

Grid Data Processing: Operations

- Rod Walker, LMU Munich, 29th Nov 2010

- MC Production, data/MC reprocessing and group reprocessing
 - mostly covered Tues&Thurs – concentrate on MC sim and interference between activities
- What's new, or assorted topics
 - CERN cloud
 - CVMFS
 - Another cloud

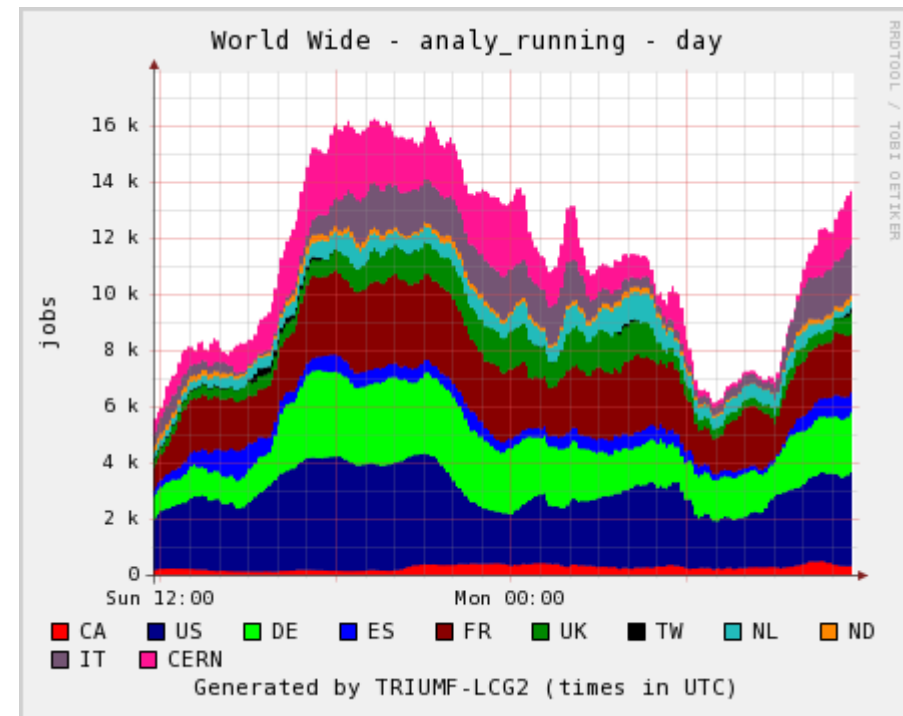

Problems and issues

- Prod is stable and shifter controlled
- Prod/Analy share is static
 - 50:50 in T2 batch system
 - ATLAS may want this changing quickly, eg HI analy
- Panda load on CERN Db
 - w/e of trouble but resolvable – see weekly meeting
- Sustaining capacity
 - Borut complains, so I go looking
 - sites offline for extended periods, maybe after problem solved
 - unclear how many jobs a site should be running
 - local analysis users. Are pledges delivered?
 - currently running jobs not the measure – some recent history sum?
 - missing SW cache – again not so obvious to shifters


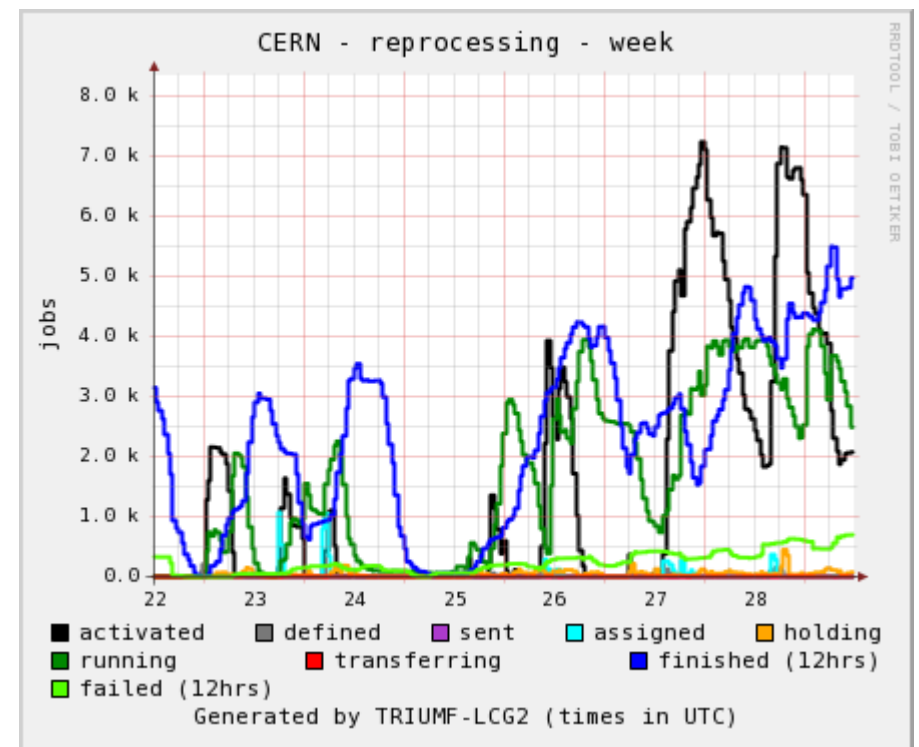
CERN cloud

- Previously used for validation, and fallback for reprocessing
 - ATLAS capacity not used, 2-4k cores available
- New: run MC production there too
- Avoid delaying validation by using single queue
 - Panda follows priorities – sim quickly makes way
- Validation jobs need different SW areas
 - UNVALID & BUILD set in task, and pilot configures
 - not quite in operation yet – capping prod pilots works too

- 11 clouds, 60k cores
- Conservation of cpu



- CERN Repro vs sim
- No sim started when repro available
- Quickly vacates
 - linearly over 10hrs
 - job kills only for the very impatient
 - takes time start 4k jobs, e.g. 1000/hr



CernVM File System(CVMFS)

- Read-only web file system with local disk cache
 - web server at CERN via local caching web proxy
- Provides exp SW, to keep VM image small
- Maybe heard about it in T3 context, but
- Some T2s have trouble providing SW area
 - space on high performance NFS disk limited
 - maybe performant NFS not available at all
 - fancy shared-FSs cannot handle ATLAS(CMT)
 - mostly Lustre, maybe gpfs too – metadata is the problem

CVMFS advantages

- All releases and caches always available(serve updated by Alessandro)
 - possibly with short delay to download missing files(local web proxy)
- Auto-SW install via WMS difficult, inevitably some sites delayed
 - works very well considering complexity(100 sites x O(100) release caches)
 - some sites may opt for cvmfs, other cannot
 - still get validated and tagged –eg. discover missing libs
- SW on local disk – scalable, fastest build performance
- Conditions data for free(also in CVMFS(for T3)) – not yet.
- Status: QMUL completely on CVMFS. RAL analysis, Wuppertal analysis, and switching Wupp prod this week
- Recommend as fallback solution for Taiwan and Lyon
 - Last w/e HI repro trouble
 - regular afs install crises, occasional performance crises

The other Cloud

- Amazon cloud, Eucalyptus, Nebulous,...
 - the way industry interprets grids
- viable to run MC sim in cloud as T2 extension, or T2 in a cloud – keep T2 for analysis
- request for interested parties to get together, I know of
 - Yoshu on Magellan cloud LBNL
 - general cloud expertise at UVIC
 - various past test projects: MPI, co-pilot BNL
- If interest then maybe mail list, or chat this week