



GridPP

UK Computing for Particle Physics

UK Tier 2s: From Service Challenge 3 to the LHC

Graeme Stewart
GridPP, University of Glasgow

Week 09 Week 10 Week 11
CPUS ■ Running Processes
History last month



- UKI Operational Review
- Current Status
- SC3 - Experiences and (ongoing) Lessons
- SC4 - Plans and Preparations
- Current Issues and Concerns
- Conclusions





GridPP

UK Computing for Particle Physics

UKI Operations

ScotGrid

LHC Era Resources

- Durham
- Edinburgh (SC3 LHCb)
- Glasgow

will have > 1000 nodes

NorthGrid

- Lancaster (SC3 ATLAS)
- Liverpool
- Manchester
- Sheffield

will have > 1000 nodes

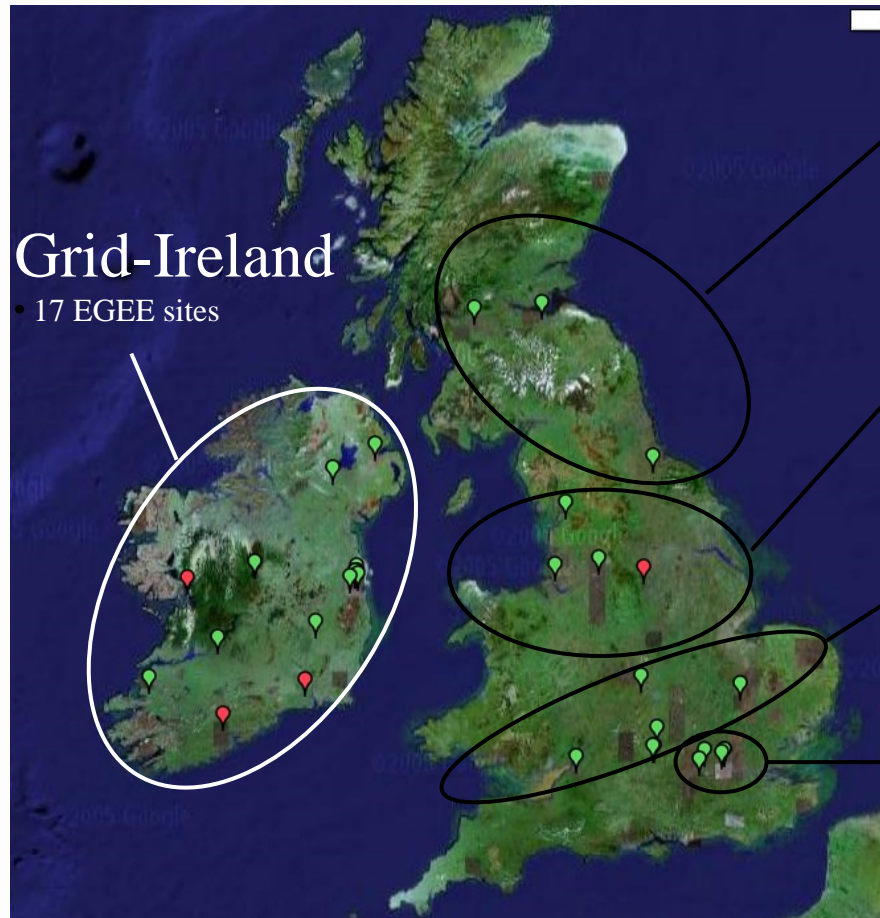
SouthGrid

- Birmingham
- Bristol
- Cambridge
- Oxford
- RAL PP

will have > 1000 nodes

London Tier-2

- Brunel
- Imperial College UL (SC3 CMS)
- Queen Mary UL
- Royal Holloway UL
- University College UL



Grid-Ireland

- 17 EGEE sites





General status

- Most Tier-2 sites running LCG middleware for 18 months.
- At the moment availability is good - most sites are green most of the time.
- More regular releases have helped achieve reasonable upgrade rates. Even when dates slipped *communication* about status was very helpful when it happened.
- Quality of release and ease of install has a strong correlation with resource availability!
- Transparent, lightweight installation tools like YAIM were very helpful - our sysadmins found YAIM easy to use, debug and customise.

AsiaPacific	BNL	CERN	Canada	CentralEurope	China	FNAL	France	Germany	Switzerland	HP
Italy	NorthernEurope	OSG	Pakistan	Russia	SouthEasternEurope	SouthWesternEurope	UKI			

ManHEP-LCG2	SHEFFIELD-LCG2	cpDIASie	csUCCie	giWTIie	scgNUIGie
RAL-LCG2	ScotGRID-Edinburgh	BHAM-LCG2	BITLab-LCG	BRISTOL-PP-LCG	CAVENDISH-LCG2
Durham	IC-LCG2	Lancs-LCG2	LivHEP-LCG2	OXFORD-01-LCG2	RALPP-LCG
RHUL-LCG2	UCL-CCC	UCL-HEP	csQUBuk	csTCDie	giAITie
gDCUie	giDTIie	giTCie	giTTAie	giTTRie	giNUImie
gRCSie	giULie	mpUCDie	obsARMuk	scotgrid-gla	





General Status

VO dteam									
Test date	St.	js	ver	wn	ca	rgna	bi	cs	
2005-11-09 19:05:57	OK	○	2 6 0	I	○	○	○	○	○
2005-11-09 16:05:50	OK	○	2 6 0	I	○	○	○	○	○
2005-11-09 13:36:36	OK	○	2 6 0	I	○	○	○	○	○
2005-11-09 13:05:31	OK	○	2 6 0	I	○	○	○	○	○
2005-11-09 10:06:18	OK	○	2 6 0	I	○	○	○	○	○
2005-11-09 07:06:03	OK	○	2 6 0	I	○	○	○	○	○
2005-11-09 04:05:22	OK	○	2 6 0	I	○	○	○	○	○
2005-11-09 01:05:11	OK	○	2 6 0	I	○	○	○	○	○
VO dteam									
Test date	St.	js	ver	wn	ca	rgna	bi	cs	sh
2005-11-15 07:05:29	CT	○	2 6 0	I	○	○	○	○	X
2005-11-15 04:05:37	OK	○	2 6 0	I	○	○	○	○	○
2005-11-15 01:05:51	CT	○	2 6 0	I	○	○	○	○	X
2005-11-14 19:05:53	OK	○	2 6 0	I	○	○	○	○	○
2005-11-14 16:05:53	OK	○	2 6 0	I	○	○	○	○	○
2005-11-14 13:05:59	OK	○	2 6 0	I	○	○	○	○	○
2005-11-14 10:05:26	OK	○	2 6 0	I	○	○	○	○	○
2005-11-14 07:05:40	JL	X	??	??	??	??	??	??	??
2005-11-14 04:05:40	JL	X	??	??	??	??	??	??	??
2005-11-14 01:05:35	JL	X	??	??	??	??	??	??	??
2005-11-13 22:05:36	JL	X	??	??	??	??	??	??	??
2005-11-13 19:05:25	JL	X	??	??	??	??	??	??	??
2005-11-13 16:05:46	JL	X	??	??	??	??	??	??	??
2005-11-13 13:05:26	JL	X	??	??	??	??	??	??	??

- Introduction of the Site Functional Tests has had clear benefits, as has the human monitoring and intervention of the Core Infrastructure Centre on Duty.
- Not all sites respond quickly to tickets but they are improving.
- The UKI management structure (by Tier) has helped coordinate the resolution of problems and get responses.
- There was confusion over OS upgrades (RH73 to SL3), which held sites back. Route from SL3 to SL4 or 5 should be better planned.
- No UK site has yet adopted Quattor - too heavy for a single T2? Much success with Kickstart+YAIM installs.





Things we found useful

- New wiki has helped gather information for sites and improved documentation within ROC (Plug: MediaWiki + GridSite is great!).
- Mailing list for UKI (TB-SUPPORT), as well as LCG lists, e.g., LCG-ROLLOUT. (But traffic on LCG-ROLLOUT is high and important things can be missed.)
- Storage support group at RAL and Edinburgh have really helped sort out SRM issues by developing a *storage support community* of middleware experts and site administrators.
- Working as a federation of sites in a T2 helps share skills in a geographical area - e.g., DPM deployment in SouthGrid was accelerated enormously.





Sites from 3 Tier-2s participated, supporting one experiment each

What worked?

- Helped push forward support of SRMs and develop support staff.
- Gained experience with sustained transfers.
- Established a lightpath between Tier-1 and Tier-2 that has allowed additional tests.
- Every Tier2 learned a lot about SRM “local issues” - a base install is only the beginning (it’s prototype to production again!).

Edinburgh, LHCb



RAL (T1)

Lancaster, ATLAS

Imperial, CMS





Concerns!

- Needed more time to continue tests.
- Would like to have pushed rates - Tier1 priorities changed in service phase.
- Components installed have needed several upgrades - was there insufficient testing before release?
- Experiments not fully engaged at sites or clear about what they want.

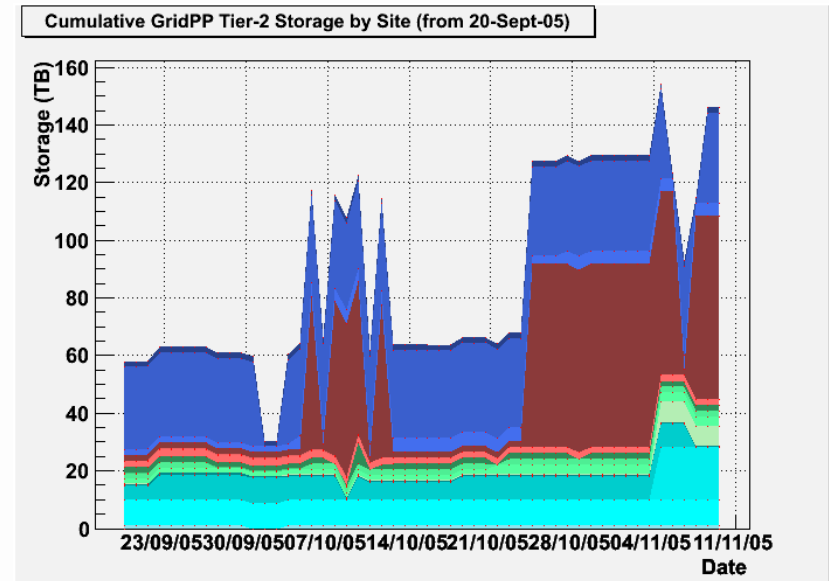
SC3 Throughputs

Site	Rate
RAL (T1)	650Mb/s
Edinburgh	250Mb/s
Lancaster	27Mb/s
Imperial	480Mb/s



Preparations

- **Storage** - strategy document produced to help sites decide route. GridPP storage group and Tier2 support very helpful. 65% sites now SRM enabled - will be 100% by SC4.
- **Networking** - slight concern with some sites (they are independent and need to be aware of regional plans).
- **Data Management** - some sites experimenting with FTS already. Local catalog deployment at all sites.
- **Mini-SC** - planning on pre-SC4 transfer testing at every site to get a heads-up on problems.



- **VO Boxes** - deployed at RAL, Lancaster and Imperial. Operationally these have been manageable, but significant concerns exist.



- **VO Boxes, VO Boxes, VO Boxes!**
 - Few will deploy for ALICE. Security concerns may prevent others.
 - Concerns over scale (EGEE has 70 VOs) and effort required.
 - Commitment as a temporary stop gap (to be addressed long term in the middleware) will help convince sites to deploy them.
- **Install Tools:** The move towards XML/python based will be hard for sites. We will get more experience by getting more sites involved in the Pre-Production Service but:
 - Python not as well understood as bash by most sysadmins.
 - gLite install modules do not even seem as modular as YAIM. How to locally customise? Can we debug them? *N.B. Every site is unique!*
 - Long term commitment to new install method essential
 - Quality mitigates all concerns: it must be high to maintain current resource availability
- **Expectations of Tier2s:** Tier 2 uptime of 95% is ambitious, and currently seems ill defined.



- **Release Schedules:** Sites need to have effort available when new features arrive, e.g., SRMv2. Upgrades work, but continuous releases mean each site can be experimenting with untested combinations of middleware components?
- **SFTs and GGUS Tickets:** Though they have been useful many things not monitored (e.g., SRMs) and too many false positives. Tickets raised often give no additional help to sites.
- **Interaction with Experiments:** Tier 2s and experiments have little experience of interacting yet - we need a model that will scale and well defined contacts, e.g., experiments haven't moved files off SEs when asked.
- **Box Tax:** New boxes for middleware are accepted (sharing for small sites is useful, e.g., LFC + MON box). VO Boxes, however, don't scale!
- **Security:** Bugs raised in security area of Savannah don't seem to make progress. This is very worrying for sites.



Conclusions

- Perceptions of stability improving. But will this continue?
- Concern about change in release strategy (though reasons for this understood). Planning effort will be more difficult.
- Deployment of new services is underway: SRM SEs well advanced, Catalogs in planning stage. VO Boxes remain contentious.
- Changes in install tools *are expected to cause problems* for sites.
- Supported OS base will have to widen - experience of not using SL3 was very painful and most sites gave up. However, this will become more necessary with sites offering large scale resources to LCG on clusters which will not run SL3.
- Sites looking forward to the challenge and gaining in confidence.

Tier 2s are a vital part of the LCG model - improvements in stability and easy of use will be repaid many times over!