



Birmingham Status Report

HEP System Managers Meeting, 29th June, 2011
Mark Slater, Birmingham University

At Birmingham we have a local cluster and storage as well as CPUs on a larger shared cluster

The specs for each are:

- 24 8-core machines (192 job slots) @ 9.61-HEP-SPEC06 (Local) ●
- 48 4-core machines (192 jobs slots) @ 7.93-HEP-SPEC06 (Shared) ●
- 177.35T of DPM storage across 4 pool nodes ●

There are 4 CEs (2 CREAM and 2 LCG) controlling the clusters with one SE controlling the storage

All machines are SL5 and are running gLite 3.2 with the exception of the LCG CEs which are still gLite 3.1 (*should these be switched off yet?*)

The only significant recent change (other than in personnel!) was shifting the MySQL server on to a VM from an old machine. There is now only one machine that's out of warranty (cfengine/DHCP head node)

(I'm afraid this will be a little lacking in detail as I'm still learning and both Lawrie and I have been on holiday the last two weeks!)

Our local cluster is on a well defined subnet (not sure the IP range, but small!)

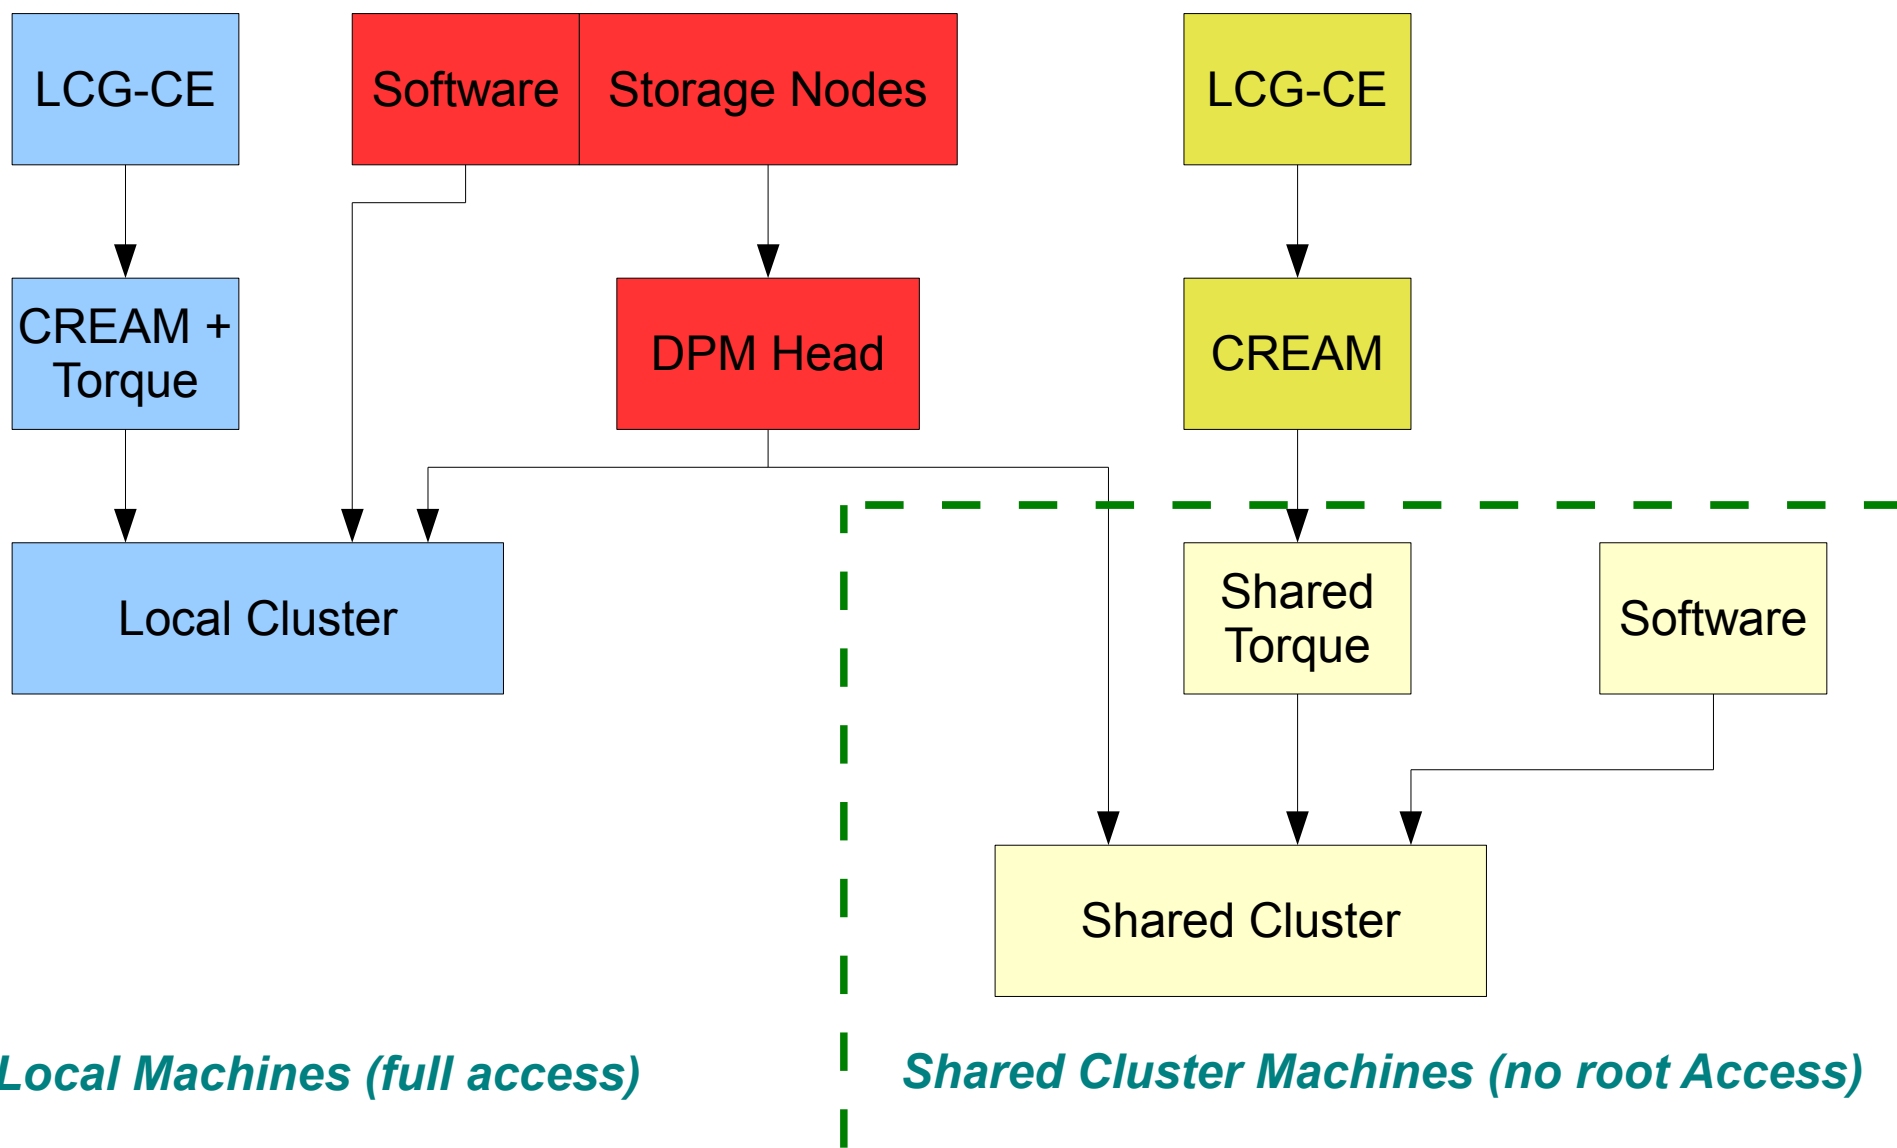
The shared cluster is also on a subnet, however, this subnet also contains other parts of the cluster

I believe we get through to JANET via the main University hub though I'm not sure on speeds

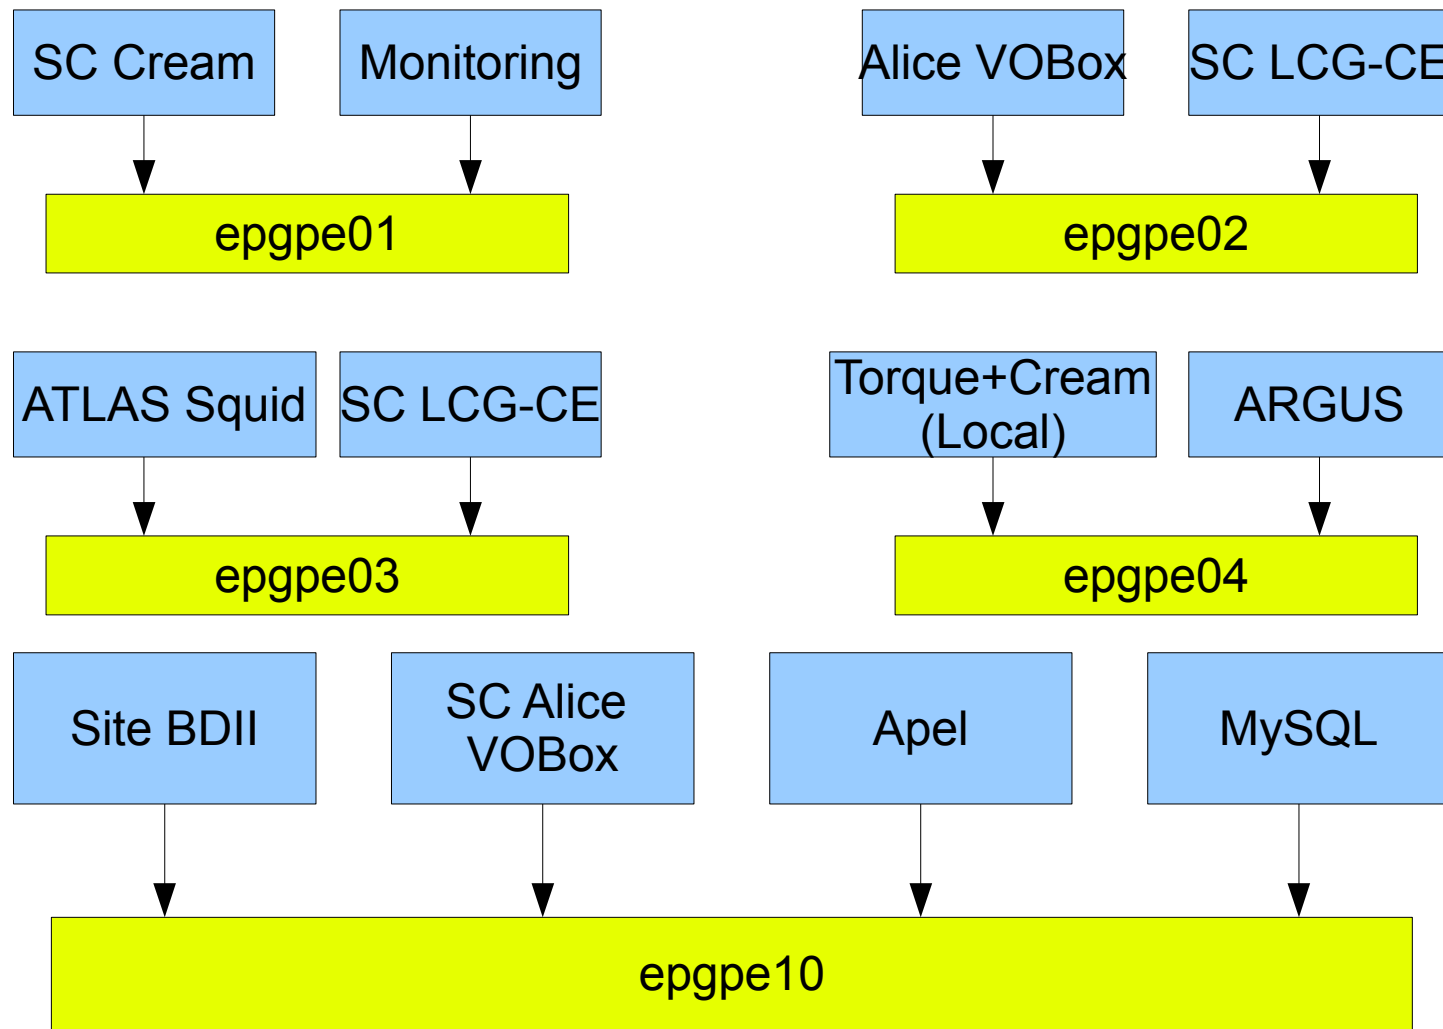
I'll go into monitoring in more detail later but essentially we use Ganglia for the majority of online network monitoring - *I'm hoping to improve this in the near future*



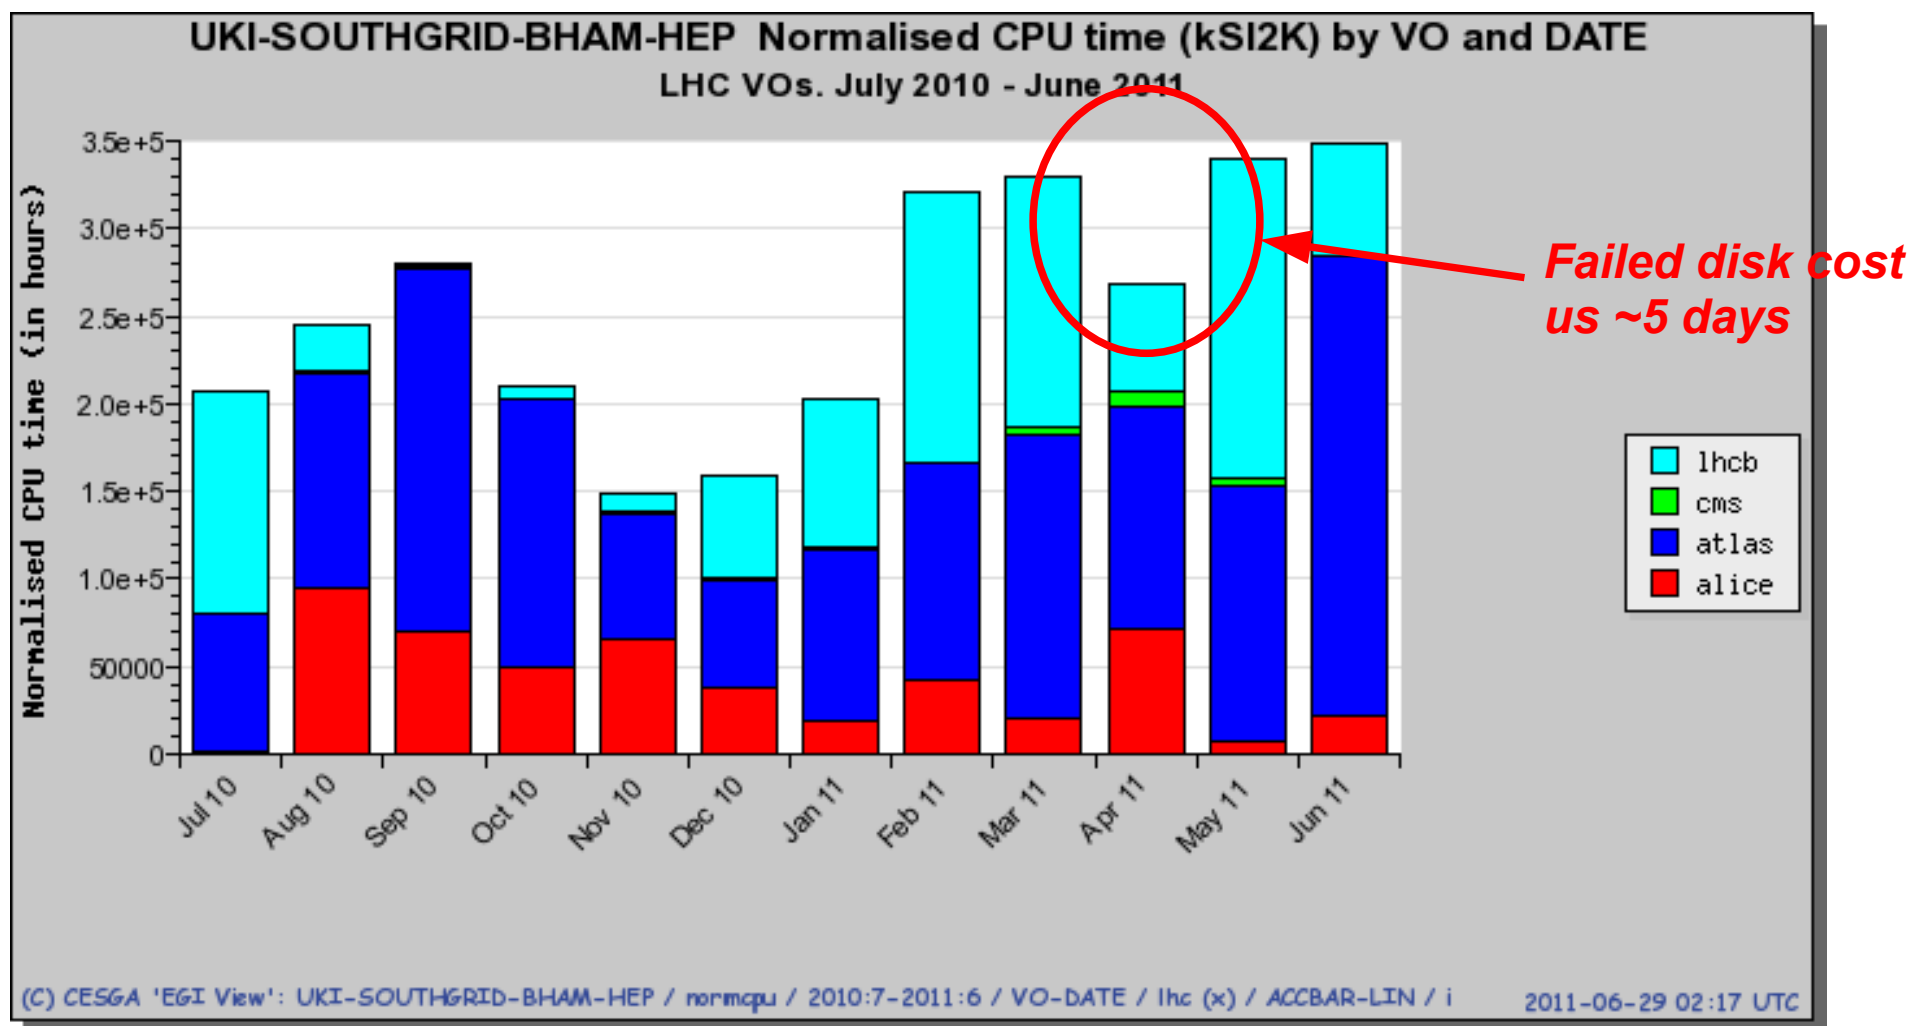
Schematic of the System (1)



We have the majority of our service nodes running as virtual machines (Xen):



We've managed to keep the site relatively stable over the last few months so we are currently enjoying the best figures for Birmingham yet:



We primarily use a combination of Ganglia and PBSWebMon for monitoring and cfengine for controlling the nodes where possible

Cluster status
2011-06-30
07:32:38
Refreshes every 30 seconds.

Show all job details
 Header always on top
 Auto-refresh

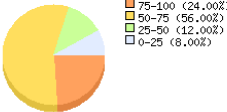
Users								Queues		
User	R	Q	E	C	H	W	Efficiency	Name	Running	Queued
bio084	0	8	0	0	0	1	0	camont	0	0
bio091	0	9	0	0	0	2	0	alice	0	1
ali071	0	1	0	0	0					
bio097	1	1	0	0	0					
fus046	1	0	0	0	0					
sgmops01	1	0	0	0	0					
prdat19	10	0	0	0	0					
prdat05	4	0	0	0	0					
bio052	4	97	0	0	0					
prdic04	40	39	0	0	0					
pilat08	7	0	5	0	0					
pilat09	7	0	10	0	0					
pillhb18	86	0	0	0	0					
Total	161	155	15	0	0					

Hosts down: 0

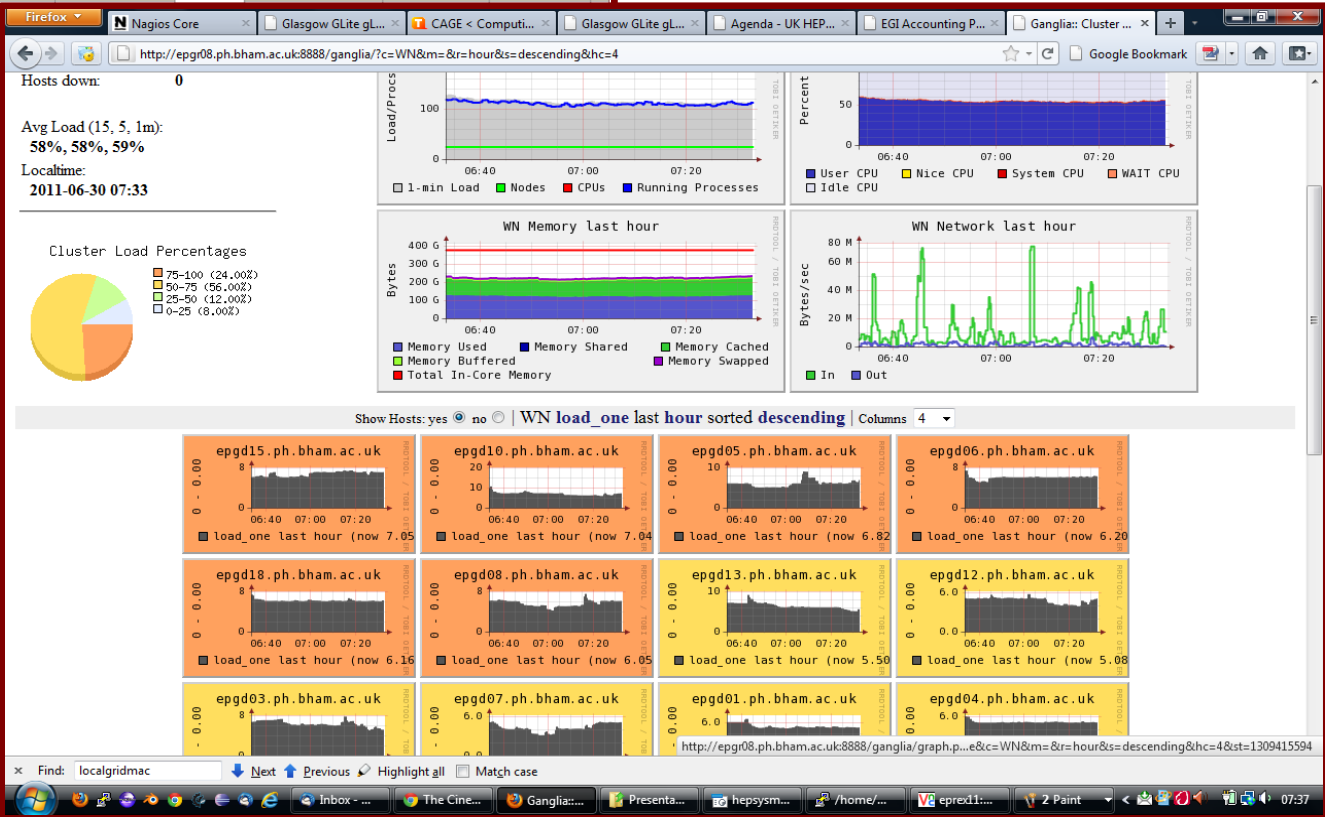
Avg Load (15, 5, 1m):
58%, 58%, 59%

Localtime:
2011-06-30 07:33

Cluster Load Percentages



<p>epgd01.ph.bham.ac.uk <input checked="" type="checkbox"/></p> <p>Show jobs</p> <p>7 jobs, 2 users, 16.44 GB, 5.09 load</p> <table border="1"> <tr><td>0: 236552 pillhb18</td><td>99.13%</td><td>1.01/0.00 GB</td></tr> <tr><td>1: 241473 pilat08</td><td>0.00%</td><td>0.03/0.00 GB</td></tr> <tr><td>2: 238006 prdilc04</td><td>74.43%</td><td>1.42/0.00 GB</td></tr> <tr><td>3: 238830 pillhb18</td><td>99.75%</td><td>1.02/0.00 GB</td></tr> <tr><td>4: 237299 pillhb18</td><td>99.39%</td><td>1.01/0.00 GB</td></tr> <tr><td>5: 238055 pillhb18</td><td>98.70%</td><td>0.97/0.00 GB</td></tr> <tr><td>6: 239627 pillhb18</td><td>99.13%</td><td>1.05/0.00 GB</td></tr> </table>	0: 236552 pillhb18	99.13%	1.01/0.00 GB	1: 241473 pilat08	0.00%	0.03/0.00 GB	2: 238006 prdilc04	74.43%	1.42/0.00 GB	3: 238830 pillhb18	99.75%	1.02/0.00 GB	4: 237299 pillhb18	99.39%	1.01/0.00 GB	5: 238055 pillhb18	98.70%	0.97/0.00 GB	6: 239627 pillhb18	99.13%	1.05/0.00 GB	<p>epgd02.ph.bham.ac.uk <input checked="" type="checkbox"/></p> <p>Show jobs</p> <p>8 jobs, 2 users, 16.44 GB, 4.00 load</p> <table border="1"> <tr><td>0: 240501 pillhb18</td><td>99.02%</td><td>0.88/0.00 GB</td></tr> <tr><td>1: 240739 pillhb18</td><td>91.36%</td><td>0.75/0.00 GB</td></tr> <tr><td>2: 240748 pillhb18</td><td>91.17%</td><td>0.77/0.00 GB</td></tr> <tr><td>3: 240636 pillhb18</td><td>98.85%</td><td>0.84/0.00 GB</td></tr> <tr><td>4: 240831 prdat05</td><td>0.00%</td><td>0.00/0.00 GB</td></tr> <tr><td>5: 240210 pilat08</td><td>0.00%</td><td>0.00/0.00 GB</td></tr> <tr><td>6: 240834 prdat05</td><td>0.00%</td><td>0.00/0.00 GB</td></tr> <tr><td>7: 241477 sgmops01</td><td>0.00%</td><td>0.00/0.00 GB</td></tr> </table>	0: 240501 pillhb18	99.02%	0.88/0.00 GB	1: 240739 pillhb18	91.36%	0.75/0.00 GB	2: 240748 pillhb18	91.17%	0.77/0.00 GB	3: 240636 pillhb18	98.85%	0.84/0.00 GB	4: 240831 prdat05	0.00%	0.00/0.00 GB	5: 240210 pilat08	0.00%	0.00/0.00 GB	6: 240834 prdat05	0.00%	0.00/0.00 GB	7: 241477 sgmops01	0.00%	0.00/0.00 GB
0: 236552 pillhb18	99.13%	1.01/0.00 GB																																												
1: 241473 pilat08	0.00%	0.03/0.00 GB																																												
2: 238006 prdilc04	74.43%	1.42/0.00 GB																																												
3: 238830 pillhb18	99.75%	1.02/0.00 GB																																												
4: 237299 pillhb18	99.39%	1.01/0.00 GB																																												
5: 238055 pillhb18	98.70%	0.97/0.00 GB																																												
6: 239627 pillhb18	99.13%	1.05/0.00 GB																																												
0: 240501 pillhb18	99.02%	0.88/0.00 GB																																												
1: 240739 pillhb18	91.36%	0.75/0.00 GB																																												
2: 240748 pillhb18	91.17%	0.77/0.00 GB																																												
3: 240636 pillhb18	98.85%	0.84/0.00 GB																																												
4: 240831 prdat05	0.00%	0.00/0.00 GB																																												
5: 240210 pilat08	0.00%	0.00/0.00 GB																																												
6: 240834 prdat05	0.00%	0.00/0.00 GB																																												
7: 241477 sgmops01	0.00%	0.00/0.00 GB																																												
<p>epgd05.ph.bham.ac.uk <input checked="" type="checkbox"/></p> <p>Show jobs</p> <p>3 jobs, 4 users, 16.44 GB, 6.24 load</p> <table border="1"> <tr><td>0: 241296 pilat08</td><td>67.89%</td><td>1.12/0.00 GB</td></tr> <tr><td>1: 239945 prdilc04</td><td>88.81%</td><td>1.38/0.00 GB</td></tr> </table>	0: 241296 pilat08	67.89%	1.12/0.00 GB	1: 239945 prdilc04	88.81%	1.38/0.00 GB	<p>epgd06.ph.bham.ac.uk <input checked="" type="checkbox"/></p> <p>Show jobs</p> <p>7 jobs, 2 users, 16.44 GB, 6.16 load</p> <table border="1"> <tr><td>0: 238602 prdilc04</td><td>70.29%</td><td>1.51/0.00 GB</td></tr> <tr><td>1: 215827 prdat19</td><td>0.00%</td><td>0.00/0.00 GB</td></tr> </table>	0: 238602 prdilc04	70.29%	1.51/0.00 GB	1: 215827 prdat19	0.00%	0.00/0.00 GB																																	
0: 241296 pilat08	67.89%	1.12/0.00 GB																																												
1: 239945 prdilc04	88.81%	1.38/0.00 GB																																												
0: 238602 prdilc04	70.29%	1.51/0.00 GB																																												
1: 215827 prdat19	0.00%	0.00/0.00 GB																																												



The screenshot shows the Ganglia monitoring interface in a browser window. It displays a dashboard with several graphs and tables. At the top, there are navigation tabs for various services. The main content area includes:

- A 'Hosts down' indicator showing 0 hosts down.
- An 'Avg Load' section showing 58%, 58%, and 59% for 15, 5, and 1m time periods.
- A 'Cluster Load Percentages' pie chart.
- Four line graphs: 'Load/Procs', 'Percent', 'WN Memory last hour', and 'WN Network last hour'.
- A grid of 16 smaller graphs showing 'load_one last hour' for various nodes (e.g., epgd15, epgd10, epgd05, epgd06, epgd18, epgd08, epgd13, epgd12, epgd03, epgd07, epgd01, epgd04).

Failed Disk in May •

One of the disks on our host machines (pe10 - hosting site BDII, APEL, VOBox and MySQL) started failing

Testing seemed to show the disk was OK so it might have been the controller

However, Dell replaced the disk and everything was restored from a direct backup of the whole disk taken before replacement

SE Overloading •

The SE started 'stalling' and would only stay up for 10mins after restart

This was found to be linked to excessive H1 jobs and direct transfers and also highlighted an issue with the MOAB config on the shared cluster.

Slow ATLAS setup •

We also have an on-going problem of multiple Atlas user jobs on a machine being very slow to setup - not sure if this is NFS related or load related

Overloading CEs? ●

Recently, I've noticed that several jobs can be stuck in waiting in Torque

This usually happens when we have a large number of jobs be submitted at once and checking Ganglia at these times reveals high load on the machine

Would splitting the Torque and Cream CE on to separate machines help?

ATLAS releases ●

In May we started getting frequently ticketed by Atlas saying that we didn't have releases installed when Panda said we did

This was due to install jobs from Atlas only going through one CE and therefore only installing on one software area (the BDII was reporting OK)

Solved by getting Allesandro de Salvo to make sure multiple jobs were sent and liaising with Alden Stradlin to update the settings for Bham in Panda

Over the next few months, we plan to do the following:

Get glxec working ●

Everything has been installed but local tests are failing at present

Improve network monitoring ●

As Ganglia doesn't give much detail, we will try to improve this, e.g. cacti

Deploy CVMFS ●

This should hopefully sort out our release issues and the 'slow setup' issue

Improve job monitoring ●

It would be nice to analyse job stats better (e.g. wall time, cpu time, etc.)