



User Analysis Support

Massimo Lamanna

Presentation at the GDB (10-JUN-2009)



Recap...

- Several areas loosely coupled but pointing to “user analysis support”
 - How to support **users**...
 - ... avoiding the **developers** to be interrupted continuously..
 - ... and be effective in collaborate with the **sites** to improve performances and become more and more dependable
- Presented at the MB more than a month ago...
 - ... contacts with all experiments going much slower than hoped/expected...
 - ...STEP09 has some responsibility in this...



Sites

- SAM is an essential tool
 - Complemented by tools from the experiments (often in the SAM framework)
- Essential tools:
 - “Robots”
 - A cocktail of realistic analysis (access to “popular” datasets)
 - Regular job submission. Usage of results more sophisticated than in straight SAM (algorithms to put in/out of experiment blacklist)
 - “LoadGenerator”
 - Similar to the Robots, but emphasis is to provide “on-demand” service (often requested by sites) to stress test a (set of) sites
 - Emphasis on performance analysis
 - Not only #failed_jobs/#jobs... Aggregate events per second



JobRobot (CMS)

Applications Places System Tue Jun 9, 17:12 Massimo Lamanna

JobSummary - Mozilla Firefox

File Edit View History Bookmarks Tools Help

http://dashb-cms-job.cern.ch/dashboard/request.py/jobsummary#user&site&ce&submissiontool&dataset&appl python links check broken

Gmail GCal Docs News Bank RW Dict Mix Old Ubuntu and Free ... Apple Google Maps YouTube Popular Python

Grid Computing:... Gmail - Inbox (1... LCG Manageme... LCG Manageme... DistributedData... JobSummary FERRYBOAT TIM...

any user any site any ce any submissiontool any application any rb JobRobot any grid any jobtype any tier

unk pend run term
 done canc abort g-unk
 succ site-fail app-fail
 all-fail a-unk donesuccess
 submitted
 terminated

from UTC 2009-06-08 15:10:09
to UTC 2009-06-09 15:10:09
sort by site
 bars in the plot

submit

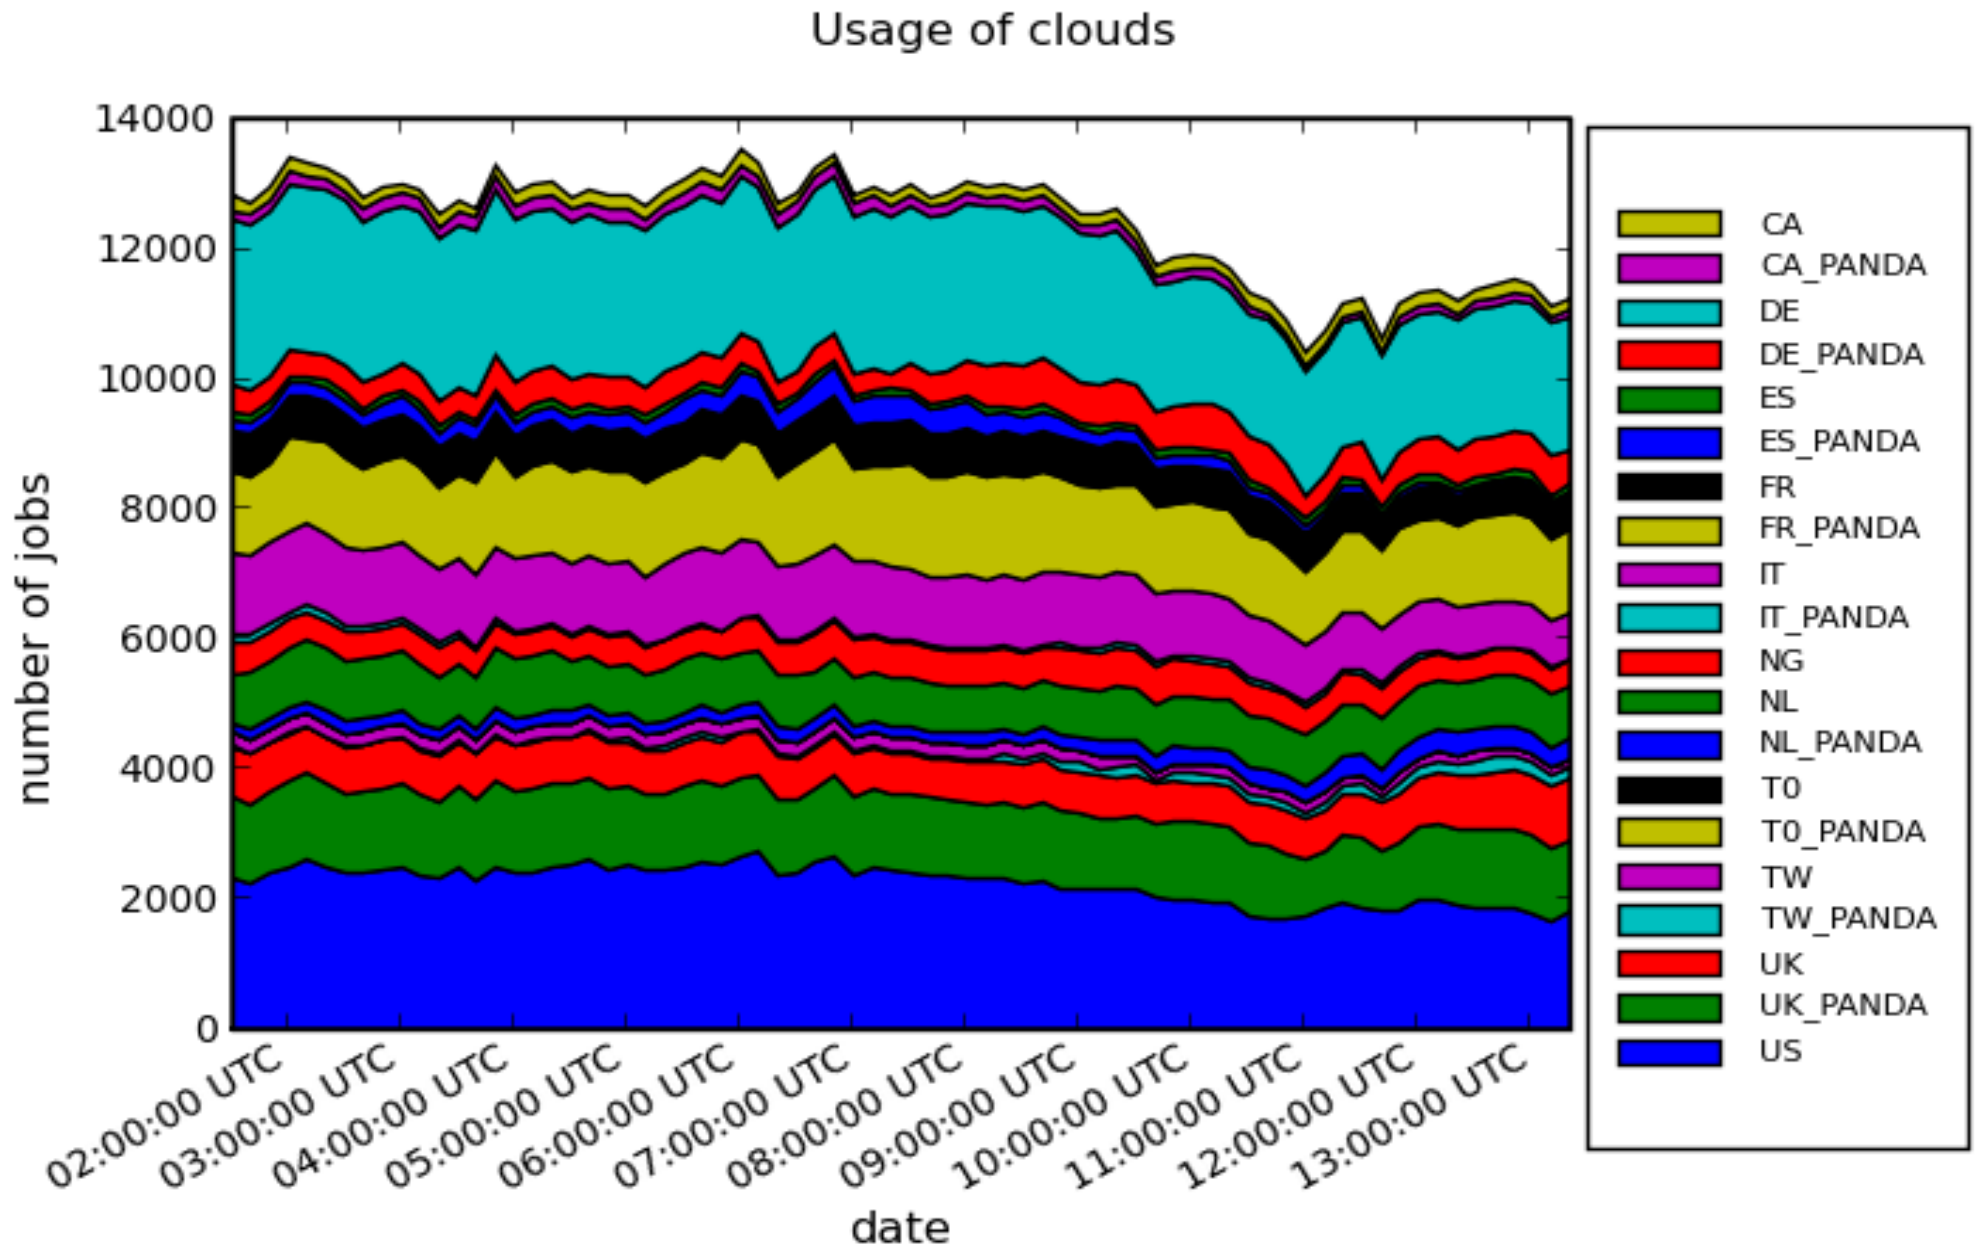
jobs per site

Site	Submitted	Application-failed	Site-failed	Aborted	Cancelled	App-unknown	Pending	Running
T2_HU_Budapest	~380	~10	~10	~10	~10	~10	~10	~10
T0_CH_CERN	~550	~10	~10	~10	~10	~10	~10	~10
T2_BE_IHE	~400	~10	~10	~10	~10	~10	~10	~10
T2_FI_HIP	~450	~10	~10	~10	~10	~10	~10	~10
T2_US_MIT	~300	~10	~10	~10	~10	~10	~10	~10
T2_RU_ITEP	~500	~10	~10	~10	~10	~10	~10	~10
T2_RU_SINP	~350	~10	~10	~10	~10	~10	~10	~10
T2_TR_METU	~580	~10	~10	~10	~10	~10	~10	~10
T1_US_FNAL	~600	~10	~10	~10	~10	~10	~10	~10
T2_BE_UCL	~600	~10	~10	~10	~10	~10	~10	~10
T2_KR_KNU	~550	~10	~10	~10	~10	~10	~10	~10
T2_FR_IPHC	~580	~10	~10	~10	~10	~10	~10	~10
T2_US_Nebraska	~550	~10	~10	~10	~10	~10	~10	~10
T2_IT_Rome	~500	~10	~10	~10	~10	~10	~10	~10
T2_US_Wisconsin	500	1	1	11	0	0	1	0
T2_UK_SGrid_RALPP	~500	~10	~10	~10	~10	~10	~10	~10
T2_DE_DESY	~500	~10	~10	~10	~10	~10	~10	~10
T2_US_UCSD	~500	~10	~10	~10	~10	~10	~10	~10
T1_UK_RAL	~500	~10	~10	~10	~10	~10	~10	~10
T2_UK_London_IC	~500	~10	~10	~10	~10	~10	~10	~10
T2_BR_UERJ	~500	~10	~10	~10	~10	~10	~10	~10
T2_ES_CIEMAT	~200	~10	~10	~10	~10	~10	~10	~10
T2_DE_RWTH	~500	~10	~10	~10	~10	~10	~10	~10
T2_US_Florida	~500	~10	~10	~10	~10	~10	~10	~10
T2_IN_TIFR	~500	~10	~10	~10	~10	~10	~10	~10

number of jobs

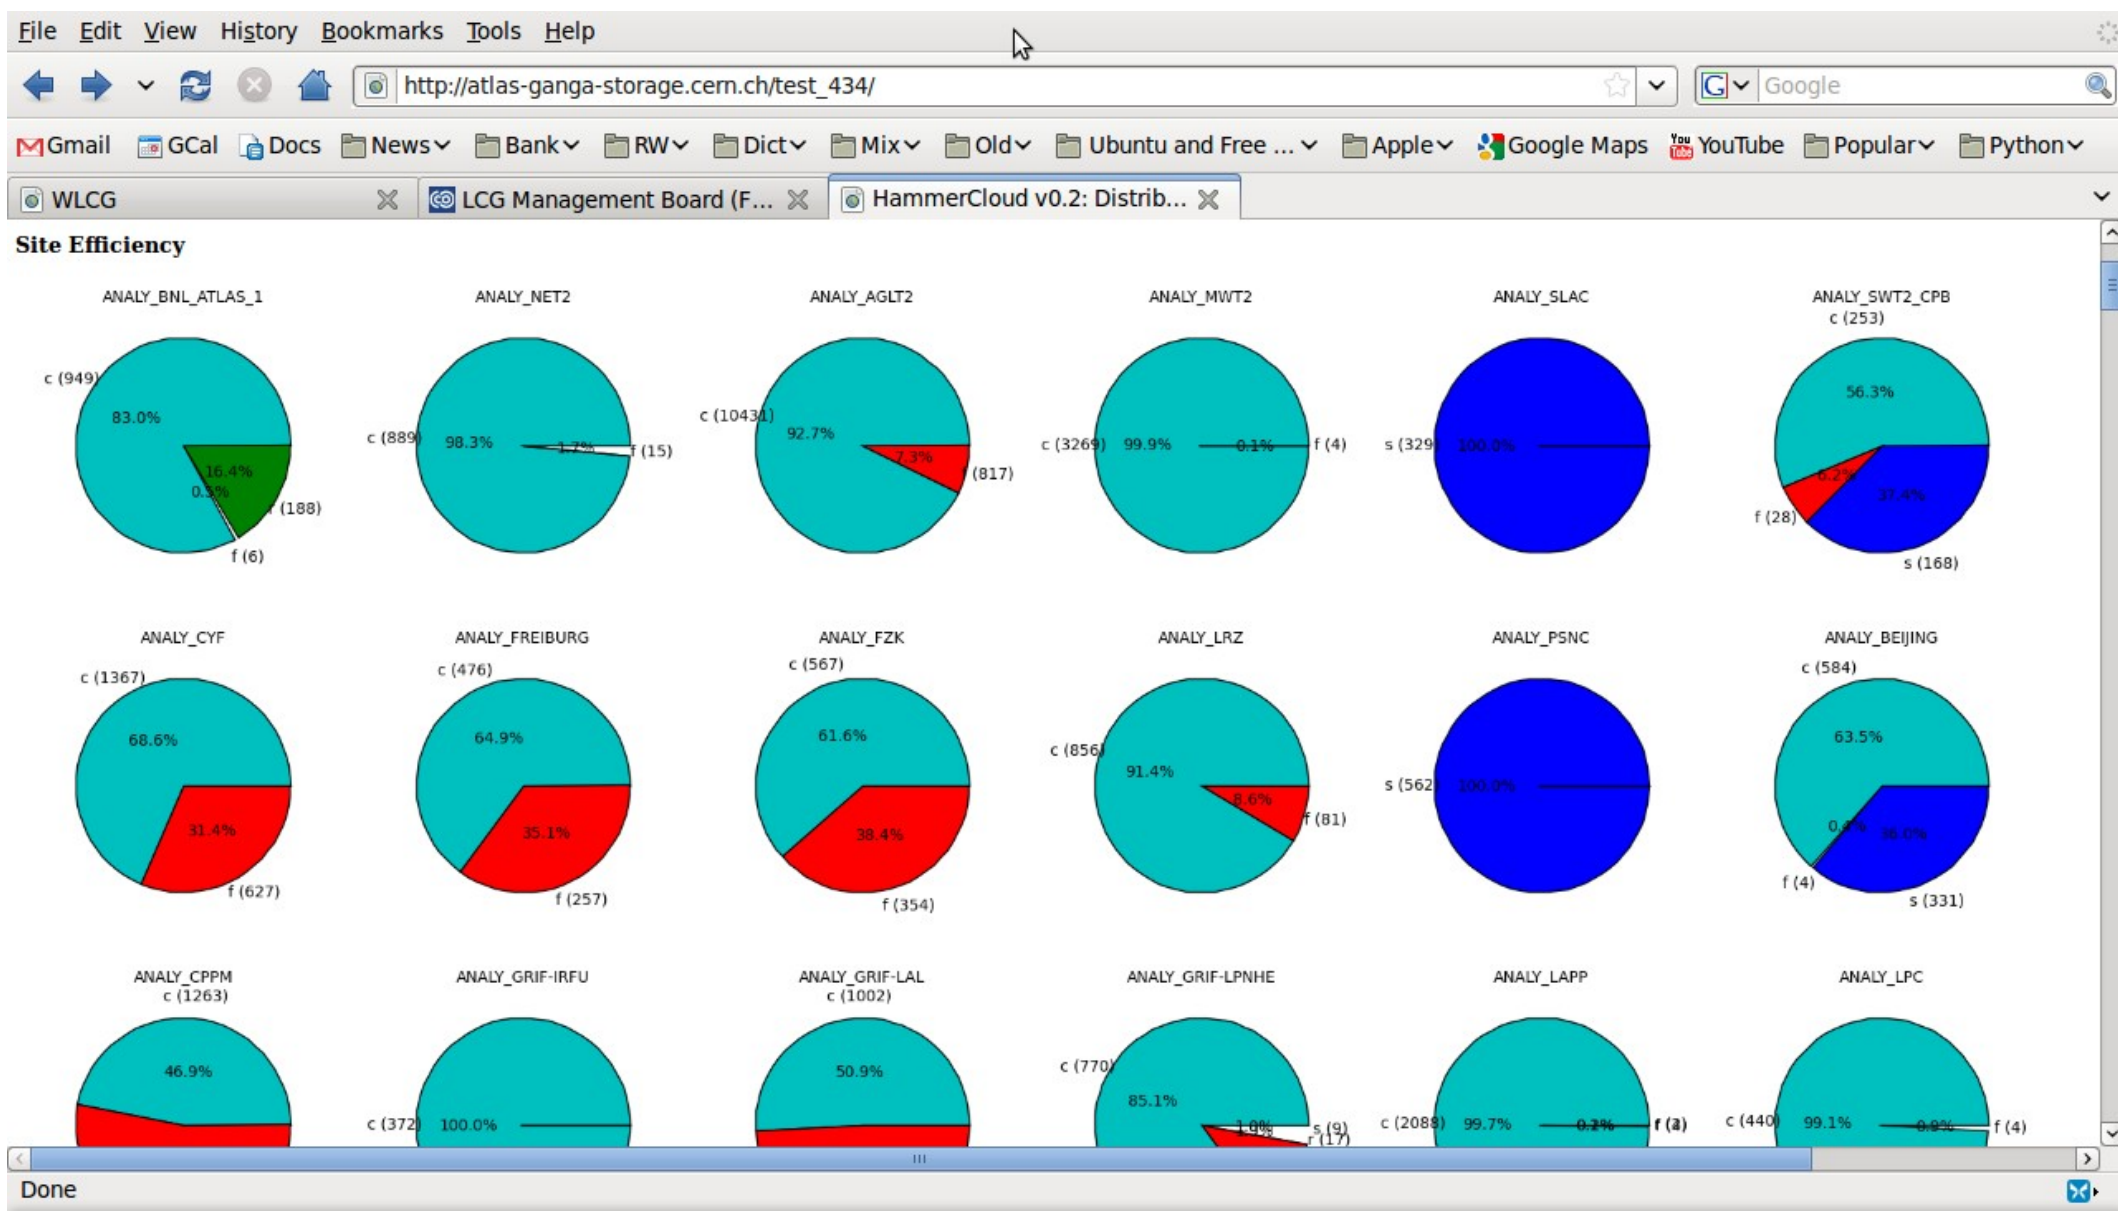
javascript:PHPbarClicked('T2_US_Wisconsin')

HammerCloud overview (last few hours)



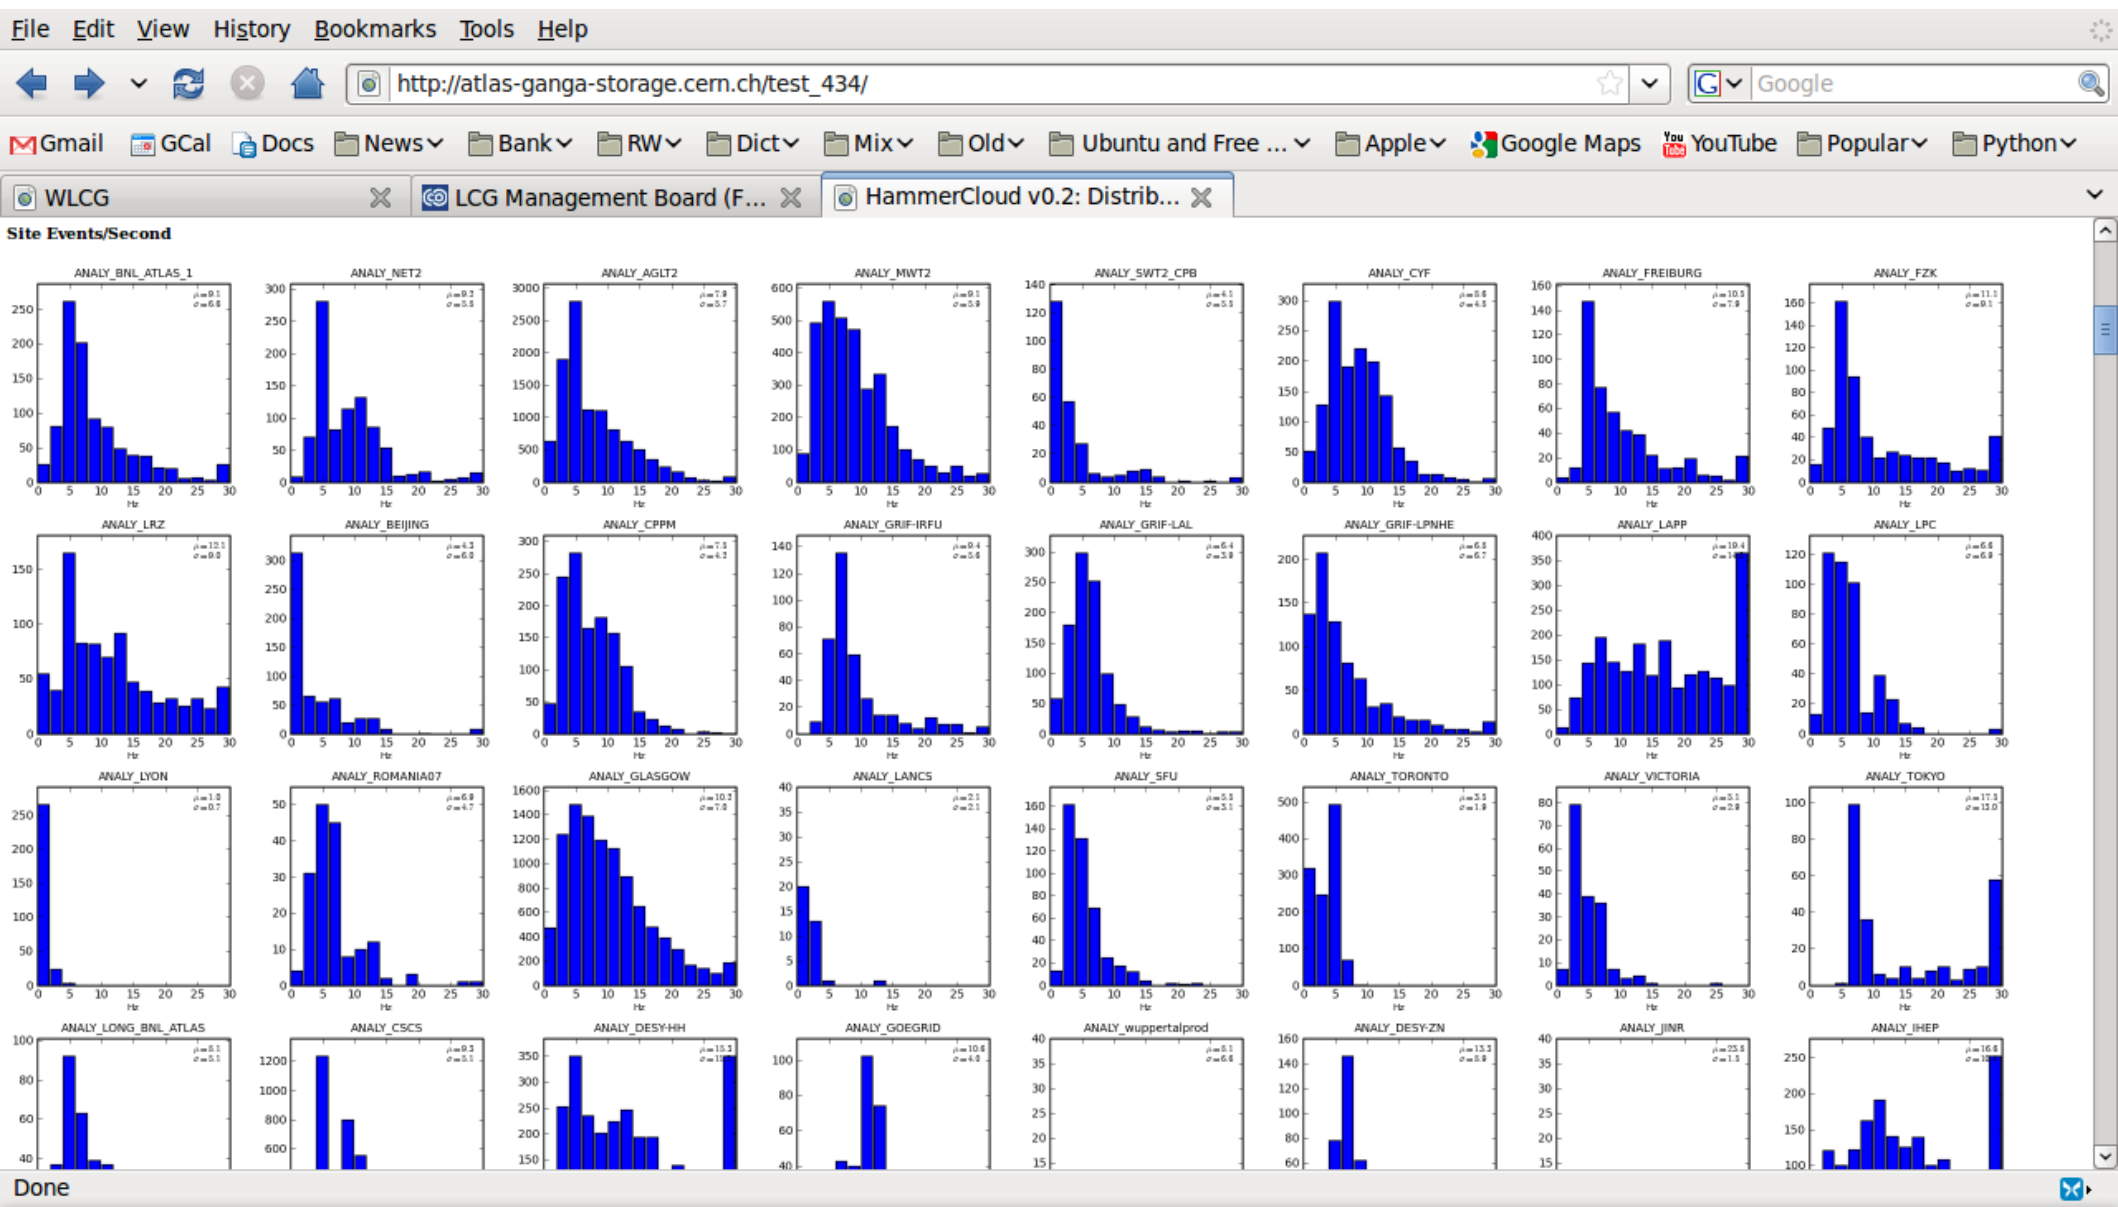


Snapshot of HammerCloud (ATLAS)





Snapshot of HammerCloud (ATLAS)





Interesting lines

- “Robots”
 - Explore commonality ATLAS/CMS (GangaRobot/JobRobot)
 - Similar systems should if not converge (on the implementation) produce similar output to be presented in a similar way
 - Put it in production for LHCb (“GangaRobot” based on the common ATLAS-LHC framework Ganga)
 - Growing activity provides “a sort of” regular test
- “LoadGenerators”
 - Put it in production for LHCb
 - Interest from CMS. Similar systems should if not converge (on the implementation) produce similar output to be presented in a similar way



What is a WLCG service?

- When the success of the experiments is at stake, no real distinction between “**baseline**” and “**experiment-specific**” services
 - Building blocks of their computing
 - Need of manage them properly (evolution, security)
 - Distribute (install) them (1-site services, Tier1-only boxes, pervasive services)

Common area of interests?

- Different approaches in the different experiments
 - Special interest for services installed at CERN
 - ATLAS has a good collaboration with FIO (installation + monitoring/alarms). Upgrade of the VOBOX model (understood and agreed procedures vz. black box)
 - CMS interest to tackle security issues. Actually this is a common concern



User support

- Even with no “site failures” the life of users is not easy...
 - Analysis submission frameworks might have issues
 - And so the data management...
 - And user software can have issues as well...
 - Just disentangling the origin can be a tough job!
 - And the (otherwise welcomed) growth of usage does not make things any easier
 - Deal with users issue effectively is key (an important component being “protect the users”)
- ATLAS (still unique?) set up a distributed shift system (DAST)



DAST lessons

- Distributed means:
 - Cover 24h with day shifts with effort
 - No need to be at CERN for this (use CERN visits to interact with colleagues, understand the detector, discuss your analysis etc...)
 - Your late-evening problem of a SLAC is taken over by shifters waking up in Europe...
 - Potentially attract more people
 - Less shifts :)
 - Require more tools!
 - You are in shift with shifters/experts in other timezones (e.g. the shifter (15pm at CERN) reproduced an user problem and discuss it with the expert (9am at BNL))
 - Information exchange, workflow control



DAST tools

- Reuse (slightly abuse) of existing tools
 - e.g. use a shared gmail as a whiteboard to handle with issues
 - gmail gets all the traffic from the egroups (hypernews)
 - Other tools are interesting and can be naturally integrated
- Although this is not deemed to be the final solution, it is an interesting prototype (and useful tool used in day-by-day DAST operations)
 - “DAST is great” (anonymous Ganga developer ;)

History Bookmarks Tools Help

https://twiki.cern.ch/twiki/bin/view/Atlas/AtlasDAST

van der ster CERN

Docs News Bank RW Dict Mix Old Ubuntu and Free ... Apple Google Maps YouTube Popular Py

Gmail - Inbox (1... LCG Manageme... LCG Manageme... DistributedData... AtlasDAST < Atl... FERRYBOAT TIM...

Gmail Issue Tracking Procedure

1. Open (unresolved) issues are to remain in the Inbox. Closed (resolved) issues are to be archived.
2. Escalated issues are to be considered open until the problem is resolved.
3. Open threads can be in one of 5 states:
 1. Requiring attention: these include any thread with no labels, or is starred, or that has an unread reply.
 2. Waiting for user response: label these WAITING, then ignore, and finally remove the label when the user has replied.
 3. Requiring urgent attention: label these URGENT, and act accordingly.
 4. Escalated: label these ESCALATED, and add a "to where" label (see (4)(1)); shifters should inform the user then close the issue when it is resolved. If an issue is escalated but still unresolved after a reasonable amount of time, we should contact "to where" for an update.
 5. Fixed in the next release: label these "FIXED in next release", and contact the user after the next release to remind/verify that the issue is fixed, and finally close. Use your judgment to decide if the thread can instead be immediately closed (i.e. if you feel that we don't need to follow up after the release).
4. Labels other than those mentioned in (3) can be used for information purposes, including:
 1. to where an issue has been escalated, e.g. Ganga Expert, GGUS, DQ2 Savannah...
 2. temporary labels used to track common issues, e.g. mc08 dataset problem.
 3. other labels for arbitrary information/tracking purposes.
5. To close an issue, remove any labels mentioned in (3), (4)(1), or (4)(2) and then archive the thread. Labels from (4)(3) can optionally remain on closed issues for reference later.

The above guidelines have the following implications:

1. To find issues needing attention, just browse to the Inbox and look for unlabeled, starred, and unread threads.
2. There will be many items in the Inbox that will not require attention. This is OK.
3. Issues will stay in WAITING until the user responds. I don't think we need to contact users if they are too uninterested to reply. I suggest we close inactive WAITING threads after 7 days.
4. Threads that we close will be automatically reopened if the user or anyone else replies to the thread. Thus, you can safely "close" a thread and it will reopen itself if the user doesn't agree with you. Perhaps this means the WAITING state is redundant, but at least I find it useful to keep these obviously open threads in view, and thus in mind.
5. If another user happens to resolve an issue without DAST intervention, just archive the thread and move on.
6. If you find a thread in an inconsistent state, try to find out its real status and correct the labels.
7. Feel free to create new labels under (4)(2) or (4)(3); please communicate their meanings to the other shifters if they are to persist, or otherwise delete them at the end of your shift week.

Shift calendar

Distributed Analysis Shifter

Today June 2009 Print Week Month Agenda

Mon	Tue	Wed	Thu	Fri	Sat	Sun
1 Jun 16:00 Sergey	2 16:00 Sergey	3 16:00 Sergey	4 16:00 Sergey	5 16:00 Sergey	6	7
8 16:00 Alden	9 16:00 Alden	10 16:00 Alden	11 16:00 Alden	12 16:00 Alden	13	14
15 16:00 Jacob	16 16:00 Jacob	17 16:00 Jacob	18 16:00 Jacob	19 16:00 Jacob	20	21

Taskbar with various application icons including Firefox, LibreOffice, and system utilities.



DAST

- One prototype exists
 - A “little” software and a lot of organisation, procedures, etc...
 - All experiments manifested interest
- Having a working prototype is a great plus since we can build on a concrete basis
 - Google seems to provide a lot of tools (in use and new component like GoogleWave)
 - We have to “rationalise” the existing tools
 - No single tool covers all use cases (e.g. Savannah, Hypernews, etc...)





Where do we go from here?

- After STEP09 start regular contacts with the experiment
- Present the progress to the MB to verify the interest and the viability of sharing/converging existing tools
 - Also new tools could be “devised” in common
 - Maybe using new technologies...
 - All in all, optimise existing effort
 - And as a minimum exchange information
 - Confident to have a good relationship with the technical coordination of WLCG (M. Litmaath)
 - If useful start a full-fledged project
 - Naturally hosted in the HEP SSC (S as Support)
 - Building on the experience and the collaborations built in IT/GS, the 4 experiments and several site experts



