

HEP Computing Matters

Group Meeting, 2nd September, 2013
Mark Slater, Birmingham University

The current list of Tier 2 kit includes:

- 816 job slots over 40 boxes ●
- Intel and AMD cores, with 12 machines at Bluebear ●
- 320TB of storage available, 120TB → Alice, 170TB → ATLAS ●
- Currently we have our own dedicated 10Gbit link to JANET ●

I'm in the process of upgrading all machines/services to use SL6 and EMI3 which I'll be doing in the next 2 days

The current list of Tier 3 kit includes:

- 152 job slots over 16 boxes + 2 login nodes ●
- AMDs on the dedicated boxes, Intels on eprexa/b ●
- 160TB storage available, mostly full (26b free for LHCb) ●
- Limited to 1Gbit off campus but some 10Gbit connectivity within ●

For the past few weeks, I've been rolling out a significant upgrade to the management of both Tier2 and Tier3 systems

The core of this is the Puppet management software and the Foreman web UI that most UK grid sites are now adopting. This handles:

- Provisioning new machines ●
- Complete install and setup from scratch ●
- Management of all required services centrally (firewall, ssh, etc.) ●

This should allow very easy application of changes across all machines, installation of new packages, monitoring, etc.

I'm also trying to improve the monitoring through the use of Graphite which will eventually monitor EVERYTHING on ALL machines, e.g. load, network traffic, services, logins, etc.

I'm now recording everything that I do as well as all purchases, quotes, machine specs, etc. in a Wordpress blog

Finally, I plan to combine all of this info with a few scripts to provide things such as:

- A single 'state of the union' page that will show if anything's broken' ●

- Monitoring of network traffic (bottlenecks, etc.) ●

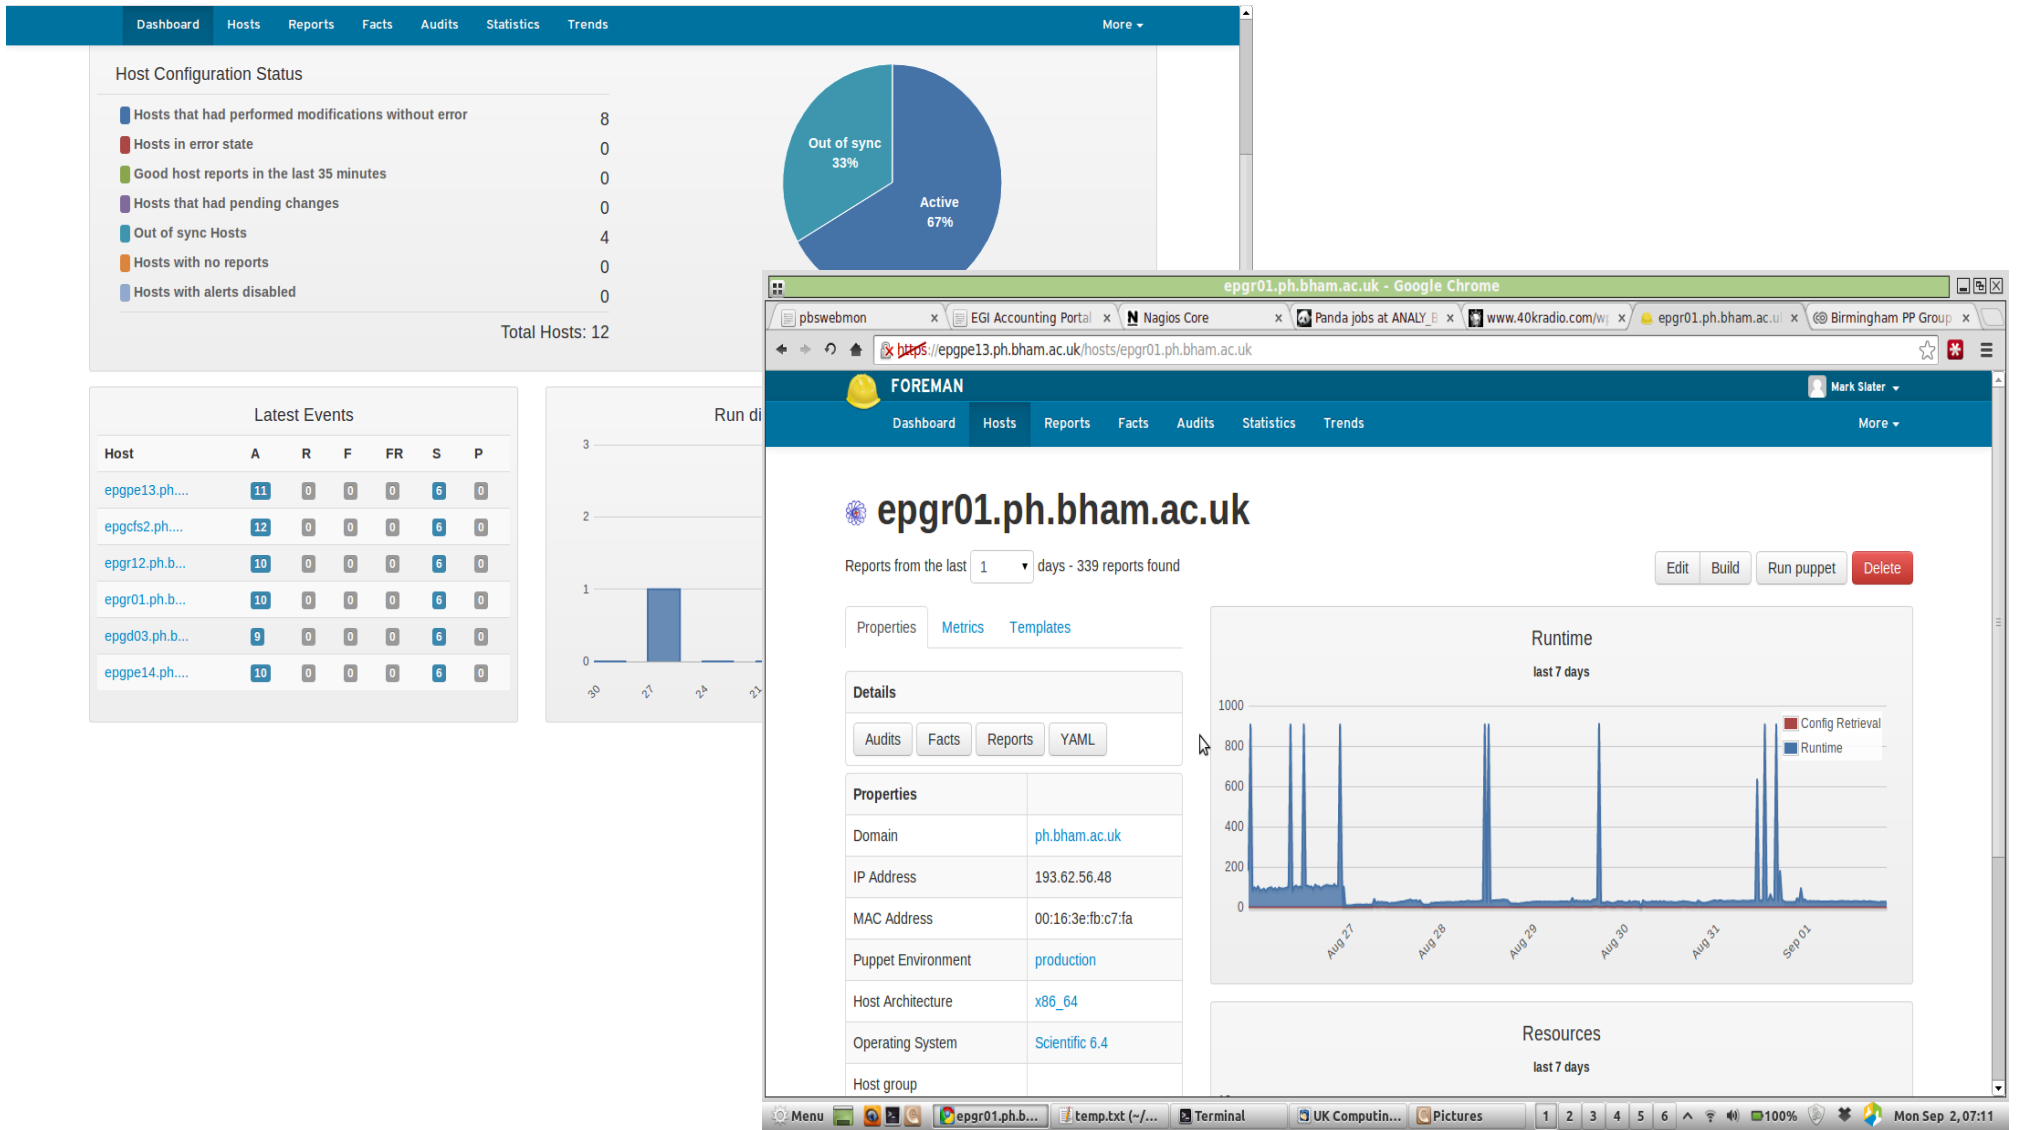
- Load of the system, Rack temperature, power consumption, etc. ●

- Easy to find machine info, warranty info, purchase orders, etc. ●

Some of this will be available to everyone so, e.g. you can easily check if you think your desktop is not behaving or the your jobs are running slow

Some users (Nigel, Pete F?) can have full access if required so more significant work/fault analysis can be done when I'm away

This can be extended to any internal/external tests people might like...



The screenshot displays the Foreman web interface for the host `epgr01.ph.bham.ac.uk`. The interface is divided into several sections:

- Host Configuration Status:** A summary of host states. A pie chart shows 67% Active and 33% Out of sync. A table lists various status categories and their counts.
- Latest Events:** A table showing recent events for several hosts, including `epgpe13.ph...`, `epgcfs2.ph...`, `epgr12.ph.b...`, `epgr01.ph.b...`, `epgd03.ph.b...`, and `epgpe14.ph...`. The table has columns for Host, A, R, F, FR, S, and P.
- Runtime Metrics:** A bar chart showing runtime metrics for the last 7 days, with a peak around 1000. The chart includes a legend for Config Retrieval and Runtime.
- Host Details:** A table listing properties for the selected host, such as Domain, IP Address, MAC Address, Puppet Environment, Host Architecture, Operating System, and Host group.

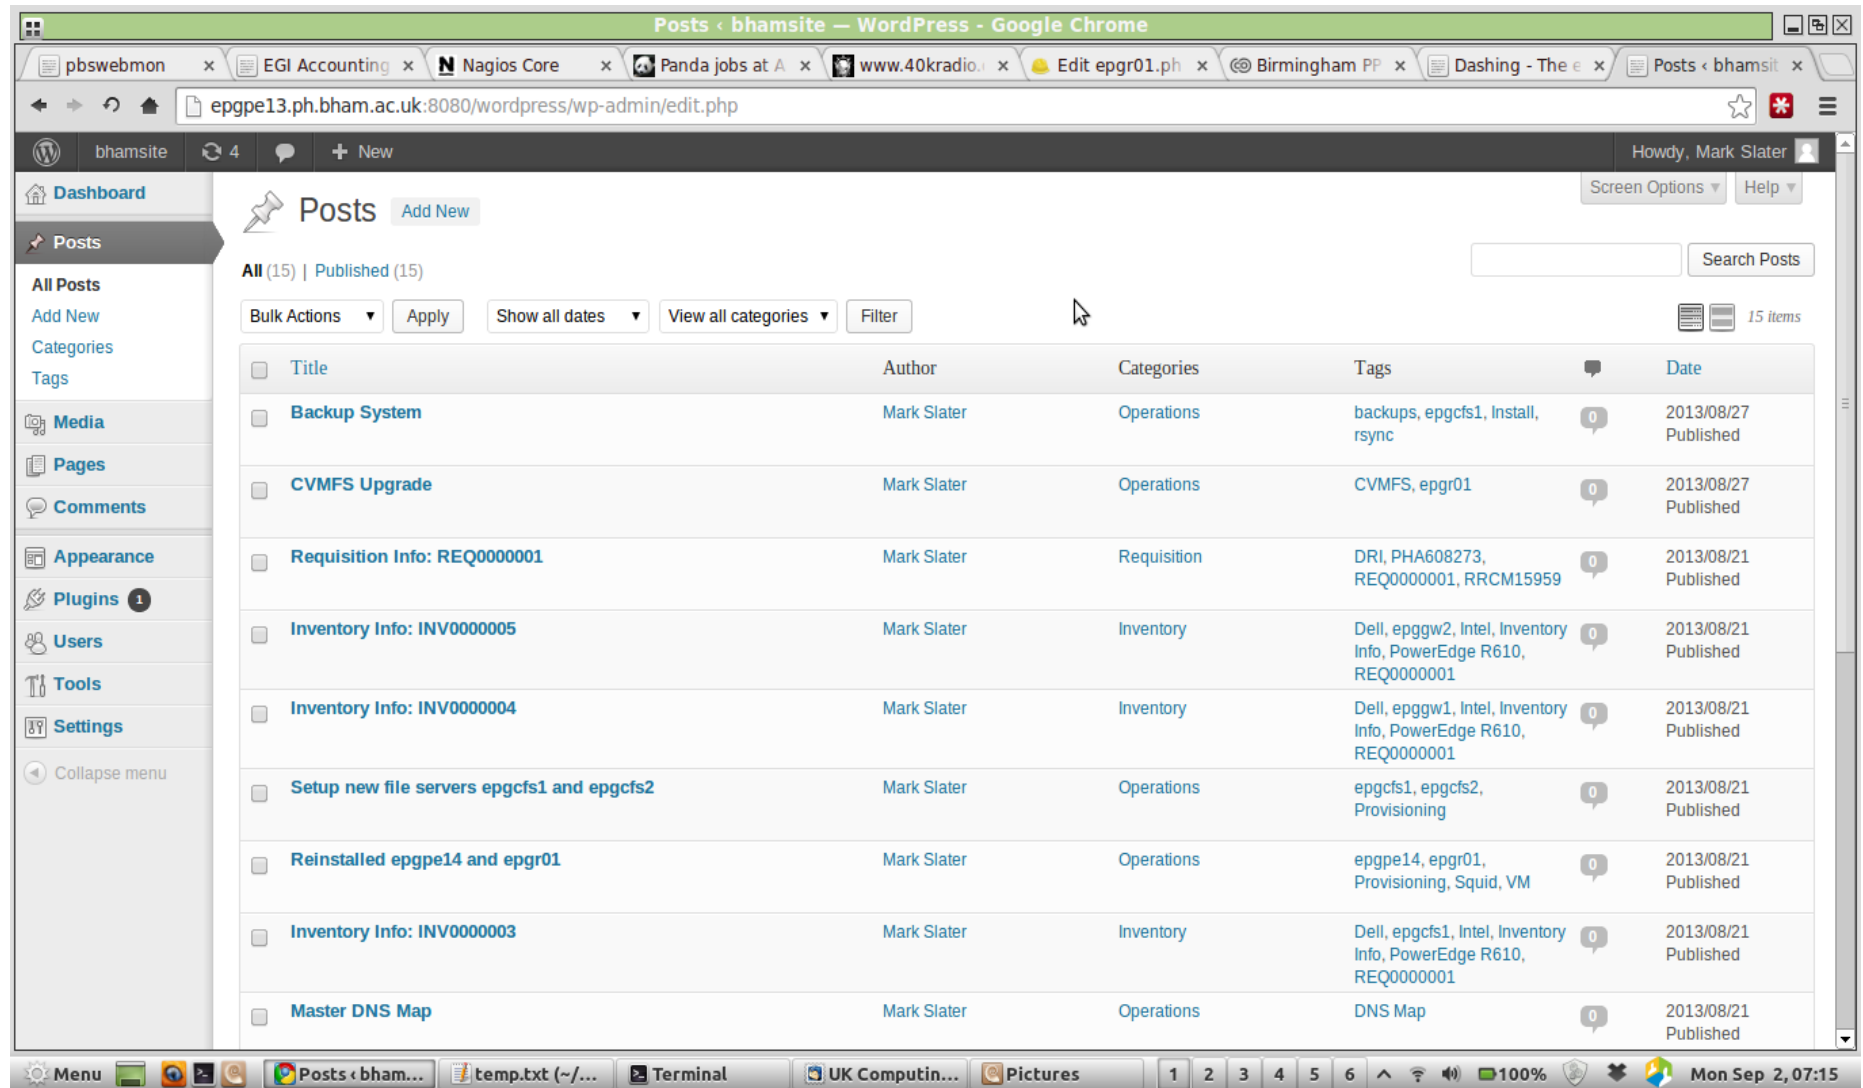
Category	Count
Hosts that had performed modifications without error	8
Hosts in error state	0
Good host reports in the last 35 minutes	0
Hosts that had pending changes	0
Out of sync Hosts	4
Hosts with no reports	0
Hosts with alerts disabled	0

Total Hosts: 12

Host	A	R	F	FR	S	P
epgpe13.ph...	11	0	0	0	6	0
epgcfs2.ph...	12	0	0	0	6	0
epgr12.ph.b...	10	0	0	0	6	0
epgr01.ph.b...	10	0	0	0	6	0
epgd03.ph.b...	9	0	0	0	6	0
epgpe14.ph...	10	0	0	0	6	0

Property	Value
Domain	ph.bham.ac.uk
IP Address	193.62.56.48
MAC Address	00:16:3e:fb:c7:fa
Puppet Environment	production
Host Architecture	x86_64
Operating System	Scientific 6.4
Host group	





The screenshot shows the WordPress admin dashboard for the 'bhamsite' website. The user is logged in as 'Mark Slater'. The main content area displays a list of 15 posts, all published on August 21 or 27, 2013. The posts are organized into categories such as Operations, Requisition, and Inventory. The interface includes a sidebar with navigation options like Dashboard, Posts, Media, Pages, Comments, Appearance, Plugins, Users, Tools, and Settings. The top of the dashboard shows the user's name and a search bar for posts.

<input type="checkbox"/>	Title	Author	Categories	Tags	<input type="checkbox"/>	Date
<input type="checkbox"/>	Backup System	Mark Slater	Operations	backups, epgcfs1, Install, rsync	0	2013/08/27 Published
<input type="checkbox"/>	CVMFS Upgrade	Mark Slater	Operations	CVMFS, epgr01	0	2013/08/27 Published
<input type="checkbox"/>	Requisition Info: REQ0000001	Mark Slater	Requisition	DRI, PHA608273, REQ0000001, RRCM15959	0	2013/08/21 Published
<input type="checkbox"/>	Inventory Info: INV0000005	Mark Slater	Inventory	Dell, epggw2, Intel, Inventory Info, PowerEdge R610, REQ0000001	0	2013/08/21 Published
<input type="checkbox"/>	Inventory Info: INV0000004	Mark Slater	Inventory	Dell, epggw1, Intel, Inventory Info, PowerEdge R610, REQ0000001	0	2013/08/21 Published
<input type="checkbox"/>	Setup new file servers epgcfs1 and epgcfs2	Mark Slater	Operations	epgcfs1, epgcfs2, Provisioning	0	2013/08/21 Published
<input type="checkbox"/>	Reinstalled epgge14 and epgr01	Mark Slater	Operations	epgge14, epgr01, Provisioning, Squid, VM	0	2013/08/21 Published
<input type="checkbox"/>	Inventory Info: INV0000003	Mark Slater	Inventory	Dell, epgcfs1, Intel, Inventory Info, PowerEdge R610, REQ0000001	0	2013/08/21 Published
<input type="checkbox"/>	Master DNS Map	Mark Slater	Operations	DNS Map	0	2013/08/21 Published

Commission new Storage

Thanks to most VOs using CVMFS now (distributed, cached read only filesystem), I plan to recover 5TB from the RAIDs and use a smaller server for this. After buying a lot of new 2TB disks I am going to populate the server with these giving us spares to keep the old kit running

Commission BlueBear provided storage

Have now got 30TB available at BB that I should be commissioning soon. Combined with above, this should give an additional 45TB of space which should help our cause in GridPP

New Servers

Before Lawrie left, he had bought several new servers for Tier 2 use. During the upgrade this week I plan to shift all services to these and either convert the old machines to workers or bin them

Improved Monitoring

After pushing out the new management system, monitoring should become a lot easier. I will hopefully be able to get the system to keep track of things and email people with they go 'out of bounds'!

Networking and other bottlenecks

This will also allow me to check for network and other bottlenecks and deal with them a lot easier and in a more permanent way. I should be able to upgrade the Tier3 subnet to 10Gbit in the near future as well

New Workers

I'll be quite a lot more computing power in the near future and retiring the ageing ones we currently have. Hopefully, I'll be able to provide more job slots on better processors.

New Storage

I also plan to buy more storage in the near future to replace older, out of warranty kit. However – be aware that long term, it will become harder to download all the data you want!

CVMFS for Atlas and LHCb

I will be adding a CVMFS caching machine in the next couple of months that should allow ALL releases to be available across all machines without any install/admin effort

New Servers

We have some rather important services running on some rather old kit, e.g. Mail Server, DHCP, etc. I'll be upgrading these at some point in the next year and try to minimise disruption when it happens.

Finally, here are a few things I'm aware of and will try to be fixing soon – please let me know any more issues or 'Quality of Life' things that you would like!

- Slow connection to disk servers sometimes
- Poor connection to the outside world sometimes
- Lack of storage space
- Problems with one of the Duets in the meeting room
- Lack of group laptops
- Scanner not great

Anything else?....