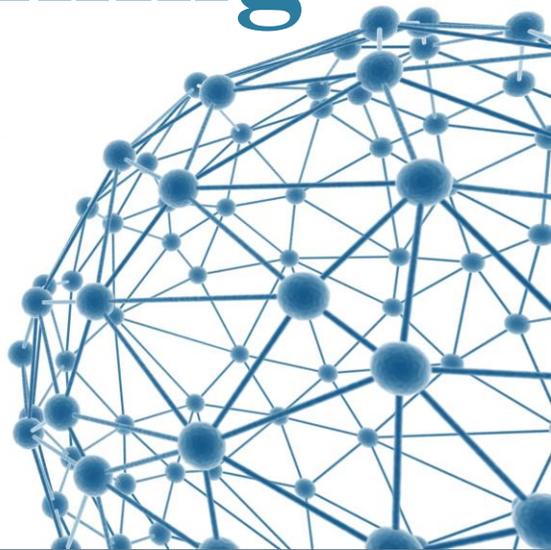




ADC Operations medium-term planning

Alessandro Di Girolamo & Ueda



Activities

- Analysis and MC production continue as in the past months
- Reprocessing campaign
 - one “small” campaign starting soon:
 - Inputs 2012 data RAW (2.2PB), output 200TB
 - Timeline:
 - Code validation in the next 3/4 weeks
 - Bulk reprocessing estimated 5 weeks (using 15% of ATLAS Grid resources)
 - To be finished before March 2014
 - Full reprocessing with new release starting in April 2014

CVMFS

- CVMFS 2.1
 - ATLAS encourage sites to migrate to the WLCG Baseline Version
 - To be able to use multiple repositories (e.g. ATLAS nightly and SFT repo of LCG sw)
 - CVMFS servers at CERN can be upgraded to 2.1 only when all the clients are 2.1
- Check-validity of CVMFS repository on WNs
 - ATLAS observed in some cases WNs with stale CVMFS repo
 - Bugs reported and mostly fixed, but still...: wouldn't be useful to have checks on WNs reported directly to the site admin to remove stale WNs from production?
 - CERN has started doing this...
 - **WLCG general interest?**

FAX

- FAX integration is progressing:
 - FAX is activated for some sites as input file failover
 - More stress tests discussed/planned

- SL6 TaskForce very useful for ATLAS
 - Most of the sites have migrated or planned their migration
 - Procedures/instructions are available for sites which did not migrate yet
 - Thanks to the SL6 Task Force

SL6

- SL6 TaskForce very useful for ATLAS
 - Sorted out obstacles to the migration
 - Procedures/instructions successfully established
 - ✓ most of the sites have migrated or planned their migration
- ! The Task Force has achieved a remarkable success !

Auto-Setup: ATLAS SW shared area

- ATLAS going in the direction of not needing anymore `$ATLAS_LOCAL_AREA`
 - Auto-Setup: define site specific information without the need of local file in the shared area
- Migration to this Auto-Setup started:
 - During migration sites using Auto-Setup can be found:
`/cvmfs/atlas.cern.ch/repo/sw/local/etc/auto-setup-switcher.conf`
- Aim to have Auto-Setup in full production in the next few weeks

WebDAV: Rucio renaming

- ATLAS needs WebDAV for Rucio renaming (or sites should agree with DDMOps on doing the renaming by themselves):
 - dcache is ok: sites with dcache should just activate it!
 - dpm: sites would need to update to 1.8.7
 - Storm: still under performance testing
 - Eos: still some bugs to be fixed
- Still more than 50% of the sites are still not providing WebDAV access for renaming!
 - Issued this week ATLAS CloudSquads tickets to track each site plan.
- Info <https://twiki.cern.ch/twiki/bin/viewauth/Atlas/WebDAVSetupForATLAS>

Rucio commissioning

- The new ATLAS DDM Rucio framework is now being tested
 - simplified full chain: load generation, data transfers, quota checks, deletion
- We have called for sites volunteers. Workload for these sites:
 - discuss the Rucio endpoint setup with DDM Ops
 - provide ~ 10TB of space for the tests
 - investigate SE logs if something (transfers) goes wrong
- If your site is willing to participate contact atlas-dq2-ops@cern.ch

SAM test: WLCG avail/rel recap

- Since October OPS SAM tests are not anymore in the WLCG availability/reliability reports:
 - There is no change respect to the ATLAS SAM tests, which were in the WLCG availability/reliability reports and are still there.
- Brief re-cap:
 - WLCG Availability/Reliability of ATLAS sites is measured with SAM ATLAS_Critical profile
 - CE: DN=ddmadmin, Role=lcgadmin . Tests: JobSubmit+SWSpace+Swtag
 - SRM: DN=ddmadmin, Role=production. Tests: VOPut, VOGet, VODel
 - Other tests on CEs and SRM (also with different credentials) do not contribute to the WLCG availability/reliability
 - ! CE and SRM needs to be defined in GOCDDB/OIM to be tested

Benchmarking Clouds (with HammerCloud)?

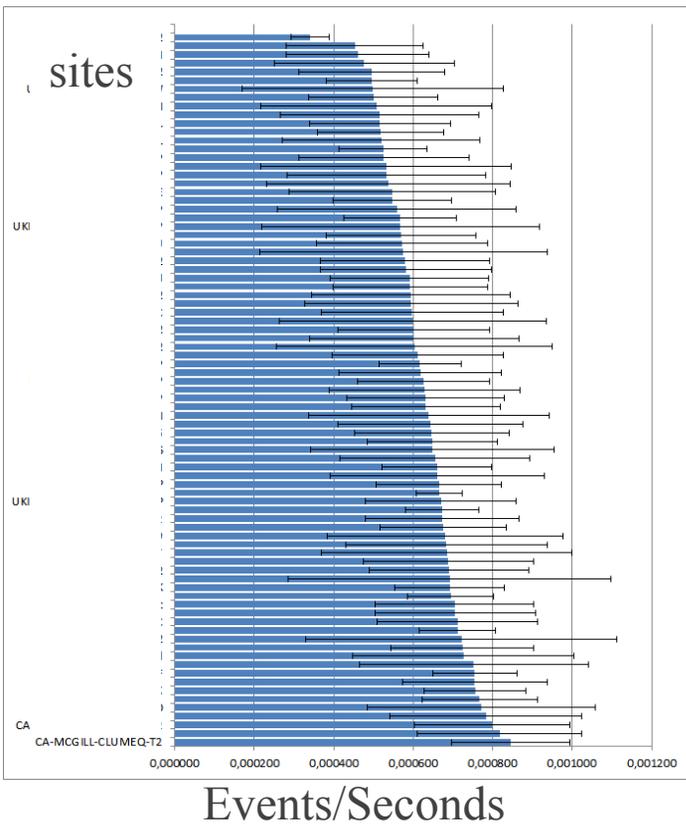
Ramon presentation ATLAS SW week attached to the agenda

- Accounting:
 - Bare metal: HS06 OK!
 - Private clouds: HS06 on HyperVisors Okish!
 - Public clouds: No HS06, no HW info: still a challenge!

- What if....
 - we run “benchmark” jobs everywhere,
 - we extrapolate the HS06 to the other places?

- ! This is NOT meant to evaluate sites!

Benchmarking Clouds (with HammerCloud)?



HC jobs with 2 events/job

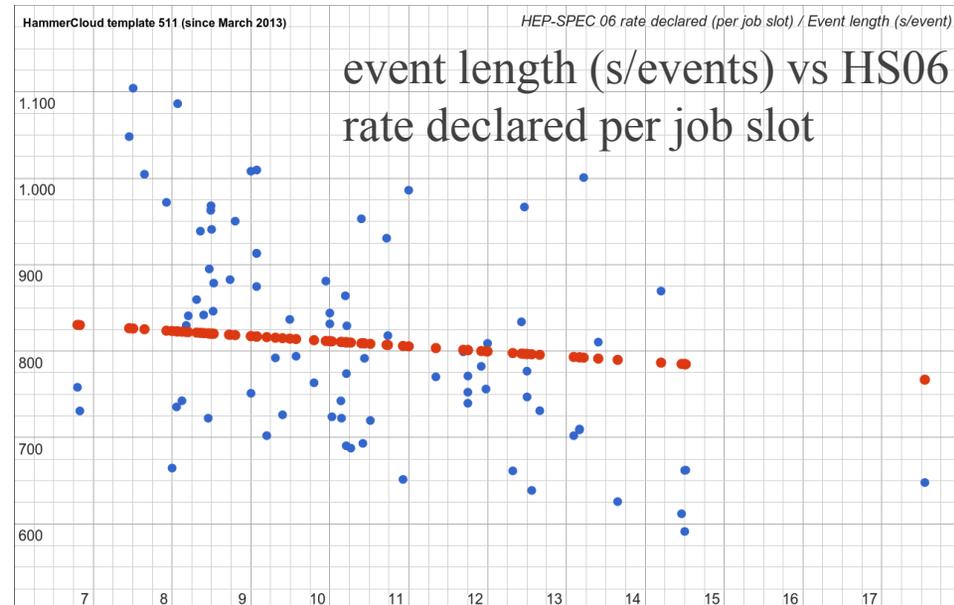
Eventrate across the board calculated as

$$\frac{\text{events}}{\text{Athena running time}}$$

The error bars a standard deviat

Typical n: 2,000

- Still not good enough! but...
- **WLCG** general interest?



MultiCore resources

- ATLAS is (almost) ready to use MCore resources for MC production activity
- Accounting: still not fully clear
(<https://twiki.cern.ch/twiki/bin/view/LCG/WLCGOpsMinutes130718>)
- Static vs dynamic allocation:
 - Few sites as of today has already provided MCore resources, some static, some dynamic.
 - We believe resource can be optimized if sites are helped to setup MCore resources in a dynamic way
- **WLCG general interest?**