



2nd DRAFT

(14/9 changes in Alice part by Federico Carminati)

Minutes of the meeting

CERN, 8 September 2004

Agenda: <http://agenda.cern.ch/fullAgenda.php?ida=a042133>

Minutes: Jeremy Coles and David Kelsey

Attendees: See appendix

1. Introduction and Short Communications

It was noted that (due to a clash with HEPiX) the next meeting will take place on 13th October starting at 09:30. A GDB dinner on the night of 12th October is planned.

2. Quattor Installation of Grid Software (Charles Loomis)

Please refer to slides on agenda page.

Summary of talk:

- *Goal of Quattor is to provide one set of tools for: installation, configuration and maintenance.*
- *Development team consists of 3 people at CERN but with no mandate for grid-specific components.*
- *Community effort came out of meeting on 26 March 04 - but with limited effort.*
- *LAL priorities are to get a full LCG-2 installation via Quattor but because of lack of effort and mounting pressure they may have to revert to LCFG.*
- *Quattor is actually a set of components not a single product:*
 - *Configuration database*
 - *Software repository*
 - *Integration cache management*
 - *Configuration of machines*
 - *Software package installation*
- *Since March, things have matured significantly. The Pan library is complete and critical configuration components exist. Easy to use with RH-like OS's (comps.xml description). Proven to work with RH7.3, SL3.01 and CEL3 (and others) originally based on LCG 2.0 and now 2.2.0.*



- *Testing status: UI, LCG-BDII and PX are all OK. The CE, WN, SE and RB have some configuration problems that still need further work, but very close. mkggridmap problem was solved yesterday.*
- *Quattor has no template facility (unlike LCFG) - plans to implement this soon.*
- *Cal's experience is that the control/checking is better than in past tools.*
- *Main issues are the lack of community effort and the longer term maintenance of configuration (LAL will not be able to support wider use).*

Conclusions

- *Now offering end-to-end installation & configuration with better control/checking than LCFG. Many things improved in port.*
- *Problem with resource in development and long term maintenance.*

Questions:

Do you expect there will need to be different modules for different institutes?

RPM configuration done but new features will require maintenance of configuration.

Different sites will require different configuration but the RPM lists are general.

Is the move to Quattor linked with the move to SL?

No. RH7.3 can be used (LAL machines too new for 7.3). Quattor supports changes of OS easier than LCFG.

How can we boost the community effort?

Many institutes willing to participate – many playing with Quattor already. Having LCG 2_2_0 available should boost things but it remains an unofficial tool - making it officially supported would help.

Can Cal help with installation on the cert. testbed and how long does it take to install?

Cal does not have time to do the install but can answer questions. Configuration should take less than a day Installation may take up to a few days. It is quite similar to LCFG.

Are all components in the central CVS server?

No. The current version is in the URL provided in the talk. Once all components have been tested CVS will be updated.

Is it possible to keep a local repository of RPMs?

This is possible and indeed LAL uses a local Apache server.

Micro requested an action on Ian Bird for clarification. Ian responded with his position:

- Main focus of the deployment team is on manual (scripted) installation
- His team will provide pointers to Quattor components
- Being careful not to focus too much on one tool (unlike the LCFG case)
- LCFGng has gone unless someone else is prepared to port it to SL
- Testbeds at CERN will use Quattor



Tony Cass continued:

- CERN intends to use Quattor and would like to see it adopted elsewhere. So there may be a way to divide up work within this community. CERN could answer bug reports but would work to CERN timescales.
- We could run training sessions to bring more people into the community

Will anyone certify Quattor installation against manual install?

Ian's response:

- CERN team does not have the manpower to do this to the same timescales as a manual release
- CERN do not want to be accused of imposing a product
- LCFG main issue was with imposition on OS
- Some sites have existing management tools (US sites, Lyon & small sites had problem of overhead)
- Components will be provided but not at the same time as manual installation
- Manual install is not entirely manual and is based on nodes. There is flexibility to make the split service rather than node based.
- Main work is not producing objects – more to do with solving problems if object problems at sites

Should there not be a commitment to the community? The risk is that many sites will wait for the Quattor components.

- A community of users should be established with joint responsibilities and explicit commitments made to support various components
- Modified components breaking will be a major issue for which CERN staff are not resourced to provide support
- With Quattor – each component change requires the server software to be reinstalled

Action: Kors to coordinate discussion outside of meeting

3. Computing MoUs (David Jacobs)

Please refer to the presentation slides for full coverage of topics.

Summary of presentation:

- *Taskforce setup to ensure coverage of Tier-1 centres*
- *LCG MoU down to last stage 4 main points to be addressed*
 - *Interface between “managed” part of LCG and smaller centres*
 - *Description of relationships between Tier centres*
 - *The resource sharing mechanism*
 - *IPR*



- *Next draft mid-September*
- *Next steps*
 - *Finalise body text by end of October*
 - *Experiments to provide requirements against time by end 2004-09-08*
 - *Jan/Feb LHCC review of requirements*
 - *Proposal document ready for April 2005 RRB*

Questions:

If funding agencies commit to experiments why sign a central MoU?

There is a need to coordinate between the various funding agencies. Also note the important role of the Scrutiny group in checking the validity of the experiment requirements

Is there not a high risk for agencies committing to experiments?

Agencies are approving use of resources. Some countries like US will only commit to specific experiments.

What is the model? The LCG MoU describes a collaboration? This needs a project plan with milestones and deliverables. Why then is the CRRB approving this?

The LCG project is already approved *but not phase 2*. The CRRB asked for the MoU. They should approve it. The MoUs come before the TDR. LCG and Computing TDR's will come in 2005 to describe the project (approved by LHCC). Could take the position that there needs to be an LCG2 position document. These are fundamental questions – should be taken off-line!

It was decided that the MoU should be presented again at the November GDB by which time the document (without annexes) will be publicly available.

David Stickland: The correct route for feedback is via representatives on the MoU Task Force.

Kors: then important that TF members consult widely.

4. The Atlas Data Challenge (Gilbert Poulard)

Please refer to the presentation slides for full coverage of topics.

Summary of presentation:

Data Challenge running now and not all problems have been analysed.

Divided DC2 into 3 parts:

- *Simulated data (July-Sept 04)*
 - *Event generation*
 - *Detector simulation*
 - *Pile-up and digitisation*



- Data transfer (to Tier-0)
- Event mixing
- Test Tier-0 (Oct 04)
- Test distributed analysis (Oct-Dec 04)

Early difficulties in LCG (items in bold still not fully resolved):

- Mis-configuration of sites: **information system** (wrong or missing information)
- Job submission and **Resource Broker**: Job ranking
- **Data management (copy & register)**: Stage in/out problems

Most problems solved with direct contact – 1 regular meeting per week. Have not used GOC.

Overall failure 40-50%. 2000 jobs.

Athena problems – job submitted to long queues but running on wrong queue so job fails LCG-RB – can not plan. Server problem.

LCG-Proxy – main problem on 1 Spanish site – distorts outlook

DMS-Output – data management system wrong information into database.

Main problem is with data management

NorduGrid failure rate of order 20%.

These stats do not match findings from certification area – but problems suggest most faults due to m/w issues!?

Summary:

- Delays in all activities have affected the schedule of DC2
- DC2 – about 80% of Geant4 simulation completed using 3 Grids which have proven to be usable for real production
- Phase I progressing slower than expected

Transfer of 35 Tbof data to CERN due to start around 15th October.

Plans 2005 – move to continuous production mode 15-20 million events. Simulated data by end February.

Questions:

How much slower are things going compared to the expectation?

Gilbert: at least factor 2 (would like at least 2000 jobs per day)

Timescale for ATLAS move to Scientific Linux?



Gilbert: not before the end of the year. is there another constraint?

5. ALICE Data Challenge (Federico Carminati)

Please refer to the presentation slides for full coverage of topics.

Summary of presentation:

PDC04:

- *Test and validate the ALICE offline computing model*
- *Three phases*
 - *Distributed production of underlying Pb+Pb events (3.8 million files in AliEn file catalogue. CPU 285 MSI-2K hrs)*
 - *Mixing (110 channels, 1M jobs, 10TB data)*
 - *Distributed analysis*
- *Jobs run using AliEn. In phase 3 – gLite+ARDA*
- *CERN LCG & Torino LCG – 18% of completed jobs*

Issues:

- *CASTOR files limit – a work-around has been found*
- *Local configuration problems of LCG sites. Frequent black holes: large sites monitor jobs on batch services but not smaller sites. Local disk space on WNs (not in JDL). Problems return (e.g. nfs mounts)*
- *Quality of information in II (e.g. time limit for queues in minutes and in seconds)*
- *WLMS does not ensure even job distribution over centres*
- *Lack of support for bulk operations (1M jobs – WMS response time problem)*
- *Lack of appropriate monitoring and reporting tools*

Phase II not yet finished.

Issues:

- *Same problems as above – followed by LCG support*
- *Transition from edg-* to lcg-* data management commands*

Phase III

- *Plan to use LCG for interactive analysis – risk*

Conclusions:

- *LCG providing increasingly coherent infrastructure with large number of resources*
- *Relations with LCG good*
- *Middleware improving but within limits of current design and architecture*
- *Willing and ready to move to gLite*



Questions:

EGEE unable to give gLite modules before end Sept. How can ALICE use it in Oct?
Only hope to be able to deploy many gLite modules. Willing to use dCache storage element

Did you make any differences between Tier-1 and Tier-2s?

No obvious difference – in phase 2 local storage occurs. Once storage full it is up to the site to move it to a Tier1. How would Tier-2 manager know files to be moved?

Alternatively, stop using site once their storage fills up. Should be transparent to change if register replicate then delete

6. HP Participation in EGEE (Michel Bernard)

Please refer to the presentation slides for full coverage of topics.

Summary of presentation:

- *Several sites either in use or testing*
- *Efficiency of order 65-80% running LHCb jobs (but will support other experiments)*
- *LHCb jobs – grid jobs run but there are firewall issues to be resolved*

Grid Benchmarks:

- *Data challenge 7GB/s for memory to memory*
- *3.3Gb/s for disk to disk*

Other results

- *All LCG middleware ported and tested on Itanium*
- *Setting up testbed for Utility Computing resource allocation model*

Questions:

3.3 Gbps – what sort of network?

Dedicated between Caltech and CERN.

7. CMS Data Challenges and Operations (David Stickland)

Please refer to the presentation slides for full coverage of topics.

Summary of presentation:

- *Predominantly data movement challenge*
- *Small amount of Tier-2 activity*



DC04 summary:

- *Achieved 25Hz reconstruction rate but only for short period*
- *RLS poor performance issues largely traced and corrected*
- *CMS TMDB over SRB worked well*
- *Solved many internal CMS problems*
- *Processed about 30M events*
- *Some problems with mass storage at Tier-1s*

New Boards:

CMS-CCC (Computing Coordination Committee)

- *Senior physicist + senior computing person from major contributors*
- *Initial task to quantify and overlook resources*

APROM – Cross-Project Analysis Coordination to coordinate CMS computing and software activities for analysis from the perspective of end users

Operations:

- *Full time production for one year – some on grid some on dedicated resources*
- *Working on areas like: enabling tasks like digitization, data movement, publication and smoothing production operations (continuous delivery)*
- *New CCS tasks of Data Management and Workload Management*
- *Continuous user access to data on the grid is a big issue*

Physics TDR scheduled for December 2005

In light of DC04 experience, DC05 is cancelled as a formal computing exercise – not possible to serve current physics requirements with data challenge environment. Expect specialised component challenges:

- *Will bring up a full time operation including data access*
- *750 CPUs continuously (~2.4GHz+)*
- *Need to reach 10M events/month level soon (Autumn)*

Questions:

What were the communication channels with Tier-1s?

Daily meetings with as many sites as possible but also used telephone and email list channels.

8. LHCb Data Challenge (Ricardo Graciani)

Please refer to the presentation slides for full coverage of topics.

Summary of presentation:



Aim: Collect information for TDR physics goals by a significant increase in number of signal and background events

Phases:

- *Production – MC simulation done: 186M events. Started with just Dirac. LCG contribution ramped up since May*
- *Stripping – event pre-selection: to start soon*
- *Analysis: in preparation*

LCG usage strategy: Use all available resources

- *Sites tested – successful sites added to production mask*
- *Continuous job submission (4RBs) and queue monitoring*
- *Worked closely with LCG-GDA team and site managers*

Are jobs checked for progress?

Clear out jobs in queues. All jobs submitted previous day that were not running were resubmitted – mainly to avoid job failures due to proxy issues. During August this kept about 3500 concurrent jobs running for LHCb.

When a job is sent to a node it only contains a small script in its sandbox. The script decides which job can run best in the environment of the worker node. Only then the necessary software is copied and installed.

Markus: interactive work could be done with running batch jobs in this way

David S: But is this appropriate for proper sharing?

Levinson: separates out resource allocation from specification of the work. Looks attractive. Can you solve this in a general way and move away from push/pull question?

IanB: gLite CE and RB address this issue.

May – 11% on LCG rest on Dirac

August – 73% on LCG

So overall – 50% on LCG

LCG efficiency – 61% (from 1st June)

211K submitted: 113K completed: 34K failed

LCG issues:

- *fraction of aborted jobs due to site mis-configuration before jobs even start to run (20%)*
- *fraction aborted after run start (18.5%)*
- *fraction of “successful” jobs did not produce Output Sandbox*
- *Job scheduling problems due to “wrong” info from CEs*
- *Resource ranking – several approaches tried*
- *Fault debugging – too many ids and logs to search*
- *Hardware, software, hacker.... problems in dedicated UI and RB machines*
- *Not tested Proxy Server, RM ...*



What were the hacker problems?

Break in to machine with both UI and RB installed.

DPK: why no Incident Response?

Ian: two separate incidents - first one not clear, second one knew the root password. Need to revisit this with new Ops Security Group.

DPK: compromised RB should at the very least result in a heads-up message.

Lessons DIRAC:

- *Improve server availability*
- *Minise need for human intervention*
- *Improve error handling and reporting*
- *Test hardware capacity of servers*

Lessons LCG:

- *Limit strong dependence on LSRM*
- *Improve OutputSandBox upload & retrieval mechanism*
- *Improve reliability of CE status collection methods (timestamps?)*
- *Add “intelligence” on CE/RB to detect and avoid large number of aborted jobs on start up*
- *Need to collect LCG-log info and tool to navigate them*
- *Need a way to limit the CPU (and wall-click time)*
- *Need to be able to allocate Local Disk Space for running job*
- *How-to manuals – clear instruction to site managers on the procedure to shutdown a site (for maintenance and or upgrade)*
- *RB performance needs to be carefully monitored*
- *Changes required to definition of “running” state*
- *Job Assignment to a resource must be delayed until previously assigned jobs are submitted to the site (or fail)*
- *For production-like tasks improve responsiveness (e.g re-use of previously uploaded Input Sandbox)*

Summary:

- *LHCb DC04 Phase 1 over*
- *Production target achieved (186M events, 50% on LCG)*
- *LHCb strategy proved correct*
- *Room for DIRAC and LCG improvements*
- *Thanks to LCG team for support*

9. Operational Issues from Data Challenges (Ian Bird)

Please refer to the presentation slides for full coverage of topics.

Summary of presentation:



*Complexity in middleware is starting to be exposed – resilience to failure.
LCG is now 73 sites: 26 sites at 2_2_0 and 33 sites at 2_1_1 29 pass all tests.
Some sites use old version of VDT. Unable to upgrade due to constraints at sites or productions.*

New version of worker node software can be installed on fly as user – would need second CE to communicate this. For experiments the WNs are most important to upgrade.

Outstanding middleware issues recorded in two documents (many to be addressed by gLite and future developments):

- <http://edms.cern.ch/file/495809/0.4/Broker-Requirements.pdf>
- http://fca.home.cern.ch/fca/DCFeedBack_gag.doc

Operational issues:

- *Slow response from sites*
- *Lack of staff available to fix problems*
- *Mis-configurations*
- *Lack of configuration management – who should provide scripts to monitor bad nodes?*
- *Lack of fabric management*
- *Lack of understanding*
- *Firewall issues*
- *PBS problems*
- *Forget to read documentation*

Addressing occasions:

- *Weekly GDA meetings (operations)*
- *Grid ops 0.5 day at HEPiX in October*
- *Operations and Fabric workshop (CERN 1-3 November) – to agree on operations model for the next year.*

Operations efforts:

- *LCG GOC (RAL) -> EGEE CICs and ROCs + Taipei*
- *Regional Operations Centres (ROC)*
- *Core Infrastructure Centres (CIC)*

Status of Scientific Linux port:

- *Worker node port is done in next public release*
- *Full SL3 by end of September*
- *Need to be able to support combinations (service nodes R73, worker node SL3)*
- *LCFGng is not ported to SL3*

Summary:



- *Biggest outstanding issues related to providing and maintaining stable operations*
- *Management buy-in to provide sufficient and appropriate effort at each site required (inc. future middleware)*

10. AOB

John Gordon: The successful LCG sysadmin course run in Oxford this summer will be turned into a sysadmin handbook.

John Gordon: Introduced Jeremy Coles as his new deputy for UK

Kors: Thanked Mirco for smooth transition in chairmanship

Next meeting in Amsterdam, 13 Oct. Web site for registration will be open soon. Theme for the meeting: input from centres (large and small) on experiences with the Data Challenges

ACTION on sites – offer talks to Kors

Meeting closed at 16:55



LCG Grid Deployment Board Meeting



Not in usual format and may not be complete – check paper list of attendees

Kors Bos	NIKHEF	NL	(Chair for afternoon)
Mirco Mazzucato		IT	(Chair for morning)
Cal Loomis	LAL		
Federico Carminati	CERN	ALICE	
David Jacobs	CERN		
Fabio Hernandez	CC-IN2P3	FR	
Jukka Klem	HIP	FI	
Michel Bernard	HP		
Ian Bird	CERN		
Markus Schulz	CERN		
Bernd Panzer	CERN		
David Foster	CERN		
John Gordon	RAL	UK	
David Kelsey	RAL	UK	
Jeremy Coles	RAL	UK	
Davide Salomoni	NIKHEF	NL	
Laura Perini	Milano	ATLAS	
Ricardo Graciani	Barcelona	LHCb	
Denis Linglin	CC-IN2P3	FR	
Gilbert Poulard	CERN	ATLAS	
Tony Cass	CERN		
Mariusz Wittek	PL		
Di Qing	ASCC	Taiwan	
F Orellana	Univ Geneve	Switzerland	
Marco Serra	INFN/CERN		
Zhongliang Ren	ASCC	Taiwan	
Lorne Levinson	Weizmann	IL	
Zdenek Sekera	CERN		
Dietmar Kuhn	Austria		
Flavia Donno	CERN		
Alberto Masoni	INFN	ALICE	
David Stickland	CERN	CMS	
Nick Brook	Bristol/CERN	LHCb	
Holger Marten	FZK	Germany	
Milos Lokajicek	Prague	CZ	

Via VRVS: Andreu Pachecho, David Groep, Manuel Sanchez, Christoph Grab, NOC
ASCC,
Vicky White, Bruce Gibbard, Jordi Garcia, Philippe Charpentier, Joel Closier, Mike
Vetterli, Manuel Delfino