

Strategy for a medium-sized site with limited effort?

GridPP 38 – University of Sussex

6-7 April 2017

Simon George



ROYAL
HOLLOWAY
UNIVERSITY
OF LONDON

Summary of hardware and support

Tier2

- 37 kHS06 in 140 WN; Cream/Torque CE x 2, moving to ARC/HTC
- 1.4 PB DPM SE in 47 servers
- 8 misc servers including 3 VM hosts running standard network services, provisioning, Grid services
- Dedicated 10 Gb/s uplink (+1 failover), 10 Gb/s network
- 8 racks in modern machine room provided by central IT service
- 1 FTE support, falling to 0.5 FTE during GridPP5
- Main VO: ATLAS


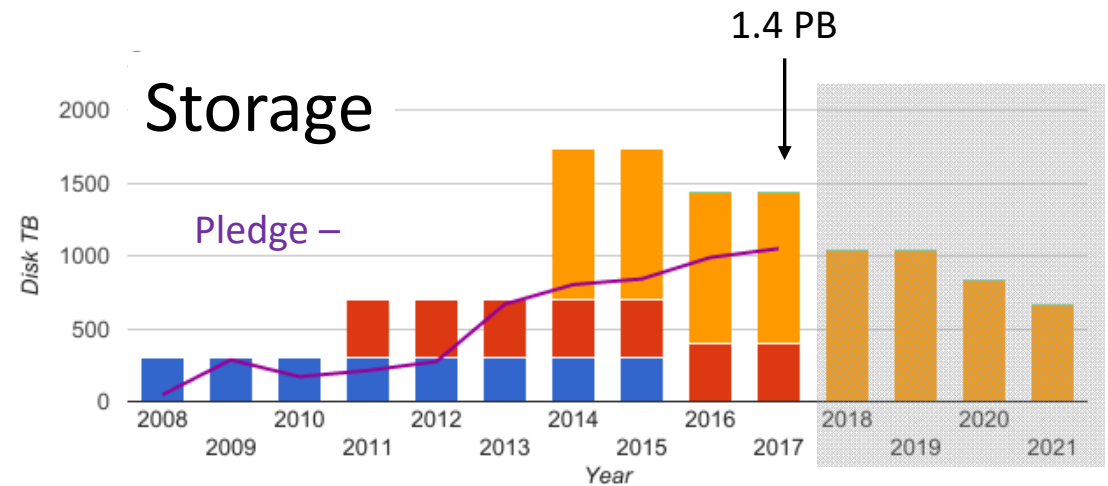
Tier3

- 6 kHS06 in ~100 old WN, mostly hand downs from Tier2, some upgraded; Torque
- Storage:
 - 60 TB Hadoop using 69 WNs
 - 170 TB NFS scratch over 7 servers
 - 10 TB NFS Home
- 11 servers running standard network services, mainly as VMs
- 8 racks packed into home-made machine room, a long way from Tier2 on 1 Gb/s link
- CG-funded 0.5 FTE sysadmin, small fractions of core research staff
- 0.5 FTE from other grants across department, with additional responsibilities
- Main users: ATLAS, dark matter, accelerator



Tier2 hardware history

- GridPP and RHUL both invested about £600k each over last 10 years
- Significant leveraging of college infrastructure funds
- RHUL also provides machine room facilities, electricity and network at no direct cost
- Machine room capacity now a limiting factor
- **All these must continue to maintain/grow site**
- But even without further hardware investment site will remain large for years
- Note: GridPP5 tranche 1 spend not shown




GridPP context

	Disk		Elapsed Time		CPU Time		Average		
	Site	TB	%	kHS06	%	kHS06	%	1	2
medium	UKI-LT2-IC-HEP	3595	14.00%	22025	9.60%	18265	9.20%	11.80%	11.60%
	UKI-SOUTHGRID-RALPP	3016	11.80%	21821	9.60%	17126	8.70%	10.70%	10.20%
	UKI-SCOTGRID-GLASGOW	2769	10.80%	23154	10.10%	17740	9.00%	10.50%	9.90%
	UKI-NORTHGRID-MAN-HEP	2334	9.10%	24892	10.90%	22093	11.20%	10.00%	10.10%
	UKI-LT2-QMUL	3237	12.60%	15858	6.90%	16371	8.30%	9.80%	10.50%
	UKI-NORTHGRID-LANCS-HEP	2251	8.80%	20325	8.90%	17672	8.90%	8.80%	8.90%
	UKI-LT2-RHUL	1453	5.70%	18566	8.10%	16143	8.20%	6.90%	6.90%
	UKI-LT2-BRUNEL	1459	5.70%	15650	6.90%	11430	5.80%	6.30%	5.70%
	UKI-SCOTGRID-DURHAM	255	1.00%	19485	8.50%	19493	9.90%	4.80%	5.40%
	UKI-NORTHGRID-LIV-HEP	1224	4.80%	10523	4.60%	9559	4.80%	4.70%	4.80%
	UKI-SOUTHGRID-OX-HEP	941	3.70%	12972	5.70%	11135	5.60%	4.70%	4.70%
	UKI-SCOTGRID-ECDF	1281	5.00%	4449	1.90%	4060	2.10%	3.50%	3.50%
	UKI-SOUTHGRID-BHAM-HEP	865	3.40%	7068	3.10%	6484	3.30%	3.20%	3.30%
	UKI-NORTHGRID-SHEF-HEP	531	2.10%	2247	1.00%	2051	1.00%	1.50%	1.60%
	UKI-SOUTHGRID-CAM-HEP	302	1.20%	3286	1.40%	2912	1.50%	1.30%	1.30%
	UKI-SOUTHGRID-BRIS-HEP	44	0.20%	5071	2.20%	4085	2.10%	1.20%	1.10%
	UKI-SOUTHGRID-SUSX	79	0.30%	918	0.40%	834	0.40%	0.40%	0.40%
	UKI-LT2-UCL-HEP	0	0.00%	78	0.00%	70	0.00%	0.00%	0.00%
	EFDA-JET	0	0.00%	40	0.00%	39	0.00%	0.00%	0.00%

ATLAS VO context

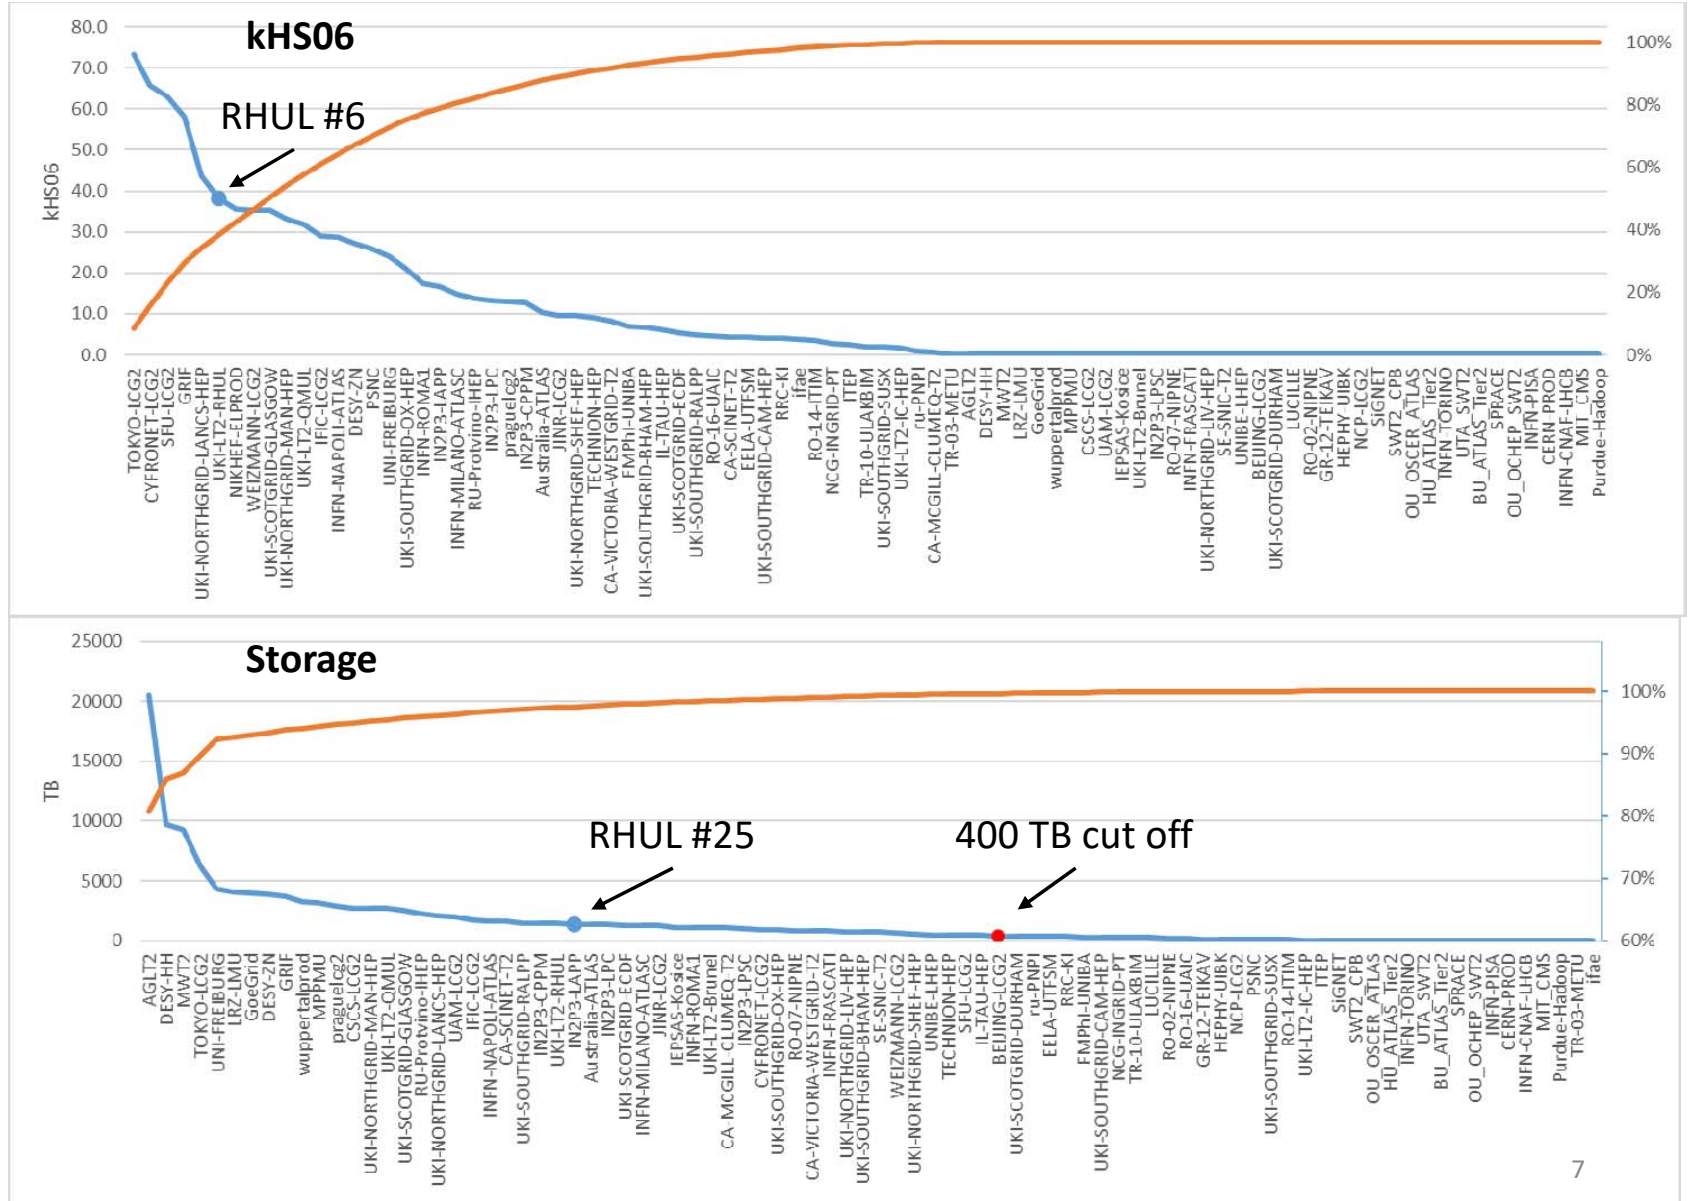
- A 1.4 PB SE is a large Tier2 site as far as ATLAS is concerned
- Sites under 0.4 TB are recommended to consolidate with regional neighbours or become cache-based – Jan 2017 sites Jamboree

Available storage at Tier 2 sites




More efficient to have larger and fewer storage end-points
2 possible categories : 'Cache based' & 'large' Tier 2s
Some Tier 2s are already larger than some Tier 1s

ATLAS Tier2 data in REBUS



Where does the effort go?

- Roughly estimated breakdown of effort of 1 person managing a Tier2 site, from RHUL experience
- How does it scale?
 - Hardware support scales with # vendors and # units
 - O/s and cluster management scales with the number of different system profiles
 - Grid middleware support scales with the number of different types of service supported
- About half of the time spent managing systems goes on significant changes, not routine maintenance
- Cutting a major service completely such as storage, including the hardware it runs on, could save 10-15%
- Reducing the scale of a service only saves on hardware support which is small
- The effort to maintain a service depends on its maturity:
- Storage, once considered a burden, is now mature and low maintenance. It does however still take us more effort to procure



Ideas I have heard about to save effort

- VAC-in-a-Box
 - Huge potential
 - Should do away with need to run CE, batch system, BDII, accounting, ARGUS, squid
 - Does it yet support all the VOs we need? Until it completely obviates the need for a traditional batch system, it only adds to complexity.
 - Would GridPP mandate it for all supported VOs?
- Cache-based storage
 - RHUL seems too big to become compute-only site with ~100 TB cache storage?
 - Network to other sites is 10 Gb/s: good but not good enough to replace local storage (> 160 Gb/s internal bandwidth)
 - Is load on other 'large' Tier2s viable?
 - What to do with current storage nodes (1.4 PB)

What else can we do?

Thinking creatively...

(several ideas shamelessly stolen)

Cross-site support of existing sites

- Some of support effort for RHUL provided by larger sites in LondonGrid Tier2
- Needs to be real significant effort from other site(s), with priority agreed by line managers
- May be worth harmonising site design to make this easier
- Usual challenges of remote working and physically distributed team

Form new collaboratively-run site

- Form a new joint LondonGrid Tier2 site at, say, the JISC datacentre in Slough
- Co-locate say storage from one site with compute from another
- Combine sysadmin effort between participating sites to manage it
- Would also address local pressure on space and easier access to higher network bandwidth
- Usual challenges of remote working and physically distributed team

Develop model for shared site support

- For both scenarios above
- Groups each take on responsibility for some of the services
- Or share admin as a team across all services
- Metrics and funding models need to reflect and motivate this model:
currently sites compete for funding

Crack squads to visit sites for planned major interventions

- Maintenance is potentially feasible with less effort but R&D, major upgrades and changes are not
- Ad hoc teams made up from large Tier 2 people, few days away, few times per year on average
- E.g. storage group
- Would strengthen GridPP as a community working together
- Sometimes work could be done remotely, as long as team remains focused

Centrally recommended system designs and technology choices leading to more in common between sites

- To make all the above easier
- Would save significant time for medium sites to have clear, researched options
- Reduce proliferation e.g. of storage designs
- Different technology choices across London effectively limit scope for cross-site support
 - StoRM/Lustre at QMUL, dCache at IC, DPM at Brunel & RHUL

My conclusions

- Medium sites do not yet have a clear path ahead
- RHUL still looks like a large site to ATLAS (need more storage)
- Potential solutions are there (ViaB, cache storage) but we need to take bolder steps to make them work
- Central support for changes might help
- I think we need to explore deeper collaboration between sites, however it is discouraged by the current proliferation of designs and the hardware funding model

Discuss...