Naf User Committee March 2010.

Report from Operation

Yves Kemp, Andreas Haupt,Kai Leffhalm
On behalf of the Naf Operators for the NUC
March 10, 2010

Usages Statistics

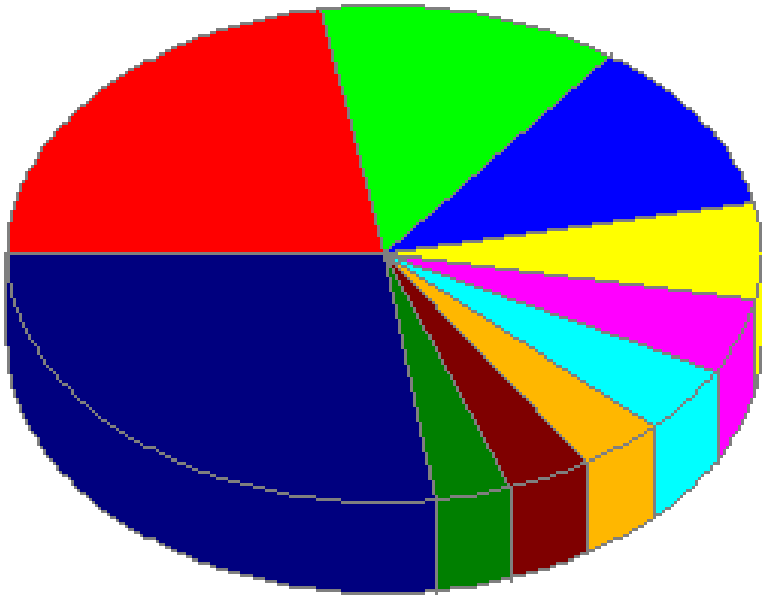

- Wallclock time in percent by project

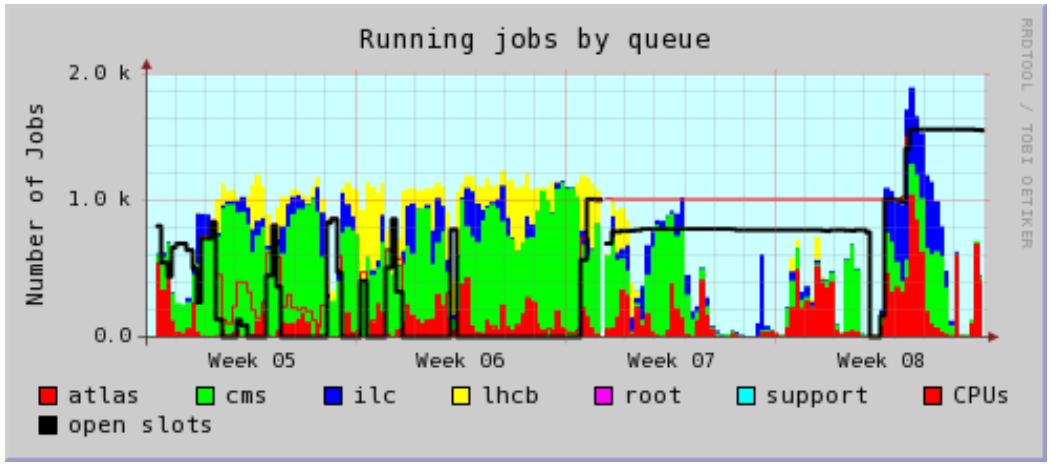

- Workload for the projects from Feb. 1st till Feb. 28th
- Some problem discovered end of Feb with the cpu count
- Slot count will change soon, configuration will be changed

Grid HH workload

CMS:

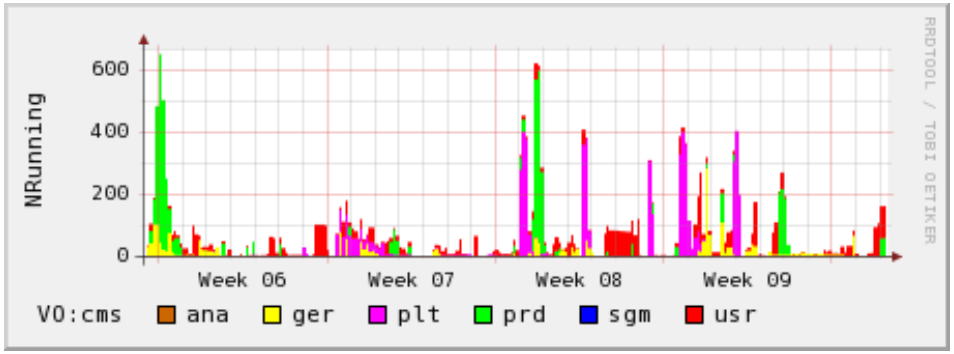

ATLAS:

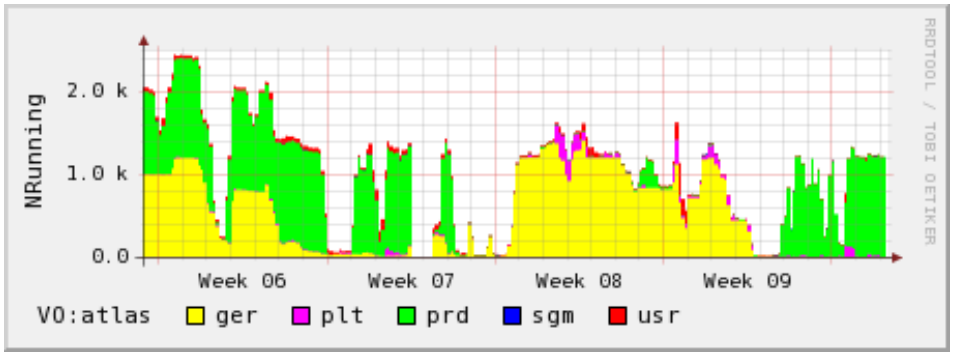

ILC:

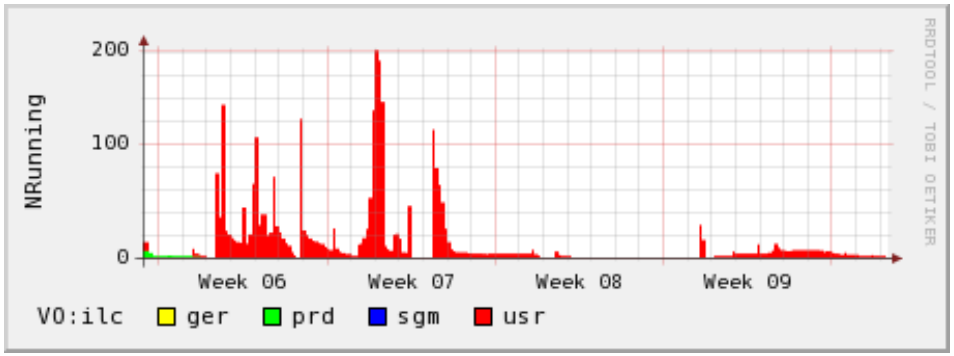

Web Stats

/interactive_login/	103
/naf_account/	78
/naf_storage/	73
/working_with_the_local_batch_system/	72
/monitoring/	64
/naf_overview/	59
/atlas/	58
/naf_storage/working_with_dcache/	55
/working_with_the_local_batch_system/best_pract...	51
/working_with_the_local_batch_system/switches_f...	47
/faq_and_support/	45
/using_the_grid/	40
/faq_and_support/login_faq/	37
/news/	36
/software/	33
/naf_storage/working_with_afs/	33
/nuc/	29
/cms/	30
/naf_storage/working_with_lustre/	26
/naf_features/autoproxy/	23
/working_with_the_local_batch_system/interactiv...	22

Web Stats cc.

/working_with_the_local_batch_system/requesting...	20
/naf_features/setup_environments/	19
/faq_and_support/hints_problem_debugging/hints_...	18
/faq_and_support/batch_sytem_faq/batch_faqs/	15
/working_with_the_local_batch_system/example_sc...	12
/faq_and_support/general_faq/	11
/ilc/	9
/lhcb/	9
/working_with_the_local_batch_system	8
/atlas	7
/nuc/admin_information/	6
/interactive_login	5
/nuc	5
/e15/index_eng.html	4
/news/index_eng.html	4
/naf_features/find_another_wgs/	3
/interactive_login/index_eng.html	3
/naf_storage/working_with_lustre	3
/using_the_grid	2
/faq_and_support/hints_problem_debugging/hints_...	2
/faq_and_support/hints_problem_debugging/index_...	1

Downtime

In general

- Took one hour longer than planned
- Should have been the last bigger downtime for a long time
- Lustre update will always need a total downtime

Upgrades

- Updated InfiniBand switch with new firmware and new line card to connect more server
- HH lustre instance has been worked on to improve stability

Lustre Problems

LBUG

- freezes processes, which are connecting to Lustre
- Patch was deployed on Feb. 26th
- . . . will be activated with reboot on every single client
- Coordination with experiments needed for IN

Processes using Lustre take 100 % cpu

- Not solved
- Proposed workaround in testing

Other Problems

Power loss

- . . . in one rack due to hardware defect
- afs file server was involved
- $\Rightarrow$ whole NAF down
- Hardware was moved to other racks, new IN installed
- Some network problems induced due to breakdown
- $\Rightarrow$ Emergency power supply will be installed
- Moving hardware shouldn't be necessary any more

DNS problems

- Some general DNS problem due to load balancer update $\Rightarrow$ is solved

Hardware problems

- Only few other broken hardware

Action Items

Recalculate batch fair share after addition of new hardware

Atlas	25.7%
CMS	43.6%
ILC	4.2%
LHCB	26.5%

Check, if SGE allows to change user priorities within an experiment by the experiment admins

- Ongoing discussion
- Will be done, but still not clear how
- Proposal: create sub groups and tell us the share for those subgroups
- Add and remove the users who should profit from the prioritized share

CMT problem (NAF invite experts/ATLAS)

- Offline discussion: new proposal from NAF is under investigation

Action Items cc.

NAF SL5 migration (running)

- LHCB and ILC only SL5 IN
- Atlas and CMS still with some SL4 IN left

deletion model for /scratch

- `find` is too slow
- Ongoing work: a tool to scan the databases directly
- The tool would be able to create a list with old files
- After verification by experiments a sublist is deleted by operators
- Still some problems with the script not solved

AFS scratch space creation

- Scratch space in afs is in work now
- New hardware is ordered

Action Items cc.

Multicore batch job monitoring

- Not big usage at all: 160d wall clock time requested

Exp.	Single core wall clock time in days
atlas	6960
cms	10611
ilc	3685
lhcb	2750

Feedback for documentation changes - Well?

- Some advice on user files storage (code development, small files) added
- glite 3.2 documentation is done

User information (news letter, motd, news section on naf web, ...)

- News section is available but not well read, news letter is kind of static

Conclusion

- Too much (unplanned) downtimes: we work on this
- Lustre still major trouble maker $\Rightarrow$ no other solution at the horizon
- Tickets from users are helpful for problem investigations

Some statistical calculation

Computation of H_Δ, the total computing time (CPU or Wall) folded with the HEP-Spec coefficients of the corresponding jobs:

$$H_\Delta = \sum_{i \text{ in } \Delta} \omega_i t_i$$

with t_i being the time spent by the job i on a node with HEP-Spec coefficient ω_i. We sum over all jobs in a time period Δ, which is e.g. one month.

$$= \sum_{i \text{ in } \Delta,a} \omega_a t_i + \sum_{i \text{ in } \Delta,b} \omega_b t_i + \ldots$$

ω_a... are constants: the HEP-Spec coefficients of nodes of category a ...

$$= \omega_a \sum_{i \text{ in } \Delta,a} t_i + \omega_b \sum_{i \text{ in } \Delta,b} t_i + \ldots$$

$$= \omega_a \eta_a n_a \Delta + \omega_b \eta_b n_b \Delta + \ldots$$

η_a is the job occupancy of the nodes of category a, of which there are n_a. Assume $\eta_a = \eta_b = \ldots = \eta$ as the batch system treats all cores identically.

$$= \eta \Delta (\omega_a n_a + \omega_b n_b + \ldots)$$

Some statistical calculation cc.

Introduce the weighted mean of $\omega_m = \frac{\omega_a n_a + \omega_b n_b + \ldots}{n_a + n_b + \ldots}$

$$= \eta \Delta \omega_m (n_a + n_b + \ldots)$$

$$= \sum_{i \text{ in } \Delta} \omega_m t_i$$

$$= \omega_m \sum_{i \text{ in } \Delta} t_i$$

$\sum_{i \text{ in } \Delta} t_i$ is a quantity that is already now computed by the batch system. ω_m is only a constant scaling factor.