



**GridPP**

UK Computing for Particle Physics

# The Tier-2 Status and Final Preparation Steps: NorthGrid


James Cullen

NorthGrid Deputy Technical Coordinator  
University of Manchester

GridPP23

Clare College, Cambridge University

Wednesday 9<sup>th</sup> September 2009





- NorthGrid Status
  - APEL Accounting
- Site News
  - Lancaster
  - Liverpool
  - Manchester
  - Sheffield
- Atlas, LHCb and SL5
- Conclusions



# NorthGrid Status

- UK SAM test success rate continues to be high.


Site	2Q07	3Q07	4Q07	1Q08	2Q08	3Q08	4Q08	1Q09	2Q09	3Q09
UKI-NORTHGRID-LANCS-HEP	52%	92%	95%	89%	94%	96%	83%	96%	100%	99%
UKI-NORTHGRID-LIV-HEP	95%	91%	98%	70%	97%	93%	98%	98%	99%	98%
UKI-NORTHGRID-MAN-HEP	99%	99%	87%	98%	87%	99%	92%	93%	96%	94%
UKI-NORTHGRID-SHEF-HEP	50%	37%	82%	98%	96%	96%	97%	97%	98%	94%

LondonGrid	79%	82%	80%	75%	72%	79%	84%	86%	86%	74%
NorthGrid	74%	80%	91%	89%	94%	96%	93%	96%	98%	96%
ScotGrid	91%	92%	73%	82%	85%	92%	94%	97%	97%	96%
SouthGrid	80%	86%	91%	94%	92%	93%	94%	95%	96%	95%

<http://pprc.qmul.ac.uk/~lloyd/gridpp/samtest.html>


	GridPP		MoU calculations	
	CPU (kSI2K)	Storage (TB)	% of MoU CPU	% of MoU Disk
Lancaster	1040	143	284.62%	78.14%
Liverpool	829	150	236.18%	226.24%
Manchester	2160	244	158.59%	56.52%
Sheffield	182.5	38	921.72%	383.84%
<b>Totals</b>	<b>4211.5</b>	<b>575</b>	<b>200.72%</b>	<b>83.22%</b>

- 09Q1 & Q2 reports - Lancaster & Manchester delivering less than agreed in the MoU.







- 12 month period (Sept '08 - Aug '09) NorthGrid provided ~15% of UKI normalised CPU Time for LHC VOs



Percentage of UKI Normalised CPU Time by Region Sept '08 - Aug '09

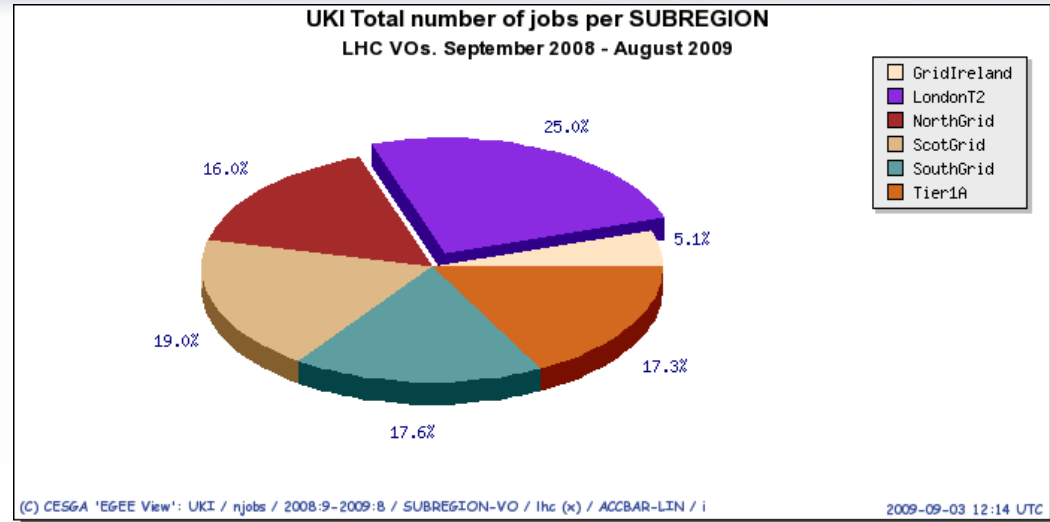


- Monthly percentage fairly steady (except first and last points!)







- Same period NorthGrid provided 16% of total UKI LHC VO jobs



Percentage Number of Jobs by Region Sept '08 - Aug '09



- Steady amount of jobs over the 12 month period





**GridPP**


UK Computing for Particle Physics

# Lancaster

Peter Love & Matt Doidge

## Current status:

- Have two CEs (~600 cores) and 186TB storage
- Old CE offline to be refurbished and upgraded, retiring troubled WNs and consolidating the remainder
- LAN partially upgraded to 10Gb/s
- 2009 procurement kit being delivered by early September
- Purchasing new hardware for DPM Headnode, going for a split SRM/Database setup using Raid 10 SAS disks on the DB node.
- Plan to move to Virtualised Servers for some services.
- Will upgrade to DPM-1-7.X soon.
- All new kit (workers + dpm pools) will be running on SL5. We'll upgrade older kit to SL5 in a rolling fashion after that (provided the users are still all up for the move).



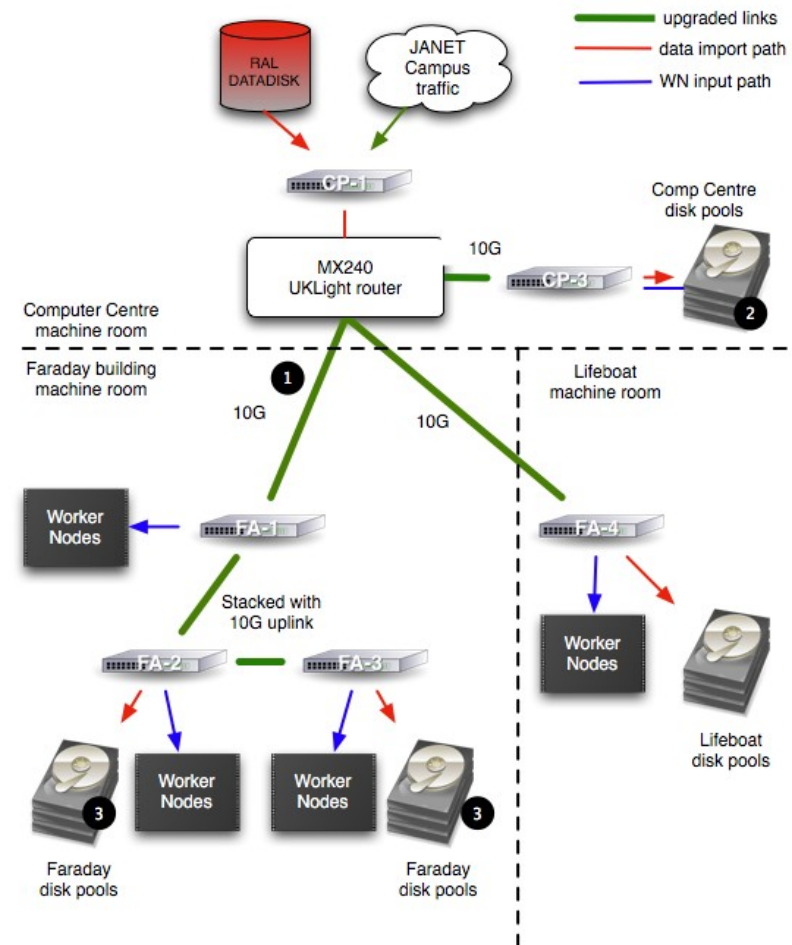
NorthGrid



## Plans for September, October:

- Complete LAN upgrade (green lines in diagram)
- commission 2009 kit (additional 256core and 280TB)
- re-assess ATLAS throughput and tweak with new kit and new LAN
- fit-out machine room in new data centre

## Lancaster's upgraded LAN





**GridPP**


UK Computing for Particle Physics

# Liverpool

John Bland, Rob Fay & Steve Jones

## Current status

- A transducer defect in our water cooling system has been repaired, but underlying issues remain. This limits us to approximately 80% of potential capacity. Apart from the cooling issue, all nodes are operational.
- The request to upgrade a significant portion of our infrastructure to SL5 before October cannot be met because 88% of our worker node resources are 32 bit, and unsuitable for upgrade.
- We have upgraded our kernels in response to the local root vulnerability (CVE-2009-2692). CE and WN middleware is up-to-date.



NorthGrid



## Plans


- We are currently creating an automatic build system for SL5, 64 bit, 8 core nodes, using kick-start, puppet and yaim etc. We have tentative plans to upgrade our BDII. We are continuing to harden and optimise our systems, using HammerCloud data.

## Weaknesses

- The vast majority of our hardware is over 5 years old. It is approaching the end of its useful life and it consumes much power relative to its SpecInt.

## Readiness


- Within the constraints outlined above, we could take data. Obviously, we'd be affected by unexpected changes to the operational baseline.






## Current Status

- dCache, which was on the WNs, has been decommissioned (June 2009):
  - Lost ~150 TB of storage space
  - WNs have gained a second system disk. SL5 WNs will use RAID 0.
- Stability problems with CEs since July. Upgraded gLite software but problems remain.



- DPM 1.6 installed on 4 SEs. All have bonded 2 x 1 Gb/s Ethernet interfaces connected directly to the main Cisco switch. One SE also has an extra dual 1 Gb/s card giving 4 Gb/s (but hasn't been tested in anger yet).





- Near Future
- Worker Nodes
  - Plan to migrate half the cluster to SL5 in mid-September, to coincide with data centre work. Second cluster in October.
  - Ageing machines:
    - single core processors with lower SpecInt2k6 values compared to modern processors (cf. <http://www.gridpp.ac.uk/wiki/HEPSPEC06> )
    - relatively high power consumption
  - Purchase new WNs. Replace half the WNs (~900 processors) with ~100 dual quad core machines.
  - Ethernet connections will be bonded to give 2 Gb/s.
- Storage and Network
  - Upgrade to DPM 1.7.x
  - Purchase 11 more 24 TB storage units (~230 TB).
  - Purchase 3 x 10 Gb switches for new storage, and 10 Gb blade for main Cisco switch.
- Install WMS




**GridPP**

UK Computing for Particle Physics

# Sheffield

Elena Korolkova & Matt Robinson

- 200 single Core 2.4 GHz Opterons (2GB), 64 bit SL4, 72 GB local disk per 2 cores
- In August several broken WNs and broken RAM were replaced with working ones
- Last week of August upgraded to the latest GLite 3.1 & secure kernel-2.6.9-89.0.9.EL.x86\_64
- CE and MONBOX were upgraded to the latest GLite 3.1 & secure kernel-2.6.9-89.0.9.EL.x86\_64
- Storage
  - SL4, RAID5, DPM 1.7.06
  - 3 X 13 TB pools
- - cluster has 97-98% availability and realability since January 2008



NorthGrid



- As a small site mainly support ATLAS VO:
  - 35 TB are reserved in ATLAS space tokens
  - 96% efficiency in ATLAS production since January 2009
  - 98% efficiency in ATLAS data transfers since 2009
  - successful participation in STEP09 in June 2009
  - successful participation in Hammer cloud tests, site throughput was measured


Site is also used by LHCb

- Future Plans
  - upgrade WNs to SL5 (the beginning of October 2009)
  - install a separate node with a mirrored 1TB disk for EXP SOFTWARE (October 2009)
  - install 2 TB disks in one pool node instead of 1 TB disks (increased storage by 11 TB) (September 2009)
- Weak Points
  - No access to the cluster out of working hours
  - Cluster availability dropped to 94% in Q309 because couldn't restart monbox remotely during Bank Holiday




- Manchester supplying much LHCb CPU usage

CPU usage by Site  
31 Days from 2009-08-05 to 2009-09-05



- |                         |                  |                       |                 |
|-------------------------|------------------|-----------------------|-----------------|
| ■ LCG.GRIDKA.de         | ■ LCG.RAL.uk     | ■ LCG.Manchester.uk   | ■ LCG.CNAF.i    |
| ■ LCG.RAL-HEPuk         | ■ LCG.PIC.es     | ■ LCG.GLASGOW.uk      | ■ LCG.Krakow.pl |
| ■ LCG.CSCS.ch           | ■ LCG.NIKHEF.nl  | ■ LCG.UKI-LT2-QMUL.uk | ■ LCG.CERN.ch   |
| ■ LCG.TCD.ie            | ■ LCG.DURHAM.uk  | ■ LCG.IN2P3-T2.fr     | ■ LCG.IN2P3.fr  |
| ■ LCG.JINR.ru           | ■ LCG.DESYde     | ■ LCG.KIAE.ru         | ■ LCG.LAPPfr    |
| ■ LCG.CNAF-T2.it        | ■ LCG.SARA.nl    | ■ LCG.DSMC.nl         | ■ LCG.CPPM.fr   |
| ■ LCG.UKI-LT2-IC-HEPuk  | ■ LCG.Legnaro.it | ■ LCG.USC.es          | ■ LCG.USC.es    |
| ■ LCG.CNAF-T2.it        | ■ LCG.Oxford.uk  | ■ LCG.UNINA           | ■ LCG.UNINA     |
| ■ LCG.UKI-LT2-RHUL.uk   | ■ LCG.EFDA.uk    | ■ LCG.LPC.fr          | ■ LCG.LPC.fr    |
| ■ LCG.UKI-LT2-Brunel.uk | ■ LCG.PDC.se     | ■ LCG.Lorino.it       | ■ LCG.LPC.fr    |
| ■ LCG.Cambridge.uk      |                  | ■ LCG.Dortmund.de     | ■ LCG.LPC.fr    |
|                         |                  | ■ LCG.PNPI.ru         | ■ LCG.LPC.fr    |
- Maximum: 18,231 , Minimum: 0.00 , Average: 5,110 , Current: 2,010  
... plus 74 r


CPU used by Site  
31 Days from 2009-08-05 to 2009-09-05




- |                       |                     |                        |                    |
|-----------------------|---------------------|------------------------|--------------------|
| ■ LCG.GRIDKA.de       | ■ LCG.Manchester.uk | ■ LCG.RAL.uk           | ■ LCG.PIC.es       |
| ■ LCG.CNAF.it         | ■ LCG.RAL-HEPuk     | ■ LCG.Krakow.pl        | ■ LCG.CERN.ch      |
| ■ LCG.IN2P3-T2.fr     | ■ LCG.NIKHEF.nl     | ■ LCG.CSCS.ch          | ■ LCG.IN2P3.fr     |
| ■ LCG.UKI-LT2-QMUL.uk | ■ LCG.JINR.ru       | ■ LCG.TCD.ie           | ■ LCG.LAPPfr       |
| ■ LCG.DESYde          | ■ LCG.USC.es        | ■ LCG.DURHAM.uk        | ■ LCG.LAPPfr       |
| ■ LCG.LPC.fr          | ■ LCG.Legnaro.it    | ■ LCG.Oxford.uk        | ■ LCG.CPPM.fr      |
| ■ LCG.PSNC.pl         | ■ LCG.EFDA.uk       | ■ LCG.UKI-LT2-IC-HEPuk | ■ LCG.EFDA.uk      |
| ■ LCG.KIAE.ru         | ■ LCG.CNAF-T2.it    | ■ LCG.CNAF-T2.it       | ■ LCG.Legnaro.it   |
| ■ LCG.SINPr           | ■ LCG.UNINA.it      | ■ LCG.UNINA.it         | ■ LCG.CBPF.br      |
| ■ LCG.SINPr           | ■ LCG.PNPI.ru       | ■ LCG.PNPI.ru          | ■ LCG.Barcelona.es |
| ■ LCG.RUG.nl          |                     |                        | ■ LCG.Barcelona.es |
- Total: 163,696 , Average Rate: 0.06 /s  
... plus 74 more



## CPU days used by Site (Sum: 163728)




- |                              |                                |                               |                           |
|------------------------------|--------------------------------|-------------------------------|---------------------------|
| ■ LCG.GRIDKA.de (10237)      | ■ LCG.Manchester.uk (10099)    | ■ LCG.RALuk (9783)            | ■ LCG.PIC.es (8509)       |
| ■ LCG.CNAF.it (7725)         | ■ LCG.GLASGOW.uk (6400)        | ■ LCG.Krakow.pl (6101)        | ■ LCG.CERN.ch (5398)      |
| ■ LCG.IN2P3-T2.fr (4733)     | ■ LCG.RAL-HEPuk (4691)         | ■ LCG.CSCS.ch (4648)          | ■ LCG.IN2P3.fr (4576)     |
| ■ LCG.NIKHEF.nl (4393)       | ■ LCG.JINR.ru (4035)           | ■ LCG.TCD.ie (3879)           | ■ LCG.LAPPfr (3840)       |
| ■ LCG.UKI-LT2-QMUL.uk (3663) | ■ LCG.USC.es (3562)            | ■ LCG.DURHAM.uk (3409)        | ■ LCG.CPPM.fr (3371)      |
| ■ LCG.DESYde (3134)          | ■ LCG.Liverpool.uk (3097)      | ■ LCG.Oxford.uk (2339)        | ■ LCG.EFDA.uk (2293)      |
| ■ LCG.LPC.fr (2233)          | ■ LCG.Dortmund.de (2039)       | ■ LCG.UKI-LT2-IC-HEPuk (1758) | ■ LCG.Legnaro.it (1717)   |
| ■ LCG.PSNC.pl (1688)         | ■ LCG.UKI-LT2-Brunel.uk (1646) | ■ LCG.CNAF-T2.it (1560)       | ■ LCG.CBPF.br (1560)      |
| ■ LCG.KIAE.ru (1535)         | ■ LCG.SINPrU (1511)            | ■ LCG.UNINA.it (1275)         | ■ LCG.Barcelona.es (1151) |
| ■ LCG.Sheffield.uk (1138)    | ■ LCG.RUG.nl (1077)            | ■ LCG.PNPI.ru (1043)          | plus 74 more              |






# Atlas Hammercloud tests

- All NorthGrid sites participating.
- Liverpool put a lot of effort in to analysis and tweaking of network to optimise performance.
- Manchester only has 4 SEs, therefore any fine tuning of disk access masked by bandwidth limitations from the servers.




Efficiency/throughput for hc test at Liverpool  
Sept 2-4 2009





# SL5 & HEPSPREC06

Site	Expected SL5 upgrade date	HEPSPREC06 Run
Lancaster	September '09 for new kit (256 CPUs), rolling upgrade of older nodes (~600 CPUs).	No
Liverpool	88% of nodes 32 bit. September '09 for eligible nodes	Yes
Manchester	ce01 September '09, ce02 October '09	Yes
Sheffield	October '09	Yes






**GridPP**

UK Computing for Particle Physics

# Conclusions

- Steady performance over last 12 months.
- All sites ready to start migrating to SL5 by October, however ageing hardware in Liverpool a concern.
- Lancaster ready to commission new kit, and move into their new data centre.
- Manchester ready to go on a spending spree.
- Outlook positive and NorthGrid ready for LHC switch on!



NorthGrid