



GridPP
UK Computing for Particle Physics



Lessons for the naïve Grid user

Steve Lloyd, Tony Doyle

[Origin: 1645-55; < F, fem. of *naïf*, OF *naif* natural, instinctive < L *nātīvus* native]



GridPP
UK Computing for Particle Physics

ATLAS User Tests



Aim: Attempting to emulate ATLAS 'User experience' of the (UK) Grid

Method: Send 1 of each job per hour to each UK site:

- 1 "Athena Hello World" - checks ATLAS code installed and available to WNs
- 2 "New Athena Package" - checks CMT, gmake etc work
- 3 "User analysis" - attempts to copy AOD data (Z->ee) from local SE to WN and calculate Z mass
- [4 Coming soon - as 3 but Posix IO from SE to WN]

Checks the correct answer from job output - 10 (events) for 1 and 2 and 80-100 GeV for 3



GridPP
UK Computing for Particle Physics

ATLAS User Tests



[Release Matrix](#) [Replicas](#) [More information](#) [History \(csv\)](#) [Blog](#)

Current status is: Running

Most recent job submitted at Tue Mar 6 2007 10:00

<http://hepwww.ph.qmul.ac.uk/~lloyd/atlas/atest.php>

ATLAS UK Grid Status at Tue Mar 6 2007 10:10

Last 10 jobs -> most recent to right

Institute	Release	Replica	Hello World										New Package										User Analysis										Last 24 hrs	Status
UKI-LT2-Brunel 40	12.0.6	DPM	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	0%			
UKI-LT2-Brunel 44	12.0.6	DPM	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	0%			
UKI-LT2-IC-HEP	12.0.6	dCache	S	S	S	S	S	S	S	S	C	C	X	F	F	F	A	F	F	F	F	C	S	S	S	S	S	S	S	S	C	57%		
UKI-LT2-IC-LeSC	12.0.6	dCache	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	98%		
UKI-LT2-QMUL	12.0.5	DPM	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	78%		
UKI-LT2-RHUL	12.0.6	DPM	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	0%		
UKI-LT2-UCL-CENTRAL	Unknown	DPM	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	0%		
UKI-LT2-UCL-HEP	12.0.6	DPM	S	S	S	S	S	S	S	C	C	C	X	S	S	S	S	C	S	C	C	C	S	S	C	S	S	S	S	C	C	93%		
UKI-NORTHGRID-LANCS-HEP	12.0.6	dCache	A	S	S	S	S	S	A	A	S	S	A	S	S	S	A	S	A	S	A	S	S	S	S	S	S	S	S	S	S	81%		
UKI-NORTHGRID-LIV-HEP	12.0.4	dCache	A	S	F	S	S	S	A	S	F	F	A	F	S	S	S	A	C	F	F	C	F	F	F	F	F	F	A	F	A	F	28%	
UKI-NORTHGRID-MAN-HEP 02	12.0.6	dCache	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	F	S	S	S	S	S	S	S	S	S	S	S	S	S	96%		
UKI-NORTHGRID-SHEF-HEP	12.0.6	No	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S										95%		
UKI-SCOTGRID-DURHAM	12.0.6	DPM	X	X	C	C	C	C	C	C	C	C	X	X	C	C	C	C	C	C	C	C	X	X	X	C	C	C	C	C	C	0%		
ScotGRID-Edinburgh	12.0.6	dCache	S	S	S	S	S	S	S	S	S	C	S	S	S	S	S	S	S	S	S	C	S	S	S	S	S	S	S	S	S	100%		
UKI-SCOTGRID-GLASGOW	12.0.6	DPM	S	S	S	S	S	S	S	S	S	C	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	100%		
UKI-SOUTHGRID-BHAM-HEP	12.0.6	DPM	S	S	S	S	S	S	S	S	C	C	F	F	S	S	S	F	S	S	S	S	F	S	S	S	S	S	S	S	S	69%		
UKI-SOUTHGRID-BRIS-HEP	12.0.6	DPM	S	S	S	S	S	S	S	S	S	C	S	S	S	S	S	S	S	S	S	C	X	S	S	S	S	S	S	S	S	47%		
UKI-SOUTHGRID-CAM-HEP	12.0.5	DPM	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	C	S	S	S	S	S	S	S	S	S	96%		
UKI-SOUTHGRID-OX-HEP	12.0.6	DPM	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	100%		
UKI-SOUTHGRID-RALPP	12.0.6	dCache	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	C	100%		
RAL-LCG2 (Tier-1)	12.0.6	dCache	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	100%		
Total																															64%			




GridPP
UK Computing for Particle Physics

ATLAS User Tests



- Many problems identified and fixed at individual sites (GridPP DTeam)
- Other 'Generic' system failures that need to be addressed before fit for widespread use by inexperienced users
- Production teams mostly 'work round' these
- **Users can't/won't**





GridPP
UK Computing for Particle Physics

Resource Broker



- Use RBs at RAL (2) and Imperial
- Break about once a week and all jobs lost or in limbo
 - Never clear to user why
- Switch to a different RB
 - Users don't know how to do this
- Barely usable for bulk submission - too much latency
- Can barely submit and query ~20 jobs in 20 mins before next submission
 - Users will want to do more than this
- Cancelling jobs doesn't work properly - often fails and repeated attempts cause RB to fall over
 - Users will not cancel jobs
- (We know EDG RB is deprecated but gLite RB isn't currently deployed)



GridPP
UK Computing for Particle Physics

Information System



- lcg-info is used to find out what version of ATLAS software is available before submitting a job to a site but it is too unreliable and previous answer needs to be kept track of
- Most unreliable and mysterious! (e.g. ldap query typically gives quick, reliable answer but lcg-info doesn't)
- The lcg-info command is very slow (querying *.ac.uk or xxx.ac.uk) and often fails completely
- Different bdiis seem to give different results and it is not clear to users which one to use (if the default fails)
- Many problems with UK SE's have made the creation of replicas painful - it is not helped by frequent bdii timeouts
- The FDR freedom of choice tool causes some problems because sites fail SAM tests because the job queues are full



User Interface

- Users need local UIs (where their files are)
- These can be set up by local system managers but generally these are not Grid experts
- The local UI setup controls what RB, BDII, LFC etc all the users using that UI get and these appear to be pretty random
 - There needs to be clear guidance on which of these to use and how to change them if things go wrong

Proxy Certificates

- These cause a lot of grief as the default 12 hours is not long enough
- If the certificate expires it's not always clear from the error messages when running jobs fail
- They can be created with longer lifetimes but this starts to violate security policies
 - Users will violate these policies
- Maybe MyProxy solves this but do users know?



- GGUS is used to report some of these problems but it is not very satisfactory
- The initial response is usually quite quick saying it has been passed to X but then the response after that is very patchy
- Usually there is some sort of acknowledgement but rarely a solution and often the ticket is never closed even if the problem was transitory and now irrelevant
- There are two particular cases which GGUS does not handle well:
 - a) Something breaks and probably just needs to be rebooted: the system is just too slow and it's better to email someone (if you know whom)
 - b) Something breaks and is rebooted/repared etc but the underlying cause is a bug in the middleware: this doesn't seem to be fed back to the developers
- There are also of course some known problems that take ages to be fixed (e.g. the globus port range bug, rfio libraries, ...)
- More generally, the GGUS system is working at the ~tens of tickets/week level but may not scale as new users start using the system



Conclusion



- The Grid is a great success for Monte Carlo production
- However it is not in a fit state for a basic user analysis
- The tools are not suitable for bulk operations by normal users
 - Current users therefore set up ad-hoc scripts that can be mis-configured
- 'System failures' are too frequent
(largely independent of the VO, probably location-independent)
- The User experience is poor
 - Improved overall system stability is needed
 - Template UI configuration (being worked on) will help
 - Wider adoption of VO-specific user interfaces may help
 - Users need more (directed) guidance
- There is not long to go
 - Usability task force required?