

Summary of Friday's talks/discussion

Jeremy & ...

Questions/comments in BLUE

Middleware - Oliver

- gLite release history and rpm lists are online along with a baseline wiki
- Baseline: FTS/SL4 (soon); CREAM (available – parallel deployment with LCG-CE). WN/SL5 (under test, deploy in parallel with SL4). Multiple versions of middleware on WN coming soon.
- Glxec/SCAS under certification. Glue2 awaiting OGF validation. Expect changes with heterogeneous resource publishing. WMS with ICE on Feb09 timescale.
- SL5: target 32 & 64-bit WN. UI will follow next.
- What policy is in place to control the middleware version deployment. (How) will SAM tests track versions installed

Storage - Flavia

- Detailed status reports and recommended versions list available. Talk looked at CASTOR (core and SRM) next 2.1.8. dCache recommend 1.9.0 (soon to provide catalogue synch; better info. Providers and pnfs->Chimera migration tool). DPM 1.6.11-3 with 1.7.0 in certification (working on quota system, monitoring, tools). StoRM 1.4 expected in December.
- Distribution plan for non-native file access libs. A: dCache will provide ROOT with compiled dcap/gsidcap.
2) Progress on common library for CASTOR and DPM?

Status of critical services (Andrea)

- Most CERN central services (fully) ready. Some services missing a piquet service and/or documentation.
- ALICE ok. ATLAS ok – but relying on experts. CMS services basically ready (but still debugging). LHCb lack procedures/documentation in many areas
- Out-of-hours Oracle service cover needs to be reviewed. VO boxes were supposed to be moved into class 2.
- The shared area for software at sites may also be considered – jobs crash when shared file system overloaded.
- WLCG maintains a list of incidents and should look at those at the weekend that are also on expt. critical list.

Analysis working group (Markus)

- Purpose is to clarify and confirm roles of Tier-1,2,3 for analysis activities. Understand types and approaches. What services required and how will they be used (access patterns)? Group will document the analysis models/overlaps/requirements (inc. review of existing experience and known problems). Report by end 2008?
- Group includes experiments, sites, data management experts. Note: No US site(s) yet in the membership. Also would like another T1 involved.
- Who will benefit from the report. A: Sites for setup. Tool developers. Some concern (esp. CMS) that this will impose new model. A: It won't, matrix will ID overlaps. Not a model scrutiny. Another "approach" would look at organised vs chaotic user analysis.
- Document could help provide a single information resource for sites – job priority settings, ACLs... Nice if group will gather canned use cases.
- Focus on hardware/resources may usefully also include how users are going to be supported – how this differs between expts. Also useful if group looks at impacts of sites/services being down

Analysis (ROOT/PROOF)- Gerado

- Review of PROOF (parallel coordination of distributed ROOT), performance considerations (T2/3 + desktops), the architecture and how it is setup and optimised, dataset handling.
- Main (likely) bottleneck is I/O. Disks ok with about 2/3 cores. Results of scalability tests – good for CPU-bound tasks and I/O with independent disk controllers. ATLAS I/O tests => 2 cores for 1 disk reasonable. Solid state disks have shown speed advantages.
- PROOF part of ROOT. Several large implementations. Dataset input is a named collection of files – uses different backends. PROOF-Lite is optimised for multi-cores.
- PROOF is an option for interactive end-user analysis at T3s. Good feedback.
- ALICE CAF (slide 19) interacts with which batch system? Discussion was around direction of hardware architecture. Many cores per CPU and disks per node. Also question about running costs vs efficiency.

Goals and metrics for data taking readiness

“challenges” – Ian/Jamie

- Summary of C-RRB/MB discussion. Should there be a CCRC'09? No but do need specific tests/validations (e.g. combined running, reprocessing ...).
- Experiment plans shown yesterday, no time coordination between activities.
- Areas to be considered listed on the agenda (based on list from operations meeting)
- What tests? How will buffers deal with multiple site outages. Real power cuts break discs, simulated outages have less problems!
- Many things happen anyway – need to document the recovery better. Follow up on things that went wrong. Expts. have to expect problems and be able to deal with them in a smooth way.
- Sites still waiting for data flow targets from the expts. CCRC08 rates are still valid. ATLAS plan on continuous load but DDM exercise in 2 weeks will provide peak load. Worry about tape rates – limits in number of available drives. CMS – alignment of smaller activities would be useful. ALICE – live or die by disk based resources.
- Reprocessing and analysis cases of particular concern. Conditions DB in ATLAS (T2 frontier/SQUID) needs testing. Other than concurrent tape rates CMS has own programme of work. When a more complete schedule exists looking at overlaps would be interesting. ATLAS hope all clouds involved with analysis by end of year. LHCb- can plan modest work in parallel – sometimes need to solve problems in production when they occur.
- Allow expts. To continue with programmes of work. When ATLAS read -> common tape recall tests. Try time when can push lots of analysis. Things should be running and then focus on resolving issues in post-mortems. Need to test things that change (e.g. FTS).

Plans and Architectural Options for Physics Data Analysis at CERN - Dirk

- Data management software status for T0: using CASTOR – simplified/improved during 2008 e.g. reduced latency and strong authentication. Working on monitoring, SRM interface and tape efficiency.
- Requirements for analysis – reduced access latency... direct file access strongly demanded, analysis on disk pools only.
- Looking at new (high performance) file system. Two main options under consideration. Mountable file system is strategic. Feedback wanted – approaches lead to different levels of consistency and data management automation.
- No direct exchange between FS and another site – would have to go via CASTOR? Yes. Local and not grid connected. Ties in with Markus's discussion – would be nice if not just tied to CASTOR. Current protocol is XROOTD – could use with others if proves successful. dCache interested in NFS4. Best first understand work patterns.
- Only LHCb used ROOTD because of timeout setting problems. Footprint on I/O server too large.

Direct user ticket “routing” to all sites - Torsten

- Reviewed standard support workflow. Alarm tickets can be sent to T1s by LHC Vos for urgent situations. Team tickets (VO experts) follow same workflow as COD tickets.
- Team tickets being extended to include T2s. Requires GOCDDB interface – estimated for early 09. Several open questions – submission via email, query GOCDDB how often, failover and inclusion of OSG.
- Groups. Better to define VOMS roles rather than just sending DNs? A: Already being considered.
- Direct routing – may happen tickets go to wrong place. This is where user makes mistake in assignment. Calculated risk. Can add authorisation later.
- Site request – see the summary of ticket and list of tickets for a given site. A: Will be possible once we have the site name. Can TPM add the site name? A: Yes. Can then filter.
- Assumption that all site admins will need a certificate to login to GGUS? No email controls. ROC still routes the ticket.

Site Status Board - Pablo

- Developed for CMS to show sites vs VO metrics. Can combine several metrics and keep ok/not ok history. Recently added X509 for adding columns, gridmap interface.
- To do – mark sites “under investigation”. Have value history. So far used mainly for site commissioning but users and other VOs can benefit.
- Do you have way to see scheduled downtime? A: Yes, check status in GOCDB.
- Why did you feel you had to develop something new? A: Many places for metrics, but not a single glance view. Not entirely new development. Sites welcome single view. Other experiments showing some interest.

A new monitoring tool from the site point of view - Elisa

- Tool developed due to feedback – site admins wanted tool to be able to compare expt. view with site monitoring view and have clear targets. Interfaces with different expt. tools.
- Database schema and collector shown to work. Some but not all metrics included.
- Where are the pledges taken from?
- Any colours between blue and red? A: yes yellow!. Biggest problem – tests from experiments are not clear. Sometimes red doesn't mean a site problem but an experiment problem. So, expts. when do we get these clean(er) tests!? A: Difficult also for the VO – some cases clear. Better error messages would help. Button to click to show status of all sites.
- Would be nice to see things as a function of time.

FTS monitoring - Ricardo

- Built using the dashboard framework. Considered various use-cases and kept integration tight. Provides activity and error summaries, FTS settings and reporting. Examples given using filtering. Tested in FTS pilot instance.
- To-do: expand coverage.
- GridPP have a tool that will be of interest.
- Real Time Monitor – in process of providing a feed of finished transfers.