

CP3 computing satisfaction survey

Results and lessons

Participation

Active users: 42 (one connection last month)

Survey participation: 27 (64%)

- 14 from CMS (80% participation)
- 11 from Pheno (60% participation)
- 2 from NA62 (67% participation)

Thank you for your time !

Results: usage

Total usage (regular)

	Ingrid-uiX	Ingrid cluster	Lxplus	Lx cluster
Pheno (%)	75 (35)	90 (50)	20 (10)	0 (0)
CMS (%)	100 (80)	100 (50)	100 (50)	60 (15)


Alternative solutions are:

- personal laptops
- LHC Grid
- Foreign computing clusters

Why do you run elsewhere ?

- CP3 user interface performances
 - big item, let's discuss that later
- Easy sharing of data with CERN collaborators
 - would xrootd on home help ?
- Access to EOS files for tests
 - xrootd could be used as an alternative
- Everything setup there before arrival
- Access to CMSSW nightlies

Unsatisfaction



Unsatisfaction causes discussed

Support efficiency (1 report):

- specific issues, discussed with reporter

Job waiting time (1 report):

- current system: short-lived priorities
- could be improved by having a “test” queue
- could be included in next year’s switch to SLURM

Documentation issues (2 reports):

- Added a “Getting started” page to the wiki
- Updated cluster related pages
- Please report outdated / misleading wiki pages

Unsatisfaction causes (II)

Size of storage (2 reports):

- we got no details (scratch ? home ?)
- home was enlarged from 3 TB to 5 TB
- specific needs can be discussed

Available software (3 reports):

- CMSSW nightlies are only available at CERN AFAIK
- SL6 = older softs, more stability
- If recent versions of softwares are needed, please ask, we can install nearly anything

Unsatisfaction causes (III)

Cluster reliability (5 reports):

- working nodes availability is usually very high
- some problems with cooling during summer
- this should improve with the new computer room in 2016
- do you have experiences to share ?

UI reliability: situation

in 2014:

- Number of active users grew from ~20 to more than 40
- Some software became more demanding (amc@nlo)

December 2014:

- Installed high-performance replacement for UI
- Storage moved to local disks

2015:

- 5 reboots of the UI in February due to memory issues
- 6 reboots since March 1

UI reliability : more details

February issues:

- memory sharing between OS and Filesystem (ZFS)
- understood and treated (ZFS memory limited, no swap)

Later issues:

- user processes using 40GB+ memory (various sources)
- linux kernel unable to manage them (why ?)
- actions taken:
 - improved cache management (more available memory)
 - hard limit on sum of users memory
 - vastly improved monitoring toolbox

UI reliability : on your side

- We reboot without notice only when there's no other solution (everything is already dead)
- Reboots are expected. Overall availability is still ~99%
- If your script requires UI availability and stability for more than O(12 hours), you may want to:
 - add a checkpointing system
 - improve job management

UI performance

Load peaks are strongly correlated with:

- MG family software (normal) use
- multiple parallel read/write to disks
- scripts using high amount of memory (15GB+)

These are normal: we seem to need (even) more hardware

- buying a second UI this summer
- separate Pheno and CMS groups on different UIs
- separate storage for Pheno and CMS
- In the meantime, heavy MG jobs and tests should run on **ingrid-ui3** as usual

Your suggestions

- Good practices and cluster tutorials:
 - usually given every year, except 2015 (yet) but:
 - very low attendance (10 non-academics)
 - no effect on remote users
- Automatic job testing before submission:
 - not applicable to any pheno soft
 - very unpractical to setup
 - lots of time lost for users

Your suggestions (II)

- Better documentation starting point:
 - See the new [Getting started](#) wiki
 - suggestions are welcome
- Emergency computers for guests
 - One windows laptop always available
 - I added one Linux laptop
 - We have a macbook that we could add

Tutorial requests

CP3 cluster etiquette:

- we will continue to do it once a year
- will be updated according to this survey

Parallel computing:

- CP3-specific, condor: same as above
- General: CISM trainings in october

Linux:

- CISM trainings

Conclusions

- High participation !
- We learned a lot about your experience
- The biggest problem are the UI
 - reliability
 - performance
- Some improvements have already been made
- Some are in progress
- I would like to do it again next year: what do you think ?