



GridPP
UK Computing for Particle Physics

Tier1 Deployment

Steve Traylen.

Your university or
experiment logo here



GridPP
UK Computing for Particle Physics


Contents

- Job Efficiencies
- CPU Fairshare
- Installation Method at the Tier1.
- Recent/Upcoming Services.

Your university or
experiment logo here



CPU Use (KSI2K*CPUMonths)



Your university or
experiment logo here




- New plots are now generated regularly.
 - <http://www.gridpp.rl.ac.uk/stats/>
- A report has been written for the Tier1 Board that describes how these are being calculated.
 - http://www.gridpp.ac.uk/tier1a/board/doc/Utilisation_OC_August_2005_V1.0.doc



GridPP
UK Computing for Particle Physics

Overall, (total cpu/total wall)

Efficiencies (2005)



Your university or
experiment logo here




GridPP

UK Computing for Particle Physics

Individual VOs


Efficiencies (2005)



Your university or
experiment logo here



ATLAS 05/2005 (23527 jobs, 468.9 CPU months)





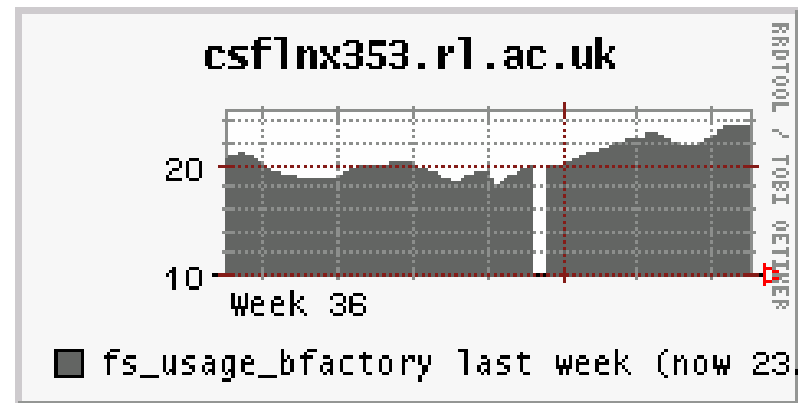

- Currently it is completely based on UNIX groupid.
- Any spare capacity is allocated to over subscribers in proportion to their allocation.
- This may change.....



- We can create accounts in Maui.
 - e.g. an LHC account and a non-LHC account.
 - e.g. a grid account and a non Grid account.
- These accounts can be given allocations.
 - But we now have competing scheduling regimes.
 - Each of the regimes can be weighted but the initial values and subsequent tuning is not obvious.



- New plots are now being constructed.
- They will plot:
 - allocated % for different VOs.
 - utilised % for different VOs.
 - allocated/utilised for different VOs.



Your university or
experiment logo here



How is tier1 deployed.

- Still done with.
 - DHCP/PXE/Kickstart and bunch of scripts.
 - YAIM is just one bit.
 - Local config packages are built.
 - e.g. tier1-ntp-config, tier1-sendmail-config, tier1-lmsensors-config, tier1-yaim-config,....
 - Packages maintained with yum/yumit.
- Batch workers:
 - Maintained via parallel SSH.
 - Down nodes are reinstalled routinely.
- Complicated service nodes:
 - Maintained by hand, often the only way.
 - Backed up to tape.



- It is simple especially for new starters.
 - Only bash and rpm building is required.
 - Flexible, extending what is there is easy.
- Problems
 - Things can get out step, though in reality not a huge problem.
 - Configuration in packages helps a lot.
 - We changed our netmask everywhere via a package upgrade.
 - Farm wide configuration changes are rare.



- <http://sf.net/pakiti/>
- Yumit renamed as Pakiti
 - Now support yum/up2date/apt-get.
- Improvements are being added now.
 - Currently only packages that could be updated published.
 - It will also display installed packages.
 - This will flag a WN with a missing/extra RPM.
 - Not a security feature but very useful.



- FTS v1.3 is now being deployed.
 - Install complete by Matt Hodges
 - Understanding is a lot more thorough than with FTS v1.1.2 when it was thrown up.
- VOBox primarily for Alice.
 - Install complete by Catalin Condurache.
 - Testing/Validation with EIS group underway.
 - Bugs found - new version released today.



- The two most critical boxes at RAL.
 - PBS scheduler.
 - Recent crashes have stopped submissions for a day or two at least twice in last three months.
 - DCache head node.
- Work is now starting on use of the HA-Linux software package to increase redundancy.
 - <http://www.linux-ha.org/>
 - Other (big) sites in the UK will probably want this.