

# GridPP

UK Computing for Particle Physics

## The UK particle physics grid

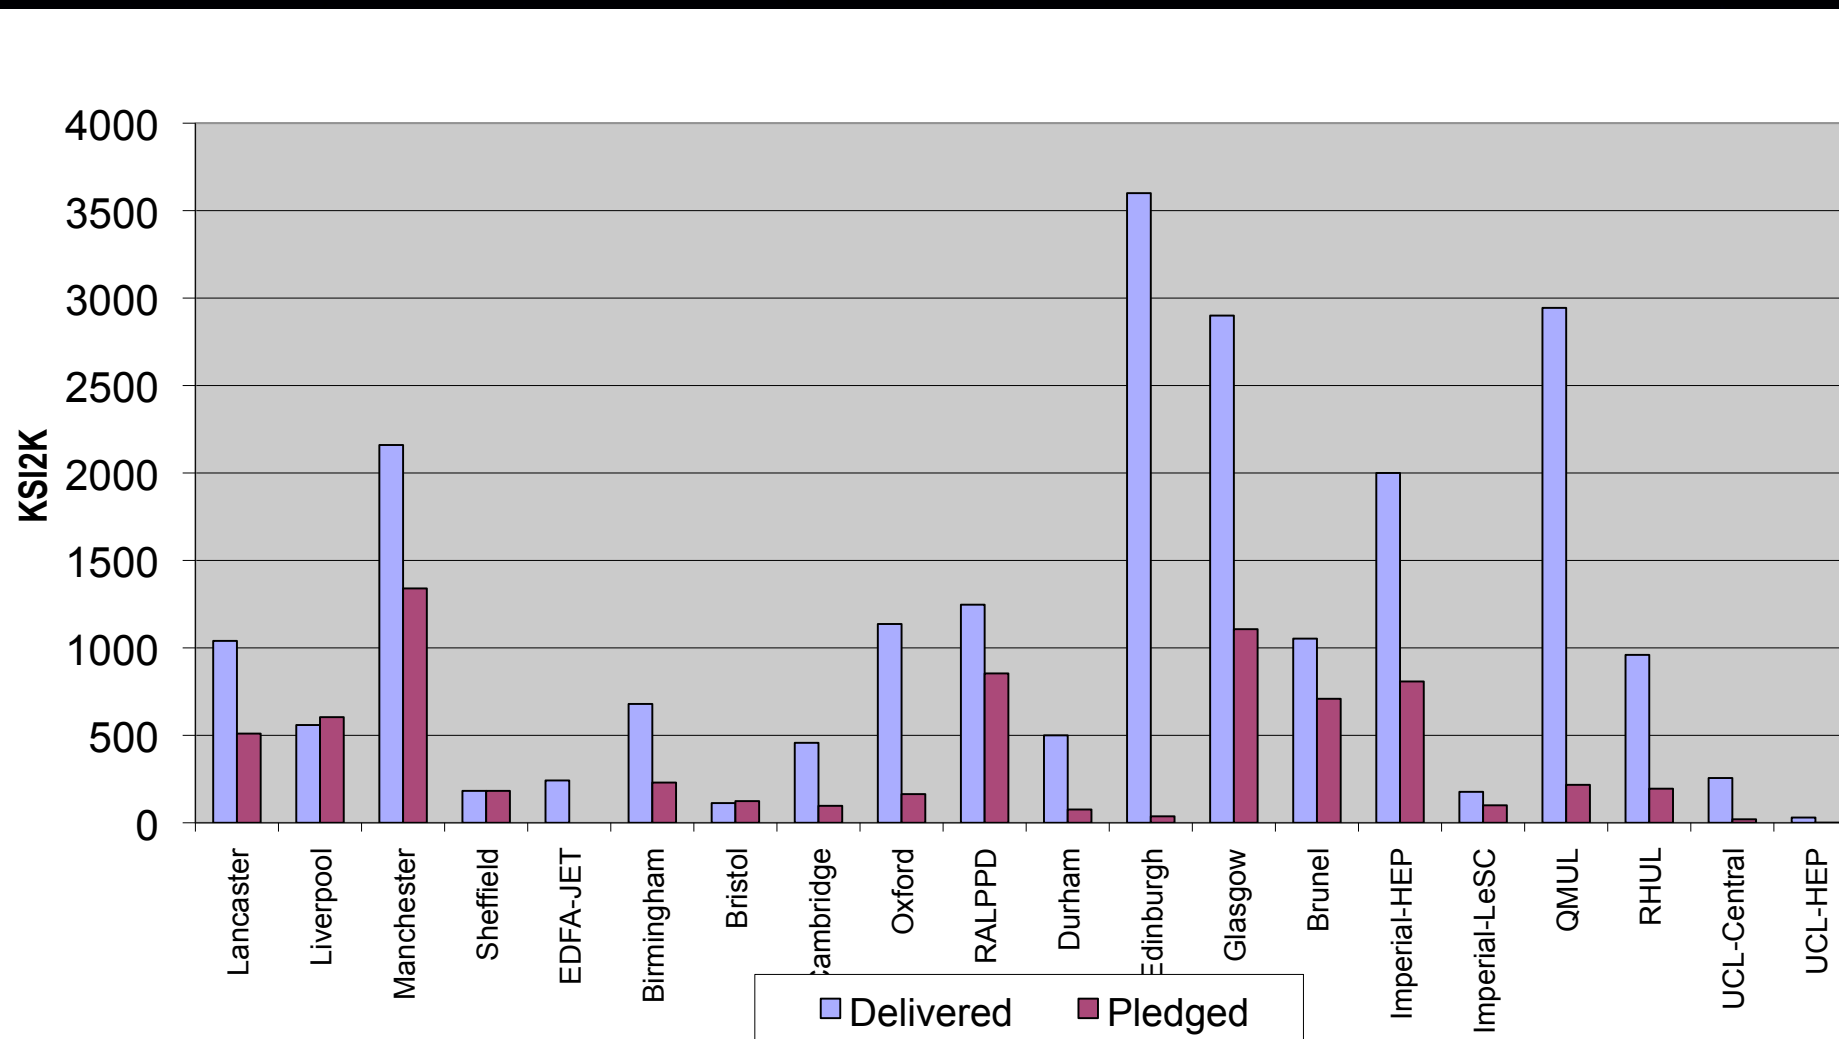
## Status & Developments

CHEP09 - Prague  
23rd March 2009

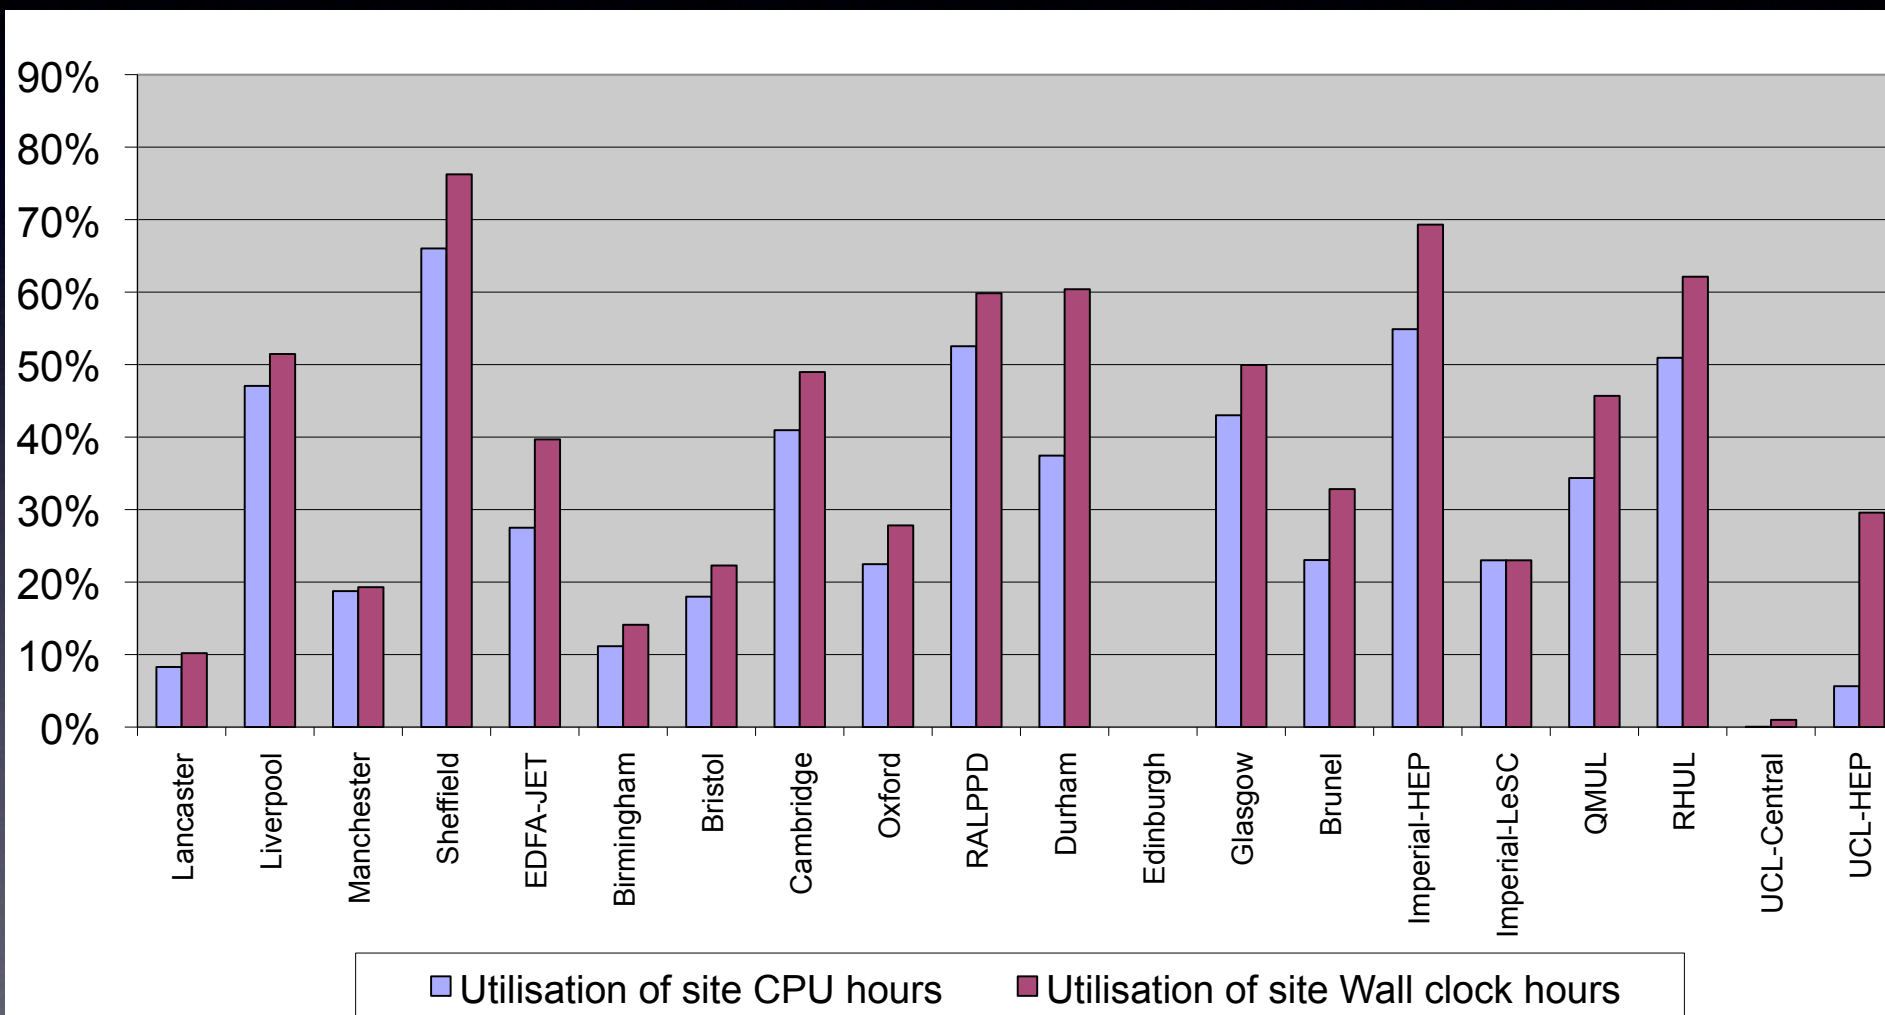
- Tier-2 resource deployments – meeting the MoU pledges
- Site stability – a GridPP perspective
- A few words on Tier-I changes
- Lessons from VO activities
- One view on improved experiment dynamics
- A GridPP theme: resilience & disasters review
- Some other stuff

## CPU at Tier-2 sites (Q408)

- CPU delivery is still well ahead of pledges across sites. Big variation in site size.

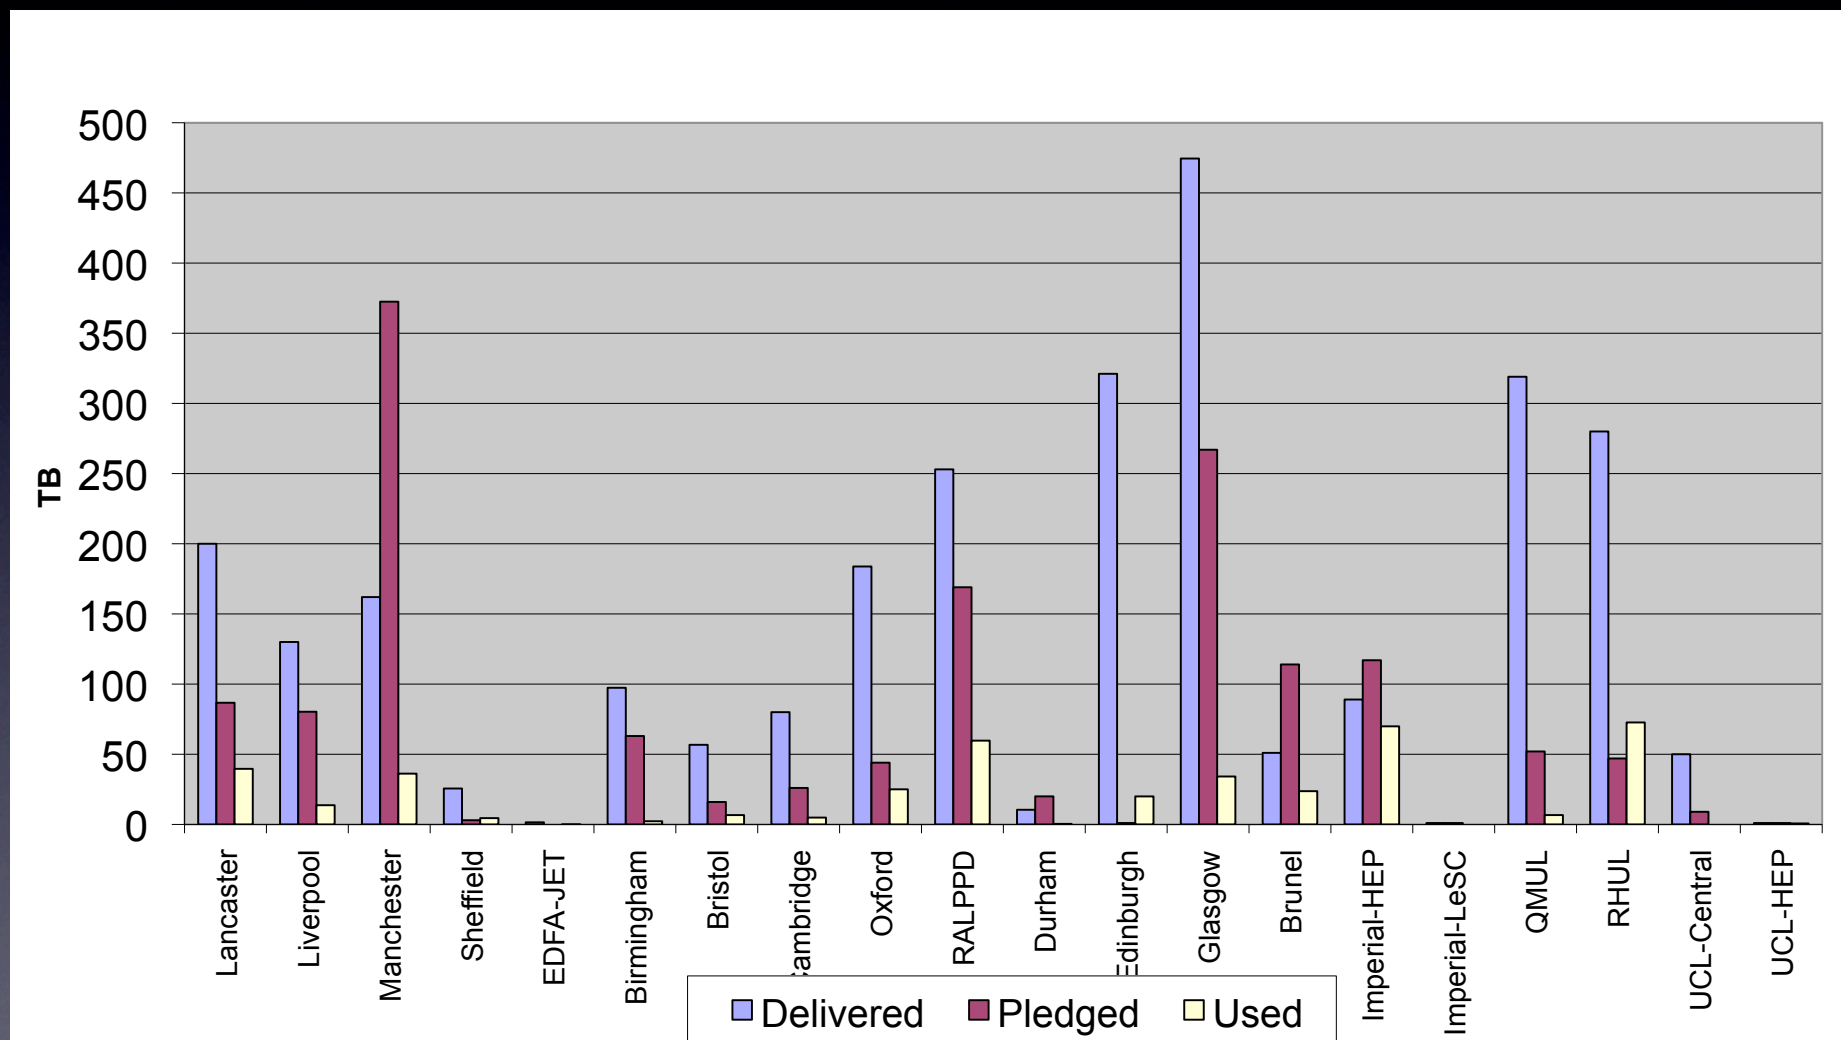


- Utilisation is typically 30-40%



## Storage at Tier-2 sites (Q408)

- In general more disk is now deployed at Tier-2s than was pledged to WLCG.



- ops test results reveal greater site stability. Not always = user experience but still good.

Hours ago:	CE										SRMv2										Availability			
	9	8	7	6	5	4	3	2	1	0	9	8	7	6	5	4	3	2	1	0	24 Hrs	Week	Month	6 Mon
EFDA-JET	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	100%	99%	100%	84%
RAL-LCG2_Tier-1	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	100%	99%	100%	99%
UKI-LT2-Brunel	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	100%	100%	100%	91%
UKI-LT2-IC-HEP	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	86%	97%	89%	90%
UKI-LT2-IC-LeSC	P	P	P	P	P	P	P	P	P	P											86%	98%	93%	89%
UKI-LT2-QMUL	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	100%	99%	98%	95%
UKI-LT2-RHUL	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	100%	99%	100%	95%
UKI-LT2-UCL-CENTRAL	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	100%	48%	87%	53%
UKI-LT2-UCL-HEP	M	M	M	M	M	M	M	M	M	M	M	M	M	M	M	M	M	M	M	M	0%	0%	26%	77%
UKI-NORTHGRID-LANCS-HEP	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	100%	91%	95%	89%
UKI-NORTHGRID-LIV-HEP	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	100%	99%	99%	98%
UKI-NORTHGRID-MAN-HEP	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	100%	100%	97%	93%
UKI-NORTHGRID-SHEF-HEP	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	100%	99%	98%	97%
UKI-SCOTGRID-DURHAM	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	100%	93%	98%	91%
UKI-SCOTGRID-ECDF	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	100%	90%	96%	97%
UKI-SCOTGRID-GLASGOW	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	95%	86%	95%	97%
UKI-SOUTHGRID-BHAM-HEP	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	100%	99%	99%	99%
UKI-SOUTHGRID-BRIS-HEP	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	100%	100%	98%	97%
UKI-SOUTHGRID-CAM-HEP	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	100%	99%	100%	94%
UKI-SOUTHGRID-OX-HEP	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	100%	80%	86%	87%
UKI-SOUTHGRID-RALPP	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	100%	94%	97%	94%
csTCDie	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	100%	99%	99%	92%
Total																					94%	90%	93%	91%

Site  
commissioning

CE space

# The longer view on availability & reliability

- Definite improvements. Sites that have (had) manpower issues or shared resources obvious.

Shared

	2Q07		3Q07		4Q07		1Q08		2Q08		3Q08		4Q08		1Q09		Recent	
EFDA-JET	75%	75%	65%	79%	85%	85%	91%	91%	69%	73%	91%	92%	77%	84%	95%	95%	84%	87%
RAL-LCG2_Tier-1	83%	84%	96%	96%	96%	96%	94%	95%	97%	97%	98%	99%	98%	99%	99%	100%	97%	97%
UKI-LT2-Brunel	98%	99%	94%	97%	99%	99%	84%	93%	92%	96%	84%	96%	91%	92%	99%	100%	91%	96%
UKI-LT2-IC-HEP	95%	95%	97%	97%	99%	99%	95%	99%	92%	94%	88%	93%	95%	95%	82%	82%	92%	94%
UKI-LT2-IC-LeSC	82%	82%	91%	91%	93%	96%	94%	99%	91%	93%	93%	95%	85%	92%	92%	92%	91%	95%
UKI-LT2-QMUL	56%	56%	54%	54%	75%	78%	40%	72%	53%	88%	92%	94%	94%	98%	96%	96%	74%	89%
UKI-LT2-RHUL	94%	96%	86%	88%	82%	94%	83%	96%	96%	98%	93%	99%	90%	93%	99%	99%	90%	97%
UKI-LT2-UCL-CENTRAL	73%	73%	65%	69%	53%	86%	49%	98%	0%		7%	85%	41%	69%	63%	69%	36%	80%
UKI-LT2-UCL-HEP	51%	51%	81%	81%	53%	81%	76%	88%	81%	86%	75%	82%	85%	85%	66%	90%	73%	86%
UKI-NORTHGRID-LANCS-HEP	54%	54%	91%	93%	95%	95%	89%	89%	93%	93%	95%	95%	83%	83%	95%	96%	91%	92%
UKI-NORTHGRID-LIV-HEP	95%	95%	91%	91%	97%	97%	70%	78%	97%	97%	93%	99%	97%	98%	98%	98%	92%	94%
UKI-NORTHGRID-MAN-HEP	99%	99%	99%	99%	87%	87%	98%	98%	86%	86%	99%	99%	92%	92%	92%	92%	92%	92%
UKI-NORTHGRID-SHEF-HEP	48%	51%	35%	50%	81%	82%	98%	99%	96%	96%	96%	96%	97%	97%	96%	97%	94%	94%
UKI-SCOTGRID-DURHAM	92%	92%	91%	92%	92%	92%	94%	96%	84%	90%	95%	95%	83%	96%	97%	97%	91%	94%
UKI-SCOTGRID-ECDF					13%	32%	66%	82%	72%	91%	82%	98%	99%	99%	95%	96%	74%	87%
UKI-SCOTGRID-GLASGOW	89%	89%	92%	93%	96%	97%	85%	86%	96%	96%	96%	96%	96%	98%	98%	97%	98%	94%
UKI-SOUTHGRID-BHAM-HEP	88%	88%	89%	89%	98%	98%	87%	91%	96%	97%	96%	97%	99%	99%	97%	99%	95%	97%
UKI-SOUTHGRID-BRIS-HEP	91%	91%	95%	95%	99%	99%	98%	99%	97%	98%	97%	97%	96%	98%	98%	99%	98%	98%
UKI-SOUTHGRID-CAM-HEP	48%	48%	75%	75%	69%	71%	93%	94%	84%	85%	84%	88%	90%	90%	97%	98%	86%	88%
UKI-SOUTHGRID-OX-HEP	79%	80%	82%	82%	89%	89%	98%	99%	90%	92%	92%	92%	87%	89%	86%	88%	91%	92%
UKI-SOUTHGRID-RALPP	93%	93%	90%	93%	94%	94%	92%	97%	91%	92%	96%	96%	94%	99%	94%	97%	93%	96%
csTCDie	75%	75%	93%	93%	91%	91%	97%	97%	91%	93%	94%	94%	88%	89%	95%	96%	92%	93%
LondonGrid	78%	79%	81%	82%	79%	91%	74%	93%	72%	93%	78%	93%	83%	91%	85%	90%	79%	92%
NorthGrid	74%	74%	79%	86%	90%	90%	89%	91%	93%	93%	96%	97%	92%	92%	95%	96%	92%	93%
ScotGrid	91%	91%	91%	93%	73%	82%	81%	88%	84%	92%	91%	96%	93%	98%	96%	97%	87%	92%
SouthGrid	80%	80%	86%	87%	90%	91%	94%	96%	92%	93%	93%	94%	93%	95%	94%	96%	93%	94%
Overall	79%	79%	83%	87%	84%	90%	87%	94%	84%	93%	89%	95%	89%	93%	92%	94%	88%	93%

# SAM results across the VOs

ops

ATLAS

CMS

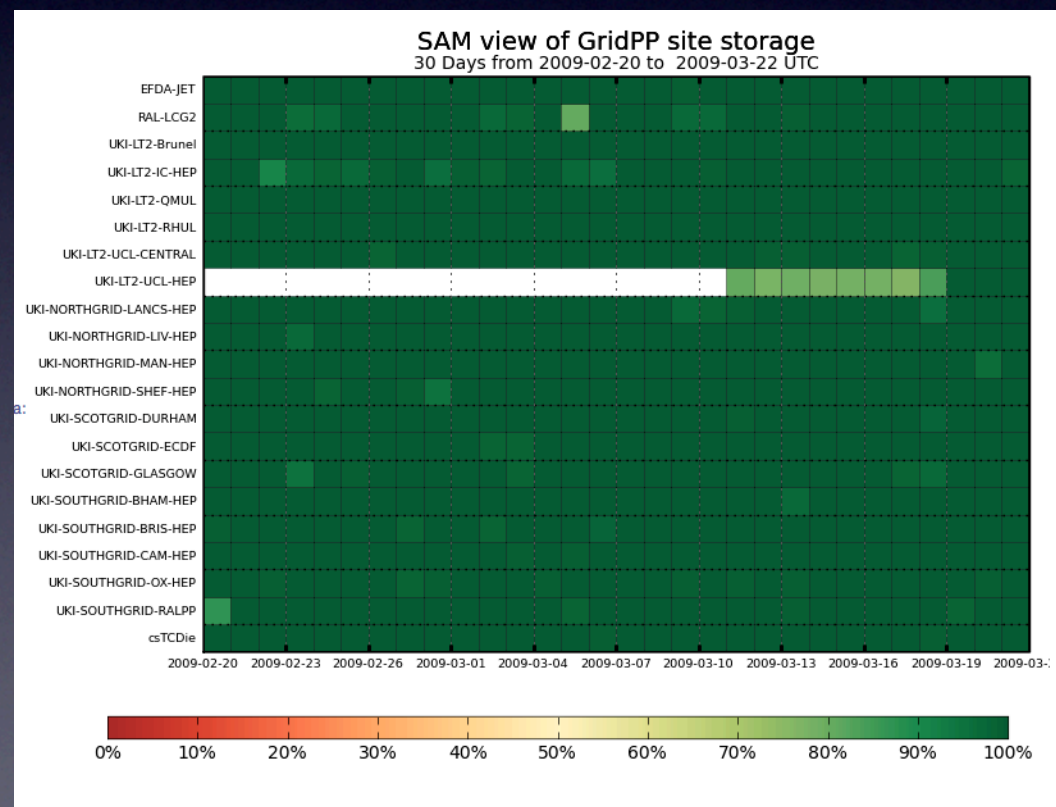
LHCb

	Availability					Availability					Availability					Availability			
	24 Hrs	Week	Month	6 Mon		24 Hrs	Week	Month	6 Mon		24 Hrs	Week	Month	6 Mon		24 Hrs	Week	Month	6 Mon
EFDA-JET	96%	99%	100%	83%		100%	98%	99%	99%		100%	100%	100%	77%		100%	100%	100%	90%
RAL-LCG2_Tier-1	96%	99%	100%	99%		96%	99%	97%	90%		100%	99%	94%	97%		100%	100%	100%	89%
UKI-LT2-Brunel	100%	100%	100%	91%		100%	99%	100%	100%		100%	100%	100%	91%		100%	100%	100%	90%
UKI-LT2-IC-HEP	96%	99%	89%	90%		96%	99%	98%	91%		91%	89%	65%	84%		100%	100%	100%	87%
UKI-LT2-IC-LeSC	100%	100%	93%	89%		100%	92%	91%	80%		100%	100%	96%	94%		100%	100%	100%	61%
UKI-LT2-QMUL	96%	99%	98%	95%		100%	88%	89%	87%		100%	100%	99%	93%		100%	75%	73%	92%
UKI-LT2-RHUL	96%	99%	100%	95%		100%	99%	100%	88%		100%	95%	86%	77%		100%	96%	99%	75%
UKI-LT2-UCL-CENTRAL	91%	48%	87%	53%		100%	32%	72%	66%		0%	0%	0%	24%		100%	62%	47%	28%
UKI-LT2-UCL-HEP	0%	0%	30%	78%		0%	0%	19%	67%		0%	0%	31%	66%		0%	0%	0%	0%
UKI-NORTHGRID-LANCS-HEP	91%	91%	95%	89%		100%	98%	94%	90%		100%	96%	98%	90%		100%	96%	99%	87%
UKI-NORTHGRID-LIV-HEP	91%	99%	99%	98%		100%	99%	100%	95%		100%	100%	100%	100%		100%	100%	99%	88%
UKI-NORTHGRID-MAN-HEP	100%	100%	97%	93%		100%	99%	98%	90%		100%	100%	97%	91%		100%	100%	99%	88%
UKI-NORTHGRID-SHEF-HEP	96%	99%	98%	97%		100%	99%	100%	94%		100%	99%	100%	99%		100%	100%	99%	89%
UKI-SCOTGRID-DURHAM	100%	93%	98%	91%		100%	97%	98%	88%		100%	99%	99%	98%		100%	96%	99%	81%
UKI-SCOTGRID-ECDF	96%	90%	96%	97%		100%	99%	98%	94%		100%	99%	93%	80%		100%	100%	98%	83%
UKI-SCOTGRID-GLASGOW	100%	88%	96%	97%		100%	93%	98%	99%		0%	0%	19%	75%		100%	96%	98%	88%
UKI-SOUTHGRID-BHAM-HEP	96%	99%	99%	99%		100%	100%	100%	98%		100%	100%	100%	100%		100%	100%	100%	85%
UKI-SOUTHGRID-BRIS-HEP	100%	100%	98%	97%		100%	97%	97%	93%		100%	100%	99%	97%		100%	100%	100%	89%
UKI-SOUTHGRID-CAM-HEP	96%	99%	100%	94%		100%	99%	96%	96%		100%	99%	100%	94%		100%	100%	100%	87%
UKI-SOUTHGRID-OX-HEP	96%	73%	85%	87%		74%	65%	90%	94%		65%	33%	57%	66%		100%	100%	98%	85%
UKI-SOUTHGRID-RALPP	96%	95%	95%	94%		100%	99%	97%	98%		100%	90%	96%	93%		100%	100%	100%	90%
csTCDie	100%	99%	98%	92%		100%	99%	98%	92%		100%	99%	99%	95%		100%	100%	100%	79%
Total	92%	90%	93%	91%		94%	89%	92%	90%		84%	82%	83%	86%		100%	96%	96%	85%

Nb. 0% can mean the VO is not enabled or supported.

Site	1Q08	2Q08	3Q08	4Q08	1Q09
UKI-IRELAND-TRINITY		38%	68%	77%	93%
UKI-LT2-Brunel-40	70%	84%	85%	87%	99%
UKI-LT2-Brunel-44	70%	84%	85%	88%	99%
UKI-LT2-IC-HEP-HEP	83%	66%	93%	94%	92%
UKI-LT2-IC-HEP-HPC	83%	66%	93%	94%	93%
UKI-LT2-IC-LeSC	82%	66%	91%	92%	92%
UKI-LT2-QMUL-01			95%	94%	88%
UKI-LT2-QMUL-02			77%	91%	98%
UKI-LT2-RHUL-01		75%	78%	92%	99%
UKI-LT2-RHUL-02		89%	81%	58%	83%
UKI-LT2-UCL-CENTRAL	92%	25%	66%	76%	93%
UKI-LT2-UCL-HEP	69%	89%	79%	56%	81%
UKI-NORTHGRID-LANCS-HEP	81%	91%	93%	94%	96%
UKI-NORTHGRID-LIV-HEP	82%	94%	94%	96%	98%
UKI-NORTHGRID-MAN-HEP-01	85%	80%	94%	93%	99%
UKI-NORTHGRID-MAN-HEP-02	94%	92%	83%	85%	96%
UKI-NORTHGRID-SHEF-HEP	94%	92%	94%	96%	97%
UKI-SCOTGRID-DURHAM	91%	76%	92%	87%	98%
UKI-SCOTGRID-ECDF	84%	89%	91%	95%	94%
UKI-SCOTGRID-GLASGOW	92%	93%	95%	97%	97%
UKI-SOUTHGRID-BHAM-ESCI		96%	93%	96%	98%
UKI-SOUTHGRID-BHAM-ESCI-04					99%
UKI-SOUTHGRID-BHAM-HEP	93%	94%	91%	96%	98%
UKI-SOUTHGRID-BRIS-HEP-01	97%	95%	96%	95%	97%
UKI-SOUTHGRID-BRIS-HEP-02	97%	95%	95%	95%	98%
UKI-SOUTHGRID-CAM-HEP	94%	87%	85%	91%	97%
UKI-SOUTHGRID-JET	96%	90%	97%	81%	96%
UKI-SOUTHGRID-OX-HEP	94%	84%	94%	93%	96%
UKI-SOUTHGRID-RALPP	92%	90%	91%	93%	96%
UKI-TIER-1-RAL-LCG2	75%	89%	41%	0%	0%
LondonGrid	78%	71%	84%	84%	92%
NorthGrid	87%	90%	92%	93%	97%
ScotGrid	89%	86%	93%	93%	96%
SouthGrid	94%	91%	93%	92%	97%

- An attempt is made to copy a 2.8MB file to each UK SE, read it back and delete it, once an hour. (Results for quarters shown on left)
- With greater experiment demands on storage (and ticketing) availability is improving



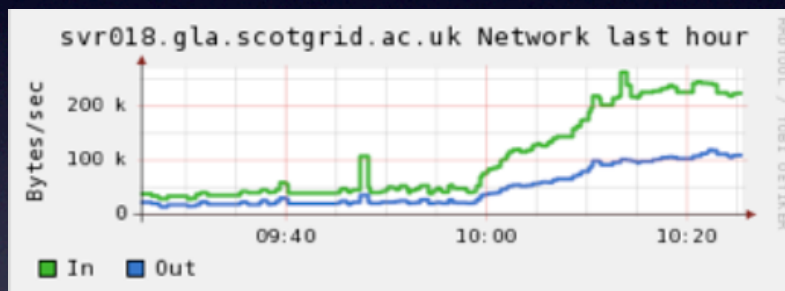
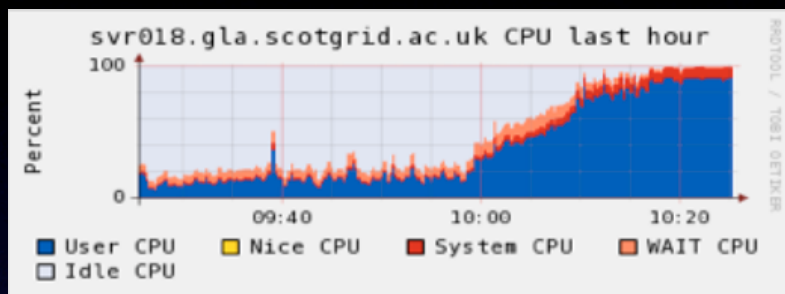
## Tier-I - status



- New machine room – almost there
- Improved CASTOR performance but v2.1.7 had many stability problems to overcome
- Operations team has grown and evolved (e.g. there is now a production manager)
- Tier-I strategy group now in place and performing a regular review of the service (appropriate metrics defined)
- Weekly storage discussion expands once per month to all user groups to discuss all Tier-I services and issues directly
- Incident reporting well developed. Response plans maturing.

## Experiment driven issues

### DPM headnode measures before optimisation



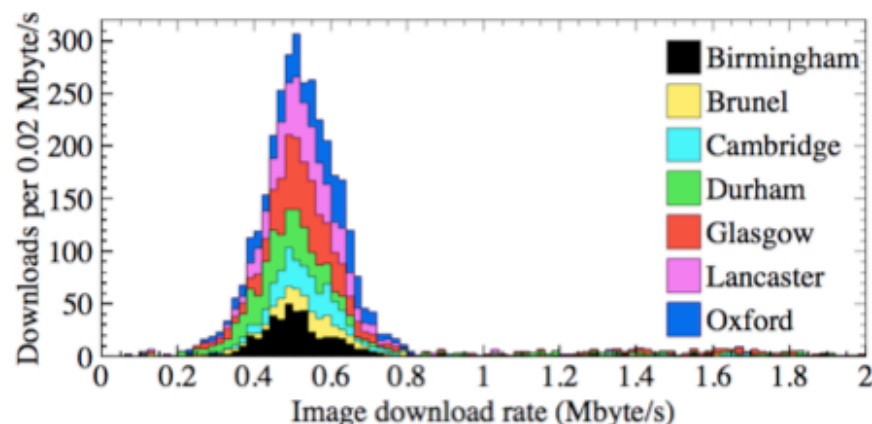
[See CHEP poster on Thursday](#): Optimised access to user analysis data using the gLite DPM

Current status of spacetokens: [http://wn3.epcc.ed.ac.uk/srm/xml/srm\\_token\\_acls\\_table](http://wn3.epcc.ed.ac.uk/srm/xml/srm_token_acls_table)

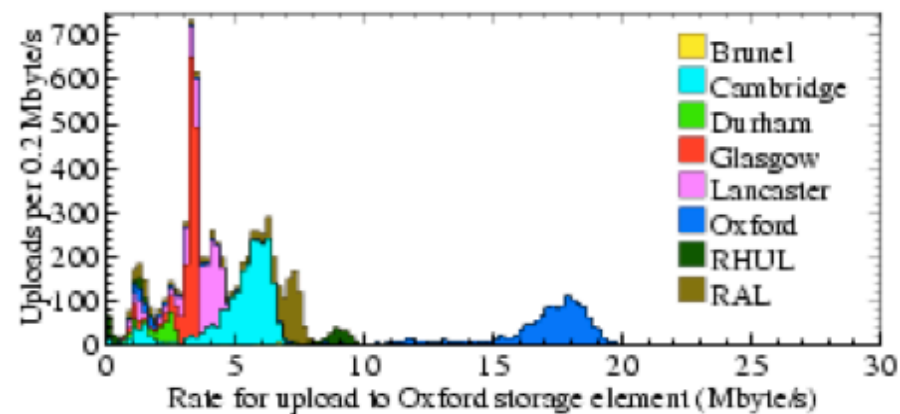
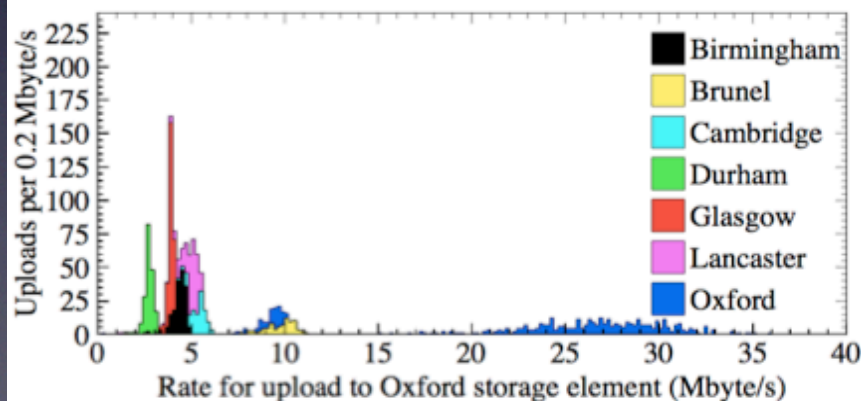
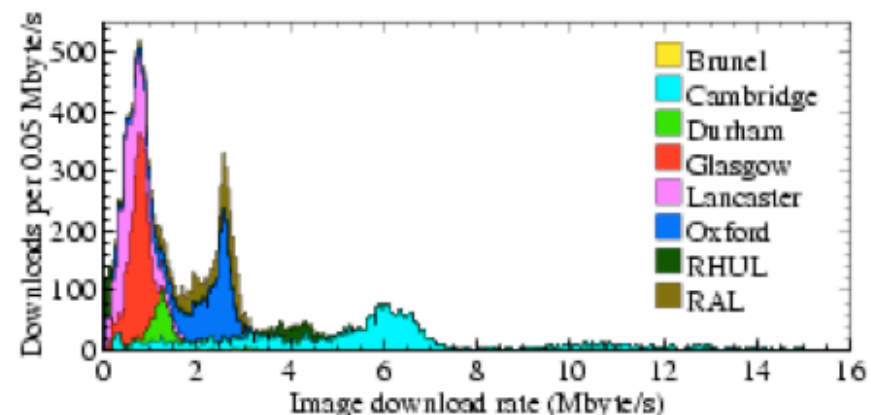
- ATLAS hammer cloud tests very useful. We always knew that concurrent file access would present new problems. Interesting to compare results across sites
- Several sites have seen “interference” from non-HEP storage requests (jobs basically pointing at a single SE). Raises issues about contention.
- Deployment of space tokens to divide storage advancing and showing up issues:
  - Still need multi-role access control
  - Monitoring has been piecemeal. Who tracks the requests?
  - Client tool integration and user education is work in progress.
  - Several management issues: draining, tools for resize & merge operations. Changing the writing group.

- It has been useful to get a view on site performance from non-LHC VO's.

July-August 2008: 3638 jobs



November-December 2008: 16733 jobs



Thanks to Karl Harrison and [Imense](#) for these findings

# Experiment blacklisting of sites

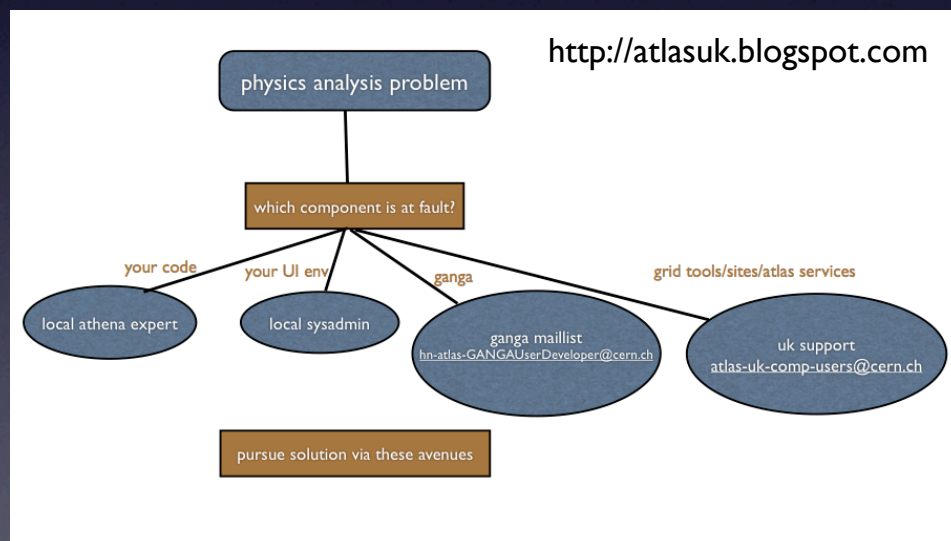
	Most Recent			24 Hrs			Week			Month			6 Months		
	atlas	cms	lhcb	atlas	cms	lhcb	atlas	cms	lhcb	atlas	cms	lhcb	atlas	cms	lhcb
UKI-IRELAND-TRINITY				0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
UKI-LT2-Brunel-40				0%	0%	0%	0%	1%	0%	0%	5%	0%	3%	7%	0%
UKI-LT2-Brunel-44				0%	0%	0%	0%	0%	0%	0%	3%	0%	5%	9%	0%
UKI-LT2-IC-HEP-HEP				0%	4%	0%	0%	1%	0%	0%	0%	0%	1%	5%	0%
UKI-LT2-IC-HEP-HPC				0%	0%	0%	0%	0%	0%	0%	0%	0%	5%	5%	0%
UKI-LT2-IC-LeSC				0%	0%	0%	0%	29%	0%	0%	67%	0%	0%	37%	0%
UKI-LT2-QMUL-01				0%	21%	0%	0%	41%	0%	0%	17%	0%	0%	16%	0%
UKI-LT2-QMUL-02				0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
UKI-LT2-RHUL-01		X		0%	100%	0%	0%	100%	0%	0%	87%	0%	1%	92%	0%
UKI-LT2-RHUL-02				0%	0%	0%	0%	5%	0%	0%	1%	0%	0%	8%	0%
UKI-LT2-UCL-CENTRAL				0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
UKI-LT2-UCL-HEP				0%	0%	0%	0%	0%	0%	0%	0%	0%	1%	0%	0%
UKI-NORTHGRID-LANCS-HEP				0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
UKI-NORTHGRID-LIV-HEP				0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
UKI-NORTHGRID-MAN-HEP-01				0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
UKI-NORTHGRID-MAN-HEP-02				0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	12%	0%
UKI-NORTHGRID-SHEF-HEP				0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
UKI-SCOTGRID-DURHAM				0%	0%	0%	0%	1%	0%	0%	0%	0%	0%	3%	0%
UKI-SCOTGRID-ECDF				0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
UKI-SCOTGRID-GLASGOW				0%	4%	0%	0%	5%	0%	0%	47%	0%	0%	9%	0%
UKI-SOUTHGRID-BHAM-ESCI				0%	0%	0%	0%	0%	0%	0%	0%	0%	14%	14%	0%
UKI-SOUTHGRID-BHAM-ESCI-04				0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
UKI-SOUTHGRID-BHAM-HEP				0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
UKI-SOUTHGRID-BRIS-HEP-01				0%	0%	0%	0%	0%	0%	0%	3%	0%	1%	7%	0%
UKI-SOUTHGRID-BRIS-HEP-02				0%	0%	0%	0%	9%	0%	0%	17%	0%	0%	21%	0%
UKI-SOUTHGRID-CAM-HEP				0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
UKI-SOUTHGRID-JET				0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
UKI-SOUTHGRID-OX-HEP				0%	35%	0%	0%	56%	0%	0%	32%	0%	0%	14%	0%
UKI-SOUTHGRID-RALPP	X	X		100%	100%	0%	82%	82%	0%	23%	30%	0%	5%	12%	0%
UKI-TIER-1-RAL-LCG2				0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%

- Freedom of choice for Resources does not present the actual site status in the experiment world
- Blacklisting can be traced if you know the right experiment links and how to read the information presented.
- The site dashboard views discussed at the WLCG workshop will help highlight problems
- ... but having Nagios probes associated with the site status per experiment is still needed.

UK coordination groups are now starting to get to grips with regional demands providing effective interaction with both users and infrastructure support. GridPP2I had a focus on grid usage: <http://www.gridpp.ac.uk/gridpp2I/>

## Discussion excerpts:

LHCb view: system instabilities were down, success rates were higher, Ganga5 had been released, users seemed happier BUT 80% success rates were not acceptable? Data access still problematic



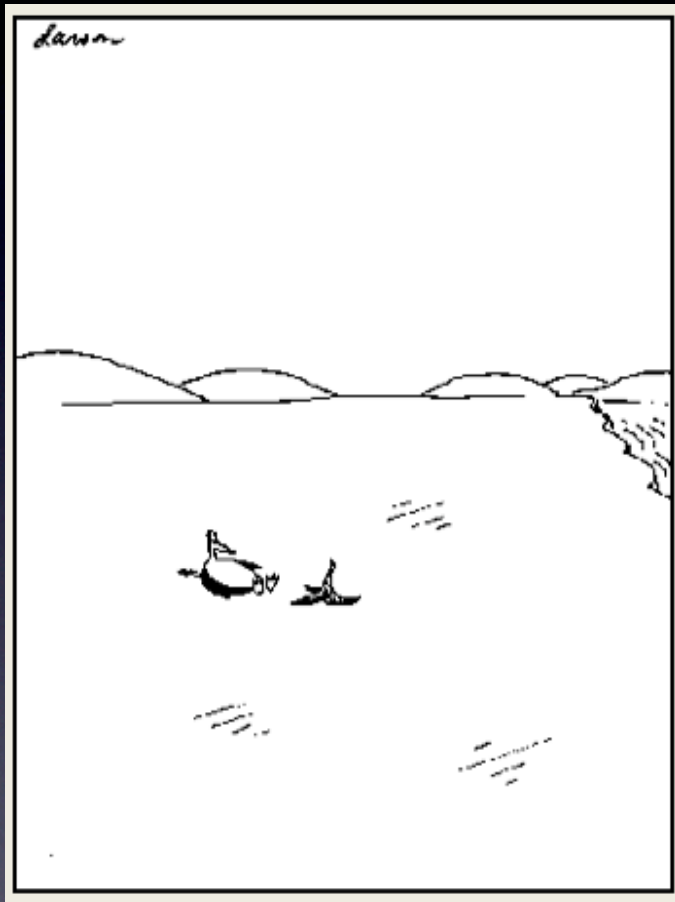
ATLAS: User tutorials very necessary. Users need improved understanding of workflows and to be taught “good practice”. For example do they know how to recover when they make a mistake. Pulling together error messages for reference was agreed to be useful. Many users try other routes than persist with the grid if problems encountered.

General: Inefficient user code (reused!) Lack of storage quotas. Coping with a “consistent low level of intermittent errors”. Will the FTS cope once user levels pick up.

We found there was a need for a UK user mailing list.



## Places we can slip up



- Knowns: New architectures - e.g. introduction of larger multi-cores and how they affect services. Can they be exploited and do some services need uncoupling?
- Known unknowns such as data management. It is going to change.
- Unknown unknowns look at LEP startup and end.
- Realities of data taking and production are bound to be different to planned exercises including many more users and the pressures of getting some results.
- Movement towards the end of EGEE and uncertainties of funding and the constraints of contract policies. Critical operation and support functions must not be disrupted.

## Infrastructure/service

1. Loss of a grid service
2. Site failure
3. Software failure - middleware
4. Software failure - load (exposes bugs/limitations)
5. Software - OS-related (external software dependencies)
6. Software - Usage (resource exhaustion, lack of quotas)
7. Loss of key staff
8. Network failure - OPN
9. Network failure - JANET
10. Security incident
11. Procurement failure

## Experiment

- Non-scalability or general failure of the Grid data transfer/placement
- Non-scalability or general failure of the Grid WMS
- Non-scalability or general failure of the metadata/bookkeeping system
- Medium-term loss of data storage resources
- Medium-term loss of CPU resources
- Long-term loss of data or data storage resources
- Long-term loss of CPU resources
- Medium- or long-term loss of wide area network
- Grid security incident
- Mis-estimation of resource requirements

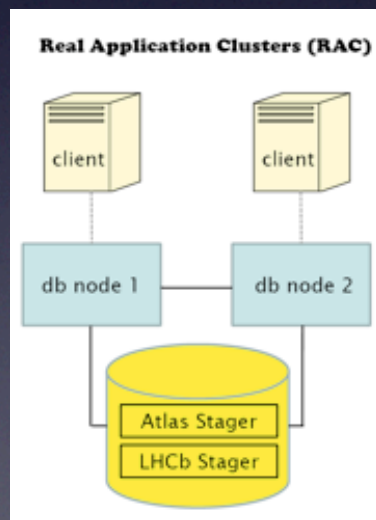


**"Well, thank God we made it out in time. ...  
'Course, now we're equally screwed."**

*Worth* 1000.com

[GridPP22 collaboration meeting](#) (next week) will focus on service resilience and disaster planning

- Increase the hardware's capacity to handle faults
- Duplicate services or machines
- Implement automatic restarts
- Provide fast intervention
- Fully investigate failures



Oracle RACs are a good solution for several service needs

FTS: awaiting automated failover. Currently 5 frontends with round robin. Oracle RAC for database.

WMS: redundant frontends with two independent LBs. Other instances at T2s

LFC: Multiple frontends and use of Oracle Server for DB. [Offsite backup?]

tBDII: Several across sites. Use load-balanced rotating alias

MyProxy: Dual machine with round robin [use of myProxy lists. DNS load balanced setup]

CE: Multiple CEs but single scheduler

WNs: Job interference and resource exhaustion

VOBoxes: Easy installs but ...

Databases: Have RACs. Oracle Data Guard is a next step

Networking: LANs moving to dual links to central switch.

OPN options often reviewed.

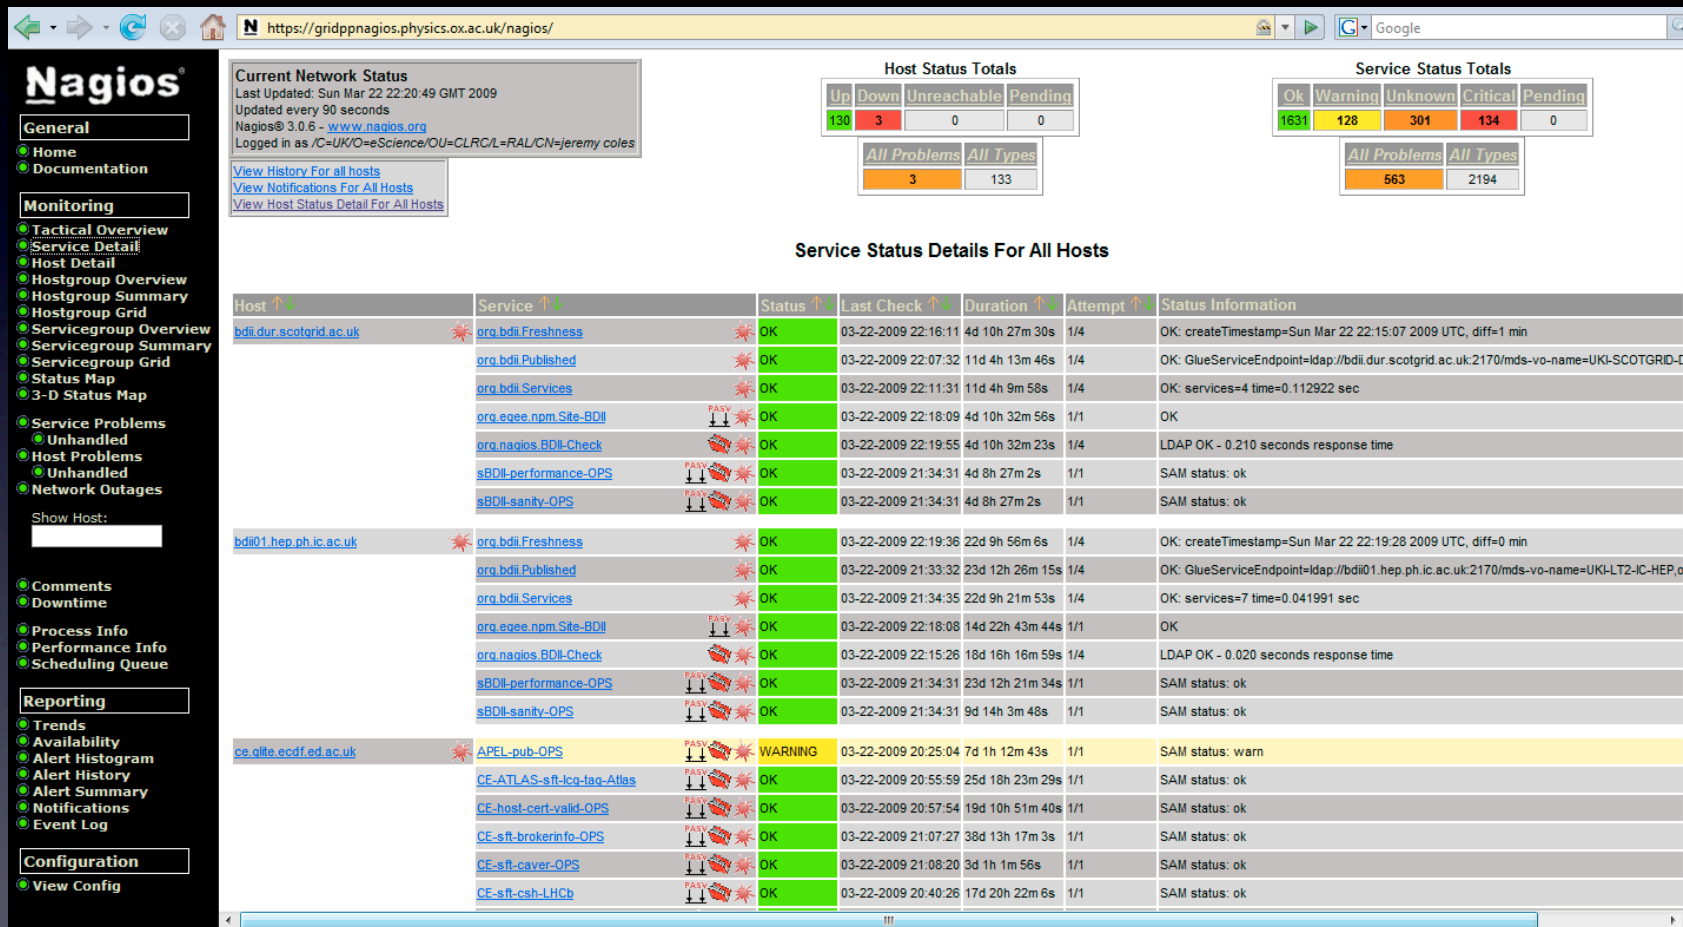
Tier-2 SEs: [distribute the headnodes vs hot spares]

GOCDDB: Oracle cluster. Replica at CNAF.

CA: [specialist components]

Etc....(Experiment) components outside of the UK...

- We now have a UK view across all sites. The amount of detail can be overwhelming and this applies to WLCG/EGEE/experiment monitoring overall.



The screenshot shows the Nagios web interface for GridPP. The left sidebar contains navigation links for General, Monitoring, and Reporting. The main content area displays 'Current Network Status', 'Host Status Totals', 'Service Status Totals', and a detailed table of 'Service Status Details For All Hosts'.

**Current Network Status**  
Last Updated: Sun Mar 22 22:20:49 GMT 2009  
Updated every 90 seconds  
Nagios® 3.0.6 - [www.nagios.org](http://www.nagios.org)  
Logged in as /C=UK/O=eScience/OU=CLRC/L=RAL/CN=jeremy.coles

**Host Status Totals**

Up	Down	Unreachable	Pending
130	3	0	0

**Service Status Totals**

Ok	Warning	Unknown	Critical	Pending
1631	128	301	134	0

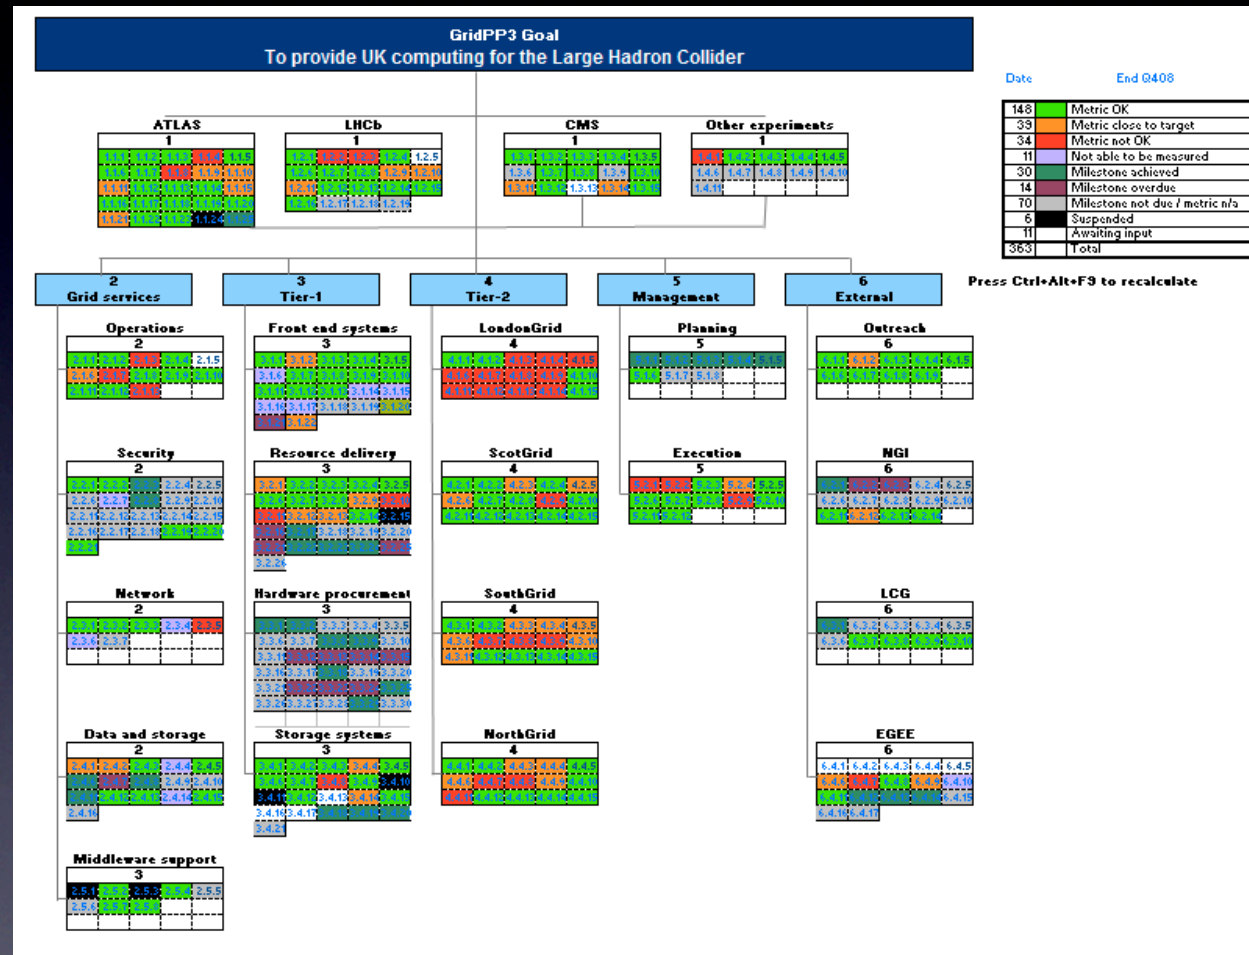
**Service Status Details For All Hosts**

Host	Service	Status	Last Check	Duration	Attempt	Status Information
bdi01.scotgrid.ac.uk	org.bdi.Freshness	OK	03-22-2009 22:16:11	4d 10h 27m 30s	1/4	OK: createTimestamp=Sun Mar 22 22:15:07 2009 UTC, diff=1 min
	org.bdi.Published	OK	03-22-2009 22:07:32	11d 4h 13m 46s	1/4	OK: GlueServiceEndpoint=ldap://bdi01.scotgrid.ac.uk:2170/mds-vo-name=UKI-SCOTGRID-I
	org.bdi.Services	OK	03-22-2009 22:11:31	11d 4h 9m 58s	1/4	OK: services=4 time=0.112922 sec
	org.egge.npm.Site-BDI	OK	03-22-2009 22:18:09	4d 10h 32m 56s	1/1	OK
	org.nagios.BDI-Check	OK	03-22-2009 22:19:55	4d 10h 32m 23s	1/4	LDAP OK - 0.210 seconds response time
	sBDI-performance-OPS	OK	03-22-2009 21:34:31	4d 8h 27m 2s	1/1	SAM status: ok
bdi01.hep.ph.ic.ac.uk	org.bdi.Freshness	OK	03-22-2009 22:19:36	22d 9h 56m 6s	1/4	OK: createTimestamp=Sun Mar 22 22:19:28 2009 UTC, diff=0 min
	org.bdi.Published	OK	03-22-2009 21:33:32	23d 12h 26m 15s	1/4	OK: GlueServiceEndpoint=ldap://bdi01.hep.ph.ic.ac.uk:2170/mds-vo-name=UKI-LT2-IC-HEP,0
	org.bdi.Services	OK	03-22-2009 21:34:35	22d 9h 21m 53s	1/4	OK: services=7 time=0.041991 sec
	org.egge.npm.Site-BDI	OK	03-22-2009 22:18:08	14d 22h 43m 44s	1/1	OK
	org.nagios.BDI-Check	OK	03-22-2009 22:15:26	18d 16h 16m 59s	1/4	LDAP OK - 0.020 seconds response time
	sBDI-performance-OPS	OK	03-22-2009 21:34:31	23d 12h 21m 34s	1/1	SAM status: ok
ce.gite.ecdf.ed.ac.uk	sBDI-sanity-OPS	OK	03-22-2009 21:34:31	9d 14h 3m 48s	1/1	SAM status: ok
	APEL-pub-OPS	WARNING	03-22-2009 20:25:04	7d 1h 12m 43s	1/1	SAM status: warn
	CE-ATLAS-sft-lcq-lag-Atlas	OK	03-22-2009 20:55:59	25d 18h 23m 29s	1/1	SAM status: ok
	CE-host-cert-valid-OPS	OK	03-22-2009 20:57:54	19d 10h 51m 40s	1/1	SAM status: ok
	CE-sft-brokerinfo-OPS	OK	03-22-2009 21:07:27	38d 13h 17m 3s	1/1	SAM status: ok
	CE-sft-caver-OPS	OK	03-22-2009 21:08:20	3d 1h 1m 56s	1/1	SAM status: ok
	CE-sft-csh-LHCb	OK	03-22-2009 20:40:26	17d 20h 22m 6s	1/1	SAM status: ok

We need to reduce the number of places a sysadmin needs to look for INFORMATION  
Our links pages keep growing: [http://www.gridpp.ac.uk/wiki/Links\\_Monitoring\\_pages](http://www.gridpp.ac.uk/wiki/Links_Monitoring_pages)

For those unfamiliar with the operations area these are some of the things we are currently looking at...

- Moving away from the LCG-CE to a gLite-CE
- Rollout of glxexec
- Moving worker nodes to SL5
- Moving to the new spec benchmarks
- Upgrading Storage Elements to SL5
- Increased memory demands (we did review site policies:  
[http://www.gridpp.ac.uk/wiki/Site\\_information](http://www.gridpp.ac.uk/wiki/Site_information))
- Learning from incidents (they are now logged  
<http://www.gridpp.ac.uk/wiki/Category:Incidents>)
- Virtualisation of nodes
- A smooth transition to the National Grid Initiatives/European Grid Initiative model



Emphasis now on running (not building) a grid. Project Map has more metrics less milestones. Structure focuses on services. Targets are set to be challenging.

Last map review highlighted issues such as: CASTOR upgrade, job efficiencies, utilisation, speed of deployment (new disk), Some procurement and manpower issues at T2s.

The project also revisited its risk register - staffing remains one of the key risks.

- GridPP sites have generally met or exceeded pledged resources
- Availability and reliability continue to improve but resolution times are sometimes too long
- The Tier-1 is making steady progress in many areas
- UK experiment operations has greatly improved in the last year
- We have caught glimpses of some potential problems from the experiment testing
- GridPP is working to improve its position with resilience and disaster planning
- Our project as a whole is now much more focused on improved grid performance