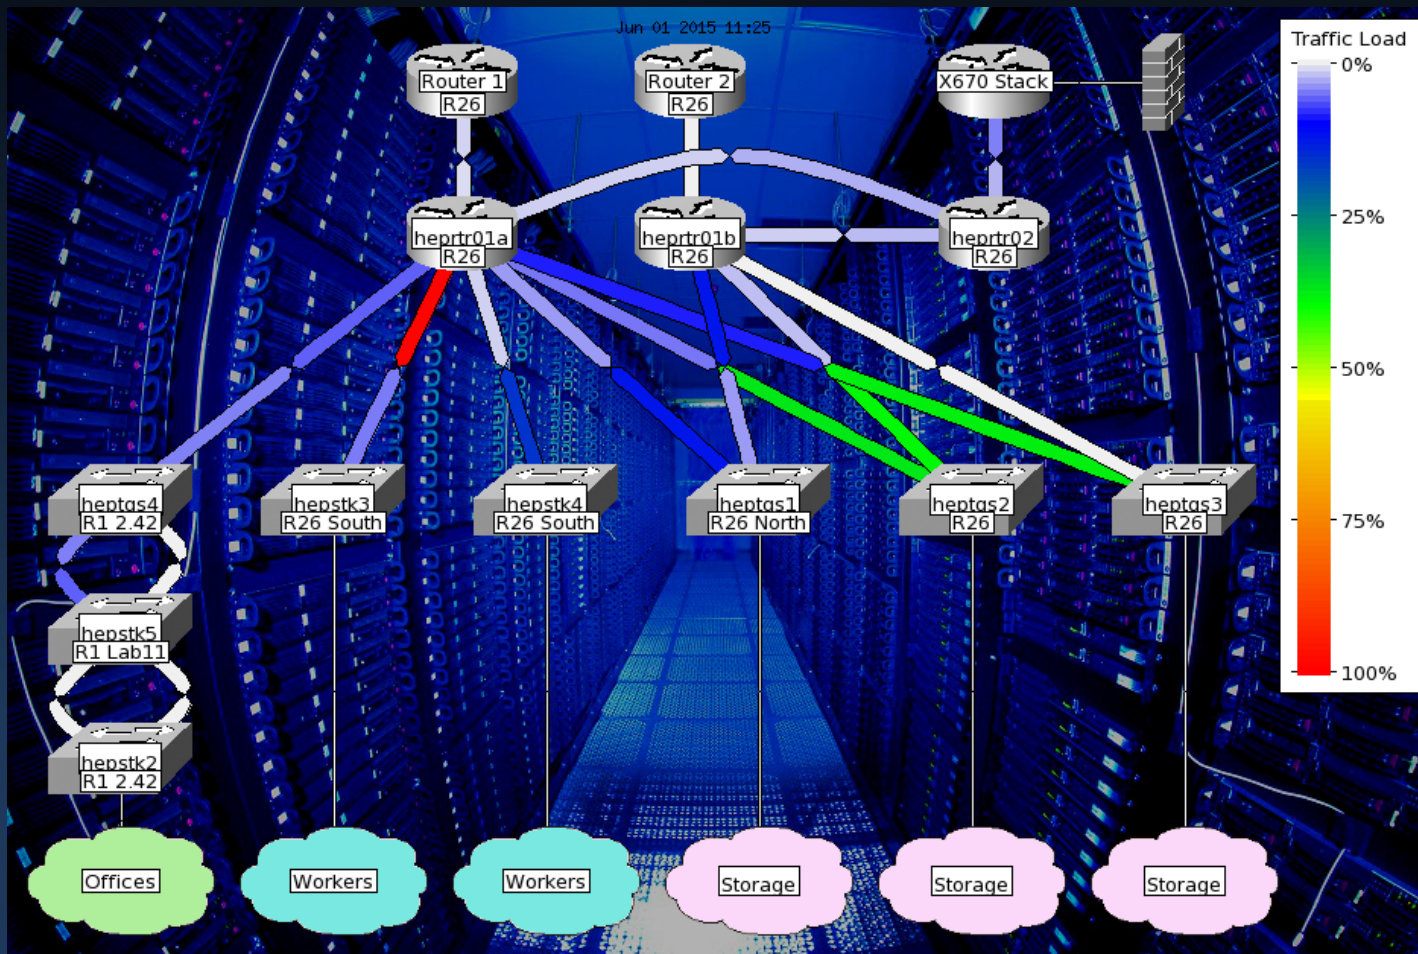


Core infrastructure - Rooms

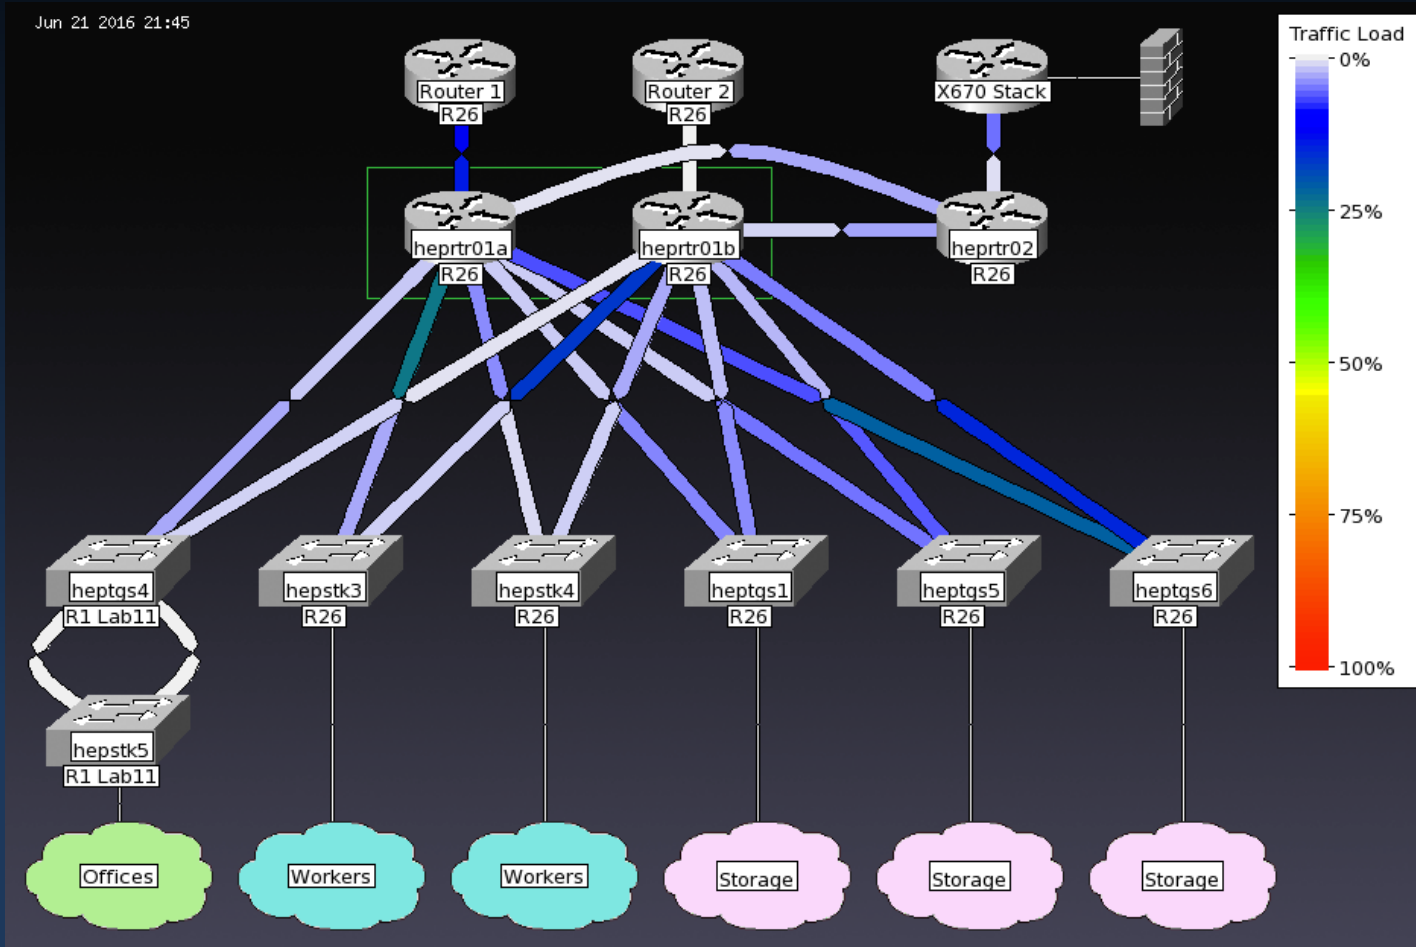
- 2 Machine Rooms
 - Remote – 18 Racks in the Atlas Centre
 - Core Network, Disk, Workers, Xrootd Reverse Proxies, PerfSONAR, HyperVisors
 - Local – 6 Racks in the basement
 - Everything else, including Windows servers, Interactive Linux, NFS, Mercury, bits, bobs, and whatsits.



Core Infrastructure - Network



Core Infrastructure - Network



Core Infrastructure - HyperVisors

- Windows 2012r2 Failover Cluster
 - Windows Server 2008 to 2012 “transition”
 - Grown from 3 to 6 HyperVisors
 - Grown from 1TB to 6TB Network Attached Storage
 - Now Hosting 70 VMs
 - Testing PXE booting Gen2 CentOS7 VMs
 - Old Hardware recycled into Dev Cluster
 - Testing SCVMM



Tier 2

- CMS Tier 2, Atlas T2D, and LHCb T2D
- Condor Batch System
 - 235 Nodes, 4,304 Cores, 46,000 HS06
 - Multicore Enabled, used by CMS and Atlas
 - Recently switched from condor_defragd to fallow + local mods
- dCache Storage
 - Migrated headnode(s) to CentOS7, dCache 2.13
 - 56 Storage Servers (SL6, dCache 2.13), 2.6PB
 - Will retire 25 Servers (500TB) and introduce 3 (~480TB)



Tier 2 Oddities


- (mostly) redundant Argus servers behind DNS alias
 - Configured with YAİM + Puppet overlay
 - Certs with Subject Alt Name
 - Relies on shared gridmap dir
- Xrootd Reverse Proxies in front of dCache
 - Stock xrootd integrates to argus and handles N2N/TFC then proxies data from dCache xroot doors



Odds and Ends

- Puppet
 - 21 local modules, 600 classes, 16,000 lines
 - PuppetDB, PuppetBoard. ENC real soon now
- Monitoring
 - Usual mix of Nagios, Ganglia, Cacti, and application specific pages





Mercury

Fast Analysis Cluster

- Pilot Project
- 5 node BeeGFS Cluster with 20TB of shared storage
- Each node 2 Intel E5-2430 v2 (6+6@2.5GHz) and 64GB RAM
- Intended for data intensive work



Accelerated Computing

- Have a number of Accelerated Computing Platforms for testing and development work.
 - DUAL NVIDIA Tesla K4
 - Xeon Phi co-processor
 - Just added an FPGA (Bittware A10PL4)



Roadmap - ELK



elasticsearch



logstash




kibana



Roadmap - OS



CentOS
7



Roadmap - Dashboards

Dashing 



Roadmap - Dashboards



Roadmap - FreeIPA



Roadmap - Tier 2 Evolution

