RAL PPD Site Report

HEPSYSMAN Summer 2017

13-15 June 2017

Ian Loader

Chris Brew

Science & Technology Facilities Council

The Department

- ~ 70 Staff
  - Joint Appointees, Visitors, Staff on LTA at CERN, UK Liaison Office at CERN
  - 3 Main LHC Experiments – Atlas, CMS, LHCb
  - Numerous “Small” Experiments e.g. Dark Matter, DUNE, T2K, Hyper-K, EDM
  - 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)
  - Ian Loader (Linux)
  - Kevin Dunford (Windows)
  - Will Pilcher (Windows)
  - Patrick Batuka (YinI, Windows)

Core infrastructure - Rooms

2 Machine Rooms (22 Racks)

Remote – 16 Racks in the Atlas Centre

- Core Network, Disk, Workers, Xrootd, Reverse Proxies, PerfSONAR, HyperVisors etc.

Local – 6 Racks in the R1 basement

- Everything else, including Windows servers, Interactive Linux, NFS, Mercury, bits & bobs

Core Infrastructure - Network

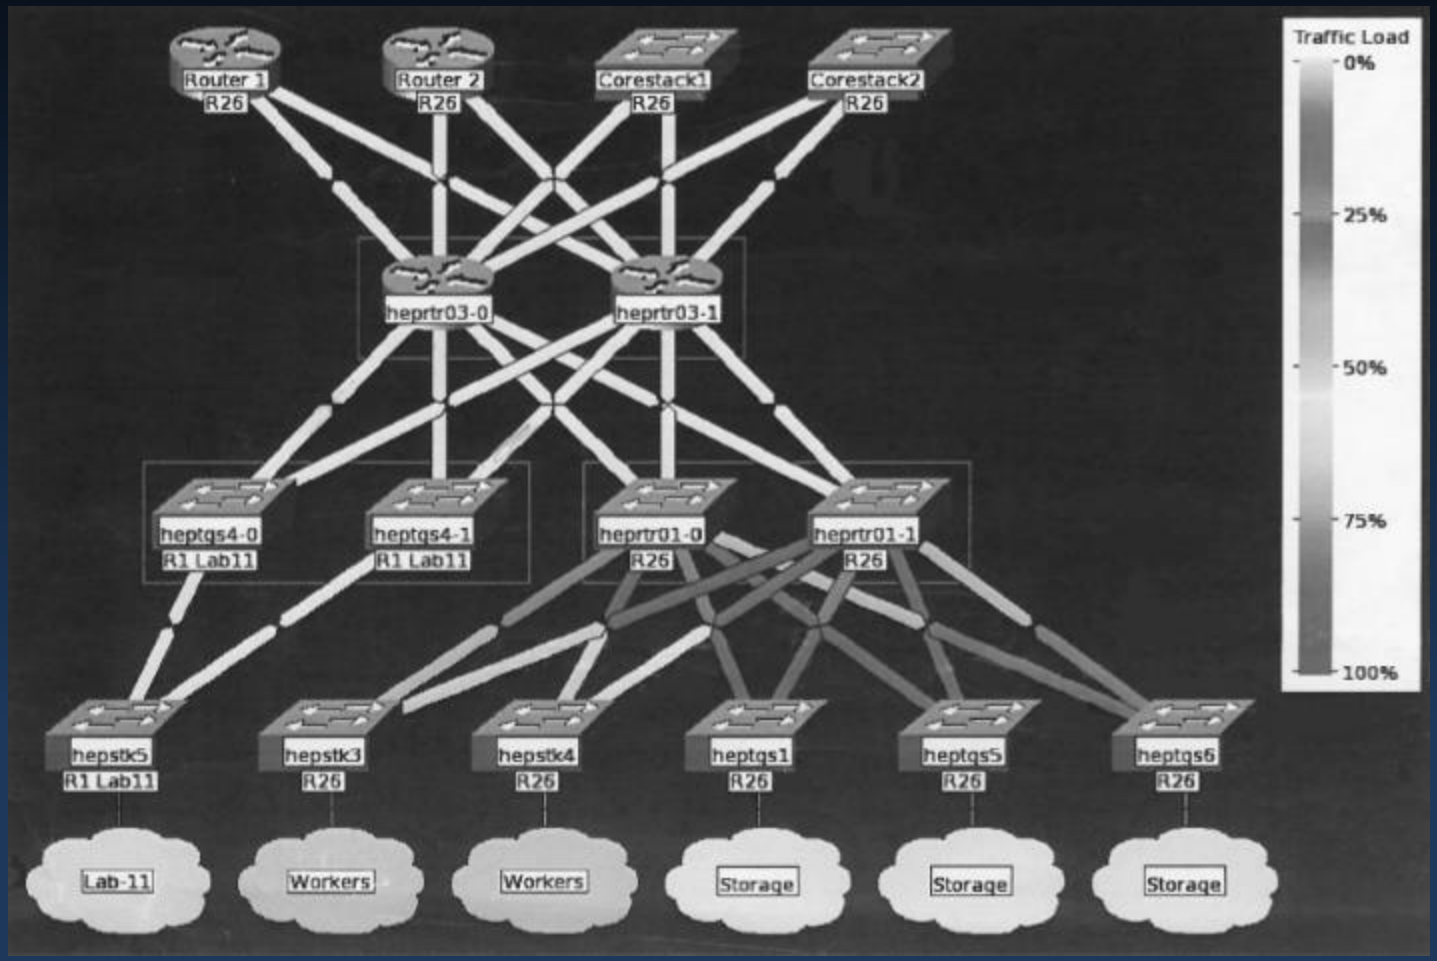

Core Infrastructure - HyperVisors

- Windows 2012r2 Failover Cluster
  - 8 Hyper-V Servers (6 Production, 2 Test)
  - 6TB Network Attached Storage
  - Now Hosting 77 VMs
  - Deploying Windows 10 by PXE and 3rd party S/W via SCCM
  - About to build test domain of 2016 windows servers for testing IPV6 (dual stack)
  - Sophos, OpenVAS system security scans, Sophos InterceptX (Ransom ware protection)

How we maintain it - *Puppet*

- >400 machines
- 22 local modules, 923 classes, 31,000 lines

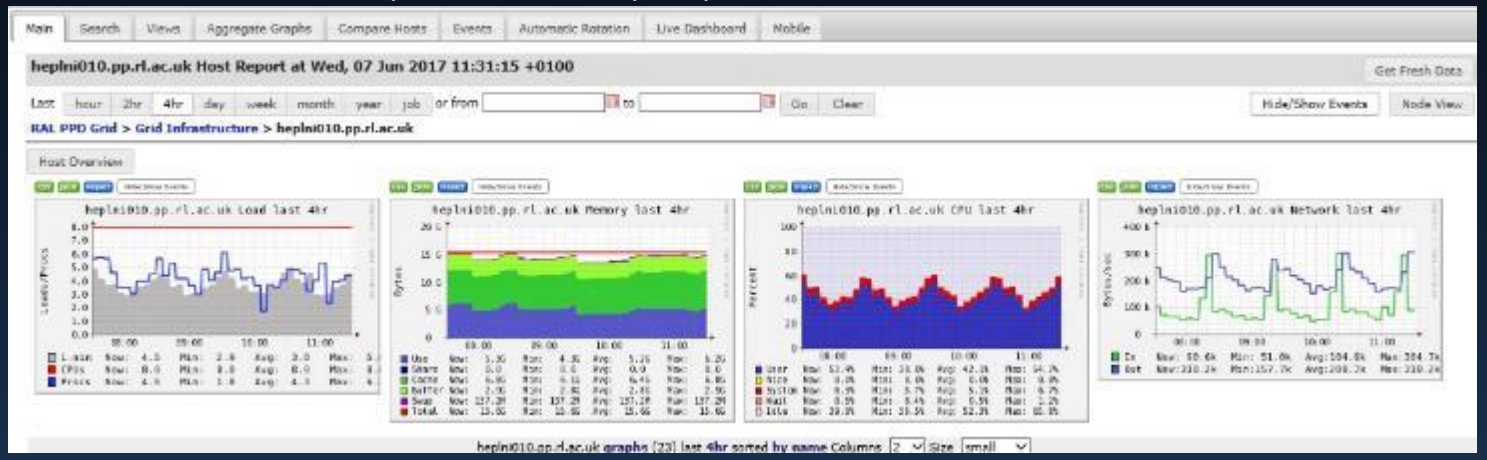

How we maintain it – Puppet+

- PuppetDB
- PuppetBoard
- PuppetENC
  - Home built, simple DB + Script + web gui

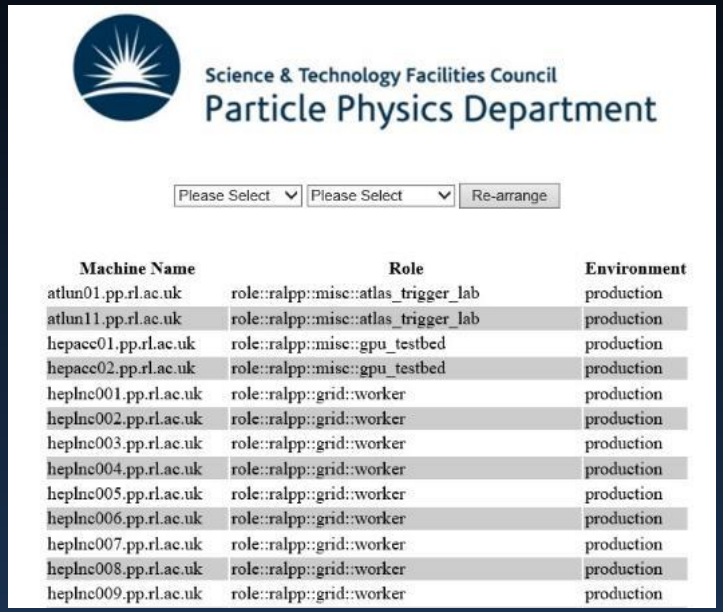

Machine Name	Role	Environment
atlun01.pp.rl.ac.uk	role::ralpp::misc::atlas_trigger_lab	production
atlun11.pp.rl.ac.uk	role::ralpp::misc::atlas_trigger_lab	production
hepacc01.pp.rl.ac.uk	role::ralpp::misc::gpu_testbed	production
hepacc02.pp.rl.ac.uk	role::ralpp::misc::gpu_testbed	production
heplnc001.pp.rl.ac.uk	role::ralpp::grid::worker	production
heplnc002.pp.rl.ac.uk	role::ralpp::grid::worker	production
heplnc003.pp.rl.ac.uk	role::ralpp::grid::worker	production
heplnc004.pp.rl.ac.uk	role::ralpp::grid::worker	production
heplnc005.pp.rl.ac.uk	role::ralpp::grid::worker	production
heplnc006.pp.rl.ac.uk	role::ralpp::grid::worker	production
heplnc007.pp.rl.ac.uk	role::ralpp::grid::worker	production
heplnc008.pp.rl.ac.uk	role::ralpp::grid::worker	production
heplnc009.pp.rl.ac.uk	role::ralpp::grid::worker	production

How we run it -

- Condor Batch System
  - Used by both Grid and Local Users
  - 234 Nodes, 4,476 Cores, 46,000 HS06
  - Multicore Enabled, used by CMS and Atlas
    - Switched from condor_defragd to local version of fallow + other mods
- dCache Storage
  - Migrated headnode(s) to CentOS7, dCache 2.13
  - 56 Storage Servers (33*SL6, 23*Centos-7, dCache 2.13), 2.7PB

Cluster Usage

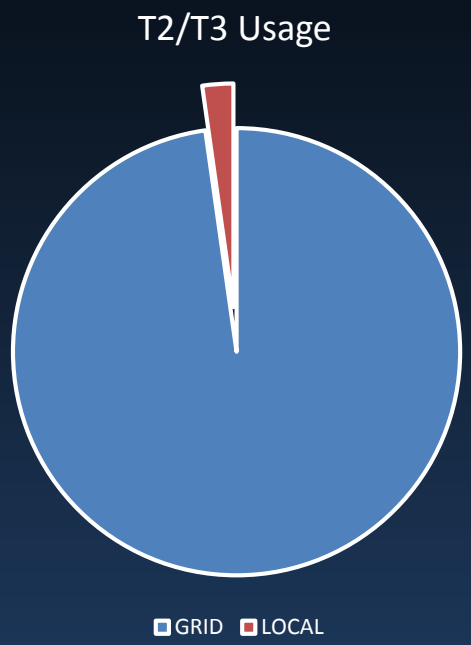

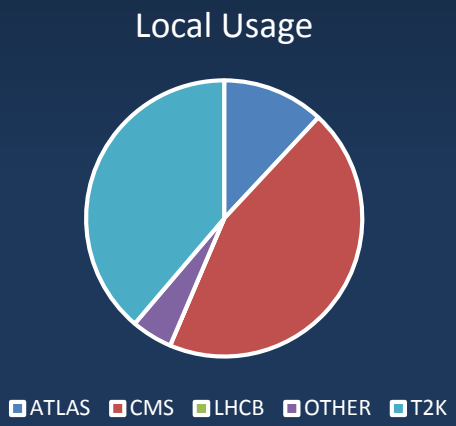

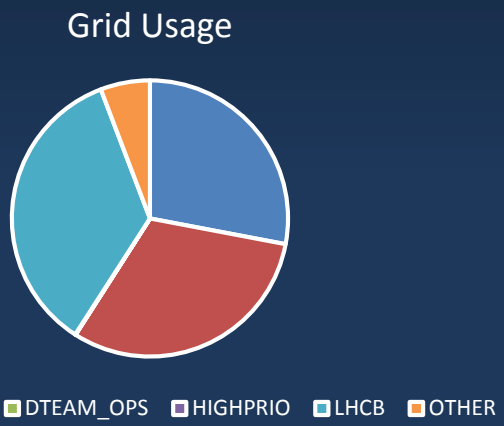

Batch system

- HTCondor

RAL-PPD Condor Pool Status

Worker Node Status

Total Nodes	Unhealthy Nodes	MultiCore Draining	Reboot Draining
234 / 234	2	28	0

Total Running Jobs

Jobs	CPUs Requested	Memory Requested
1445	1	2000
1393	1	3000
145	8	16000
20	8	24576

Total Pending Jobs

Jobs	CPUs Requested	Memory Requested
429	1	2000
239	1	3000
400	8	16000
25	8	24576

Worker	CPU Usage	MEM Usage	Status	Health	Reboot
heplnc001	12 of 16	25000 of 25000	Accepting Jobs	Ok	...
heplnc002	9 of 16	25000 of 25000	Not Available	Draining-MC	...
heplnc003	9 of 16	22000 of 25000	Not Available	Draining-MC	...
heplnc004	9 of 16	21000 of 25000	Not Available	Draining-MC	...
heplnc005	10 of 16	22000 of 22000	Accepting Jobs	Ok	...
heplnc006	9 of 16	24000 of 25000	Accepting Jobs	Ok	...
heplnc007	9 of 16	24000 of 25000	Not Available	Draining-MC	...
heplnc008	9 of 16	24000 of 25000	Not Available	Draining-MC	...
heplnc009	9 of 16	22000 of 25000	Not Available	Draining-MC	...
heplnc010	10 of 16	25000 of 25000	Not Available	Draining-MC	...
heplnc011	9 of 16	22000 of 25000	Not Available	Draining-MC	...
heplnc013	10 of 16	23000 of 25000	Not Available	Draining-MC	...
heplnc014	11 of 16	25000 of 25000	Accepting Jobs	Ok	...
heplnc015	10 of 16	25000 of 25000	Accepting Jobs	Ok	...
heplnc016	10 of 16	25000 of 25000	Not Available	Draining-MC	...
heplnc017	14 of 20	34000 of 49000	Not Available	Draining-MC	...
heplnc018	20 of 20	46000 of 49000	Accepting Jobs	Ok	...
heplnc019	14 of 20	32000 of 41000	Not Available	Draining-MC	...

Monitoring

- Ganglia
- Nagios
- Swatch
- PuppetDashboard
- ELK
- Pakiti
- Cacti

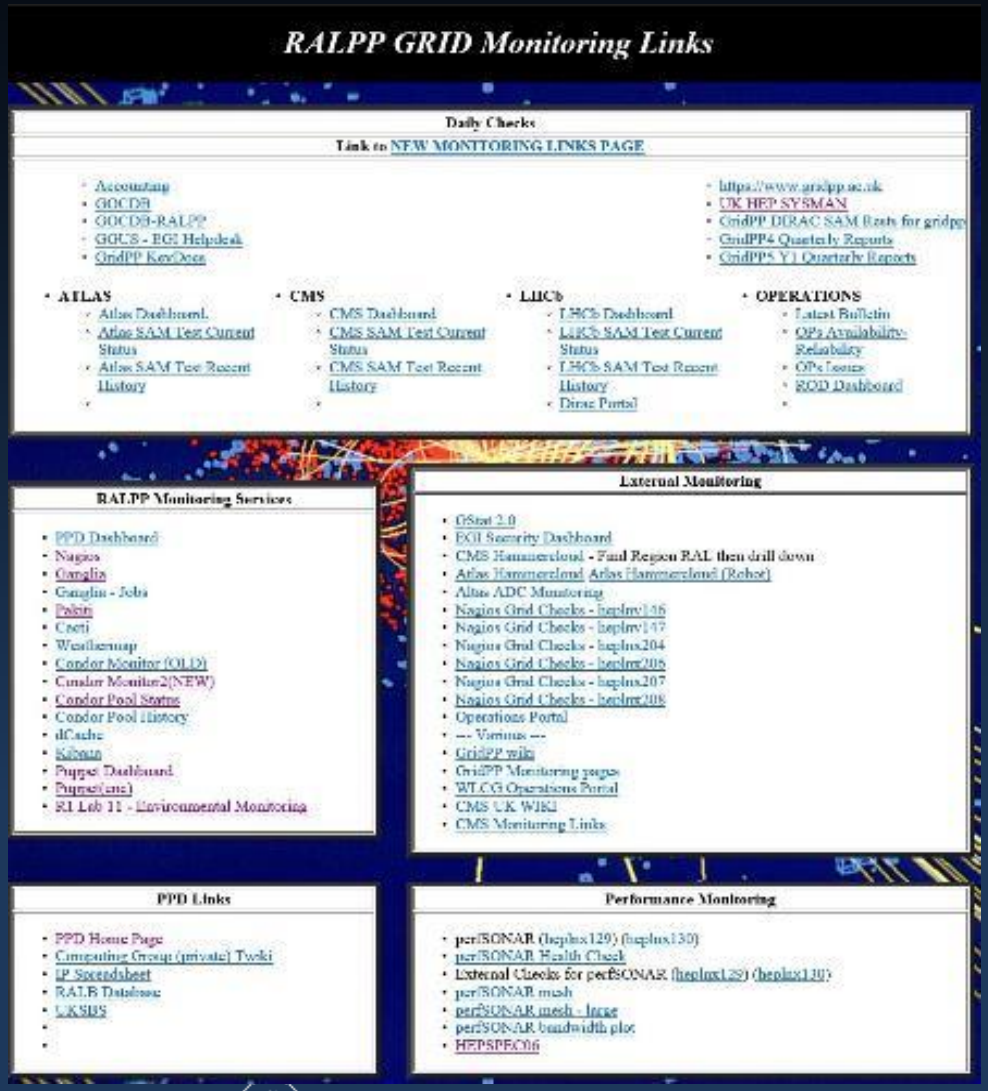

Desktops and OS's

- SL6 (.9)
- Centos-7 (.3)
- Ubuntu 16.04

- Desktop Environments:
  - Gnome
  - Mate
  - Unity

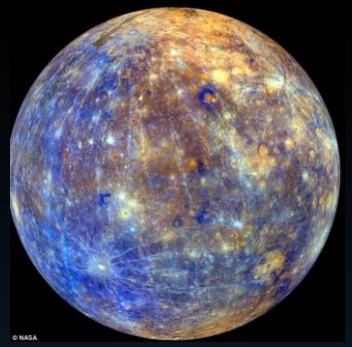

Mercury
Fast Analysis Cluster

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

Storage

- Local Storage was 2 SL5 NFS Servers for Home and Software and two very old (2006) ex-dCache NFS Disk servers for scratch
- New Home Area – VM with iSCSCI NAS
- New Software Area – NFS mount of ZFS volumes rsync'd to second server (want to look at DRBD)
- New Data Area – 400TB of old dCache space, still in dCache, NFS4 mounted, 2 copies of files

Accelerated Computing

Early Access to New Multi-core Processors, Parallel Programming Models and Tools

- Have a number of Accelerated Computing Platforms for testing and development work using CUDA, OpenCL or OpenACC

- **HEPACC02 -** Tech Specs:
  - DUAL E5-2680v4 - Intel 14 Core Xeon 2.4GHz
  - 128 GB DDR4 Memory
  - DUAL NVIDIA TITAN X -GE Force PASCAL GPUs each with
    - 12 GB GDDR5X memory, 10Gpbs
    - 3584 CUDA Computing Cores (1.5 GHz) delivering around 11Tflops

- **HEPACC01 -** Tech Specs:
  - DUAL Intel Xeon E5-2670 2.6GHz 8 core processors
  - 64 GB Memory
  - DUAL NVIDIA Tesla K40M GPU Computing Cards each with
    - 12 GB memory
    - 2880 CUDA Computing Cores (745 MHz) delivering around 4.2Tflops

- **HEPGPUW101 -** Tech Specs:
  - DUAL Intel Xeon E5620 2.4 GHz QUAD Core 64Bit Processors
  - 24 GB Memory
  - NVIDIA Tesla C2050 'Fermi' 3GB GPU Computing Card
    - 448 GPU Computing Cores (1.15 GHz)
  - NVIDIA Quadro 600 1GB Workstation Graphics Card
    - 96 GPU Cores for Visualisation and Rendering
  - 1TB Main System Hard Disk (SATA2)
- **HEPGPUW104 -**
  - DUAL AMD Opteron 6276 2.3 GHz SIXTEEN Core 64Bit Processors
  - 64 GB Memory
  - AMD Firepro V9800 4GB GPU Computing Card
    - 1600 GPU Computing Cores (850 MHz)
  - 1TB Main System Hard Disk (SATA3)
- **HEPGPUW105 -**
  - DUAL Intel Xeon E5-2670 V2 2.5 GHz, 10 core processors
  - 64 GB Memory
  - Xeon Phi co-processor (57 cores, 1.1 GHz, 512 bit wide vector unit, 6 GB memory).

Other Recent Developments & Future Work

- Diskless CMS T2 work with Oxford
- SL5 Retirement
- Upgrading our network infrastructure
  - Replacing Nortel 55XX stack with individual Dell N1548 switches
- Upgrading to condor 8.6.3
- Upgrading dCache to 2.16
- STFC now have Cyber Essentials accreditation

List of things to try

- CentOS7 + HTCondor + Docker + Singularity WNs
- FreeIPA
- Expand GPU test boxes into a production service
  - CentOS7 + Docker to allow for different users
- PerfSONAR Refresh – One box to rule them all
- Containerized reconfigurable cluster
- Puppet4
- NoMachine
- LHCONE

Current Issue!

Ghost/Zombie jobs!!

heplnc224	7 of 8	16000 of 16000	Accepting Jobs	Ok	...
heplnc341	20 of 20	49000 of 49000	Accepting Jobs	Ok	...
heplnc342	20 of 20	48000 of 49000	Accepting Jobs	Ok	...
heplnc343	19 of 20	49000 of 49000	Accepting Jobs	Ok	...
heplnc344	19 of 20	49000 of 49000	Accepting Jobs	Ok	...
heplnc345	19 of 20	49000 of 49000	Accepting Jobs	Ok	...
heplnc346	19 of 20				
heplnc347	20 of 20				
heplnc348	20 of 20				
heplnc349	20 of 20				
heplnc350	20 of 20				
heplnc351	20 of 20				
heplnc352	20 of 20				
heplnc353	19 of 20				
heplnc354	20 of 20				
heplnc355	19 of 20				
heplnc356	20 of 20				
heplnc357	20 of 20				
heplnc358	20 of 20				
heplnc359	20 of 20				
heplnc360	20 of 20				
heplnc361	20 of 20				
heplnc362	20 of 20				
heplnc363	20 of 20				
heplnc364	19 of 20				
heplnc365	20 of 20				
heplnc366	19 of 20				
heplnc367	20 of 20	49000 of 49000	Accepting Jobs	Ok	...
heplnc368	20 of 20	47000 of 49000	Accepting Jobs	Ok	...
heplnc369	19 of 20	49000 of 49000	Accepting Jobs	Ok	...
heplnc370	19 of 20	49000 of 49000	Accepting Jobs	Ok	...
heplnc371	20 of 20	47000 of 49000	Accepting Jobs	Ok	...
heplnc372	20 of 20	49000 of 49000	Accepting Jobs	Ok	...
heplnc373	20 of 20	49000 of 49000	Accepting Jobs	Ok	...
heplnc374	20 of 20	45000 of 45000	Accepting Jobs	Ok	...
heplnc375	20 of 20	49000 of 49000	Accepting Jobs	Ok	...
heplnc376	20 of 20	49000 of 49000	Accepting Jobs	Ok	...

```
[root@heplnv146 condor]# check_condor_status

JOB SLOTS WITH NO ACTIVITY FOR 00d:00h:21m:52s
==============================================

slot1_15@heplnc040 LINUX      X86_64 Claimed   Busy      6.300 3000  8+17:44:16
slot1_17@heplnc100 LINUX      X86_64 Claimed   Idle      0.000 2000  0+00:00:09
slot1_14@heplnc342 LINUX      X86_64 Claimed   Busy      1.000 3000  1+16:28:30
slot1_14@heplnc368 LINUX      X86_64 Claimed   Busy      1.100 2000  1+00:38:32
slot1_15@heplnc368 LINUX      X86_64 Claimed   Busy      1.420 3000  1+09:26:38
slot1_16@heplnc368 LINUX      X86_64 Claimed   Busy      1.030 2000  0+05:54:15
slot1_17@heplnc368 LINUX      X86_64 Claimed   Busy      1.020 2000  1+19:39:09
slot1_18@heplnc368 LINUX      X86_64 Claimed   Busy      1.200 2000  0+23:51:45
slot1_19@heplnc368 LINUX      X86_64 Claimed   Busy      1.060 2000  1+00:52:53
[root@heplnv146 condor]#
```

The End…