



CC-IN2P3

Contribution to DC2013



- Storage
- Computing
- Cloud
- Database

Rachid Lemrani

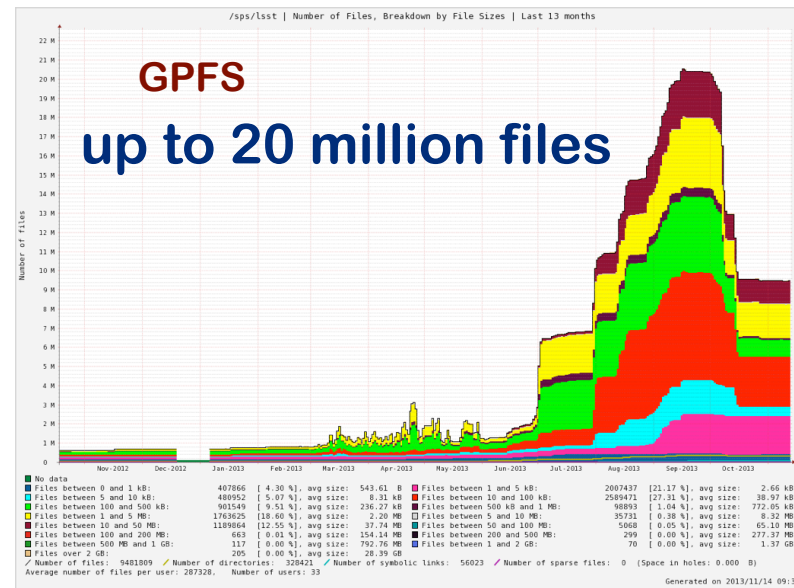
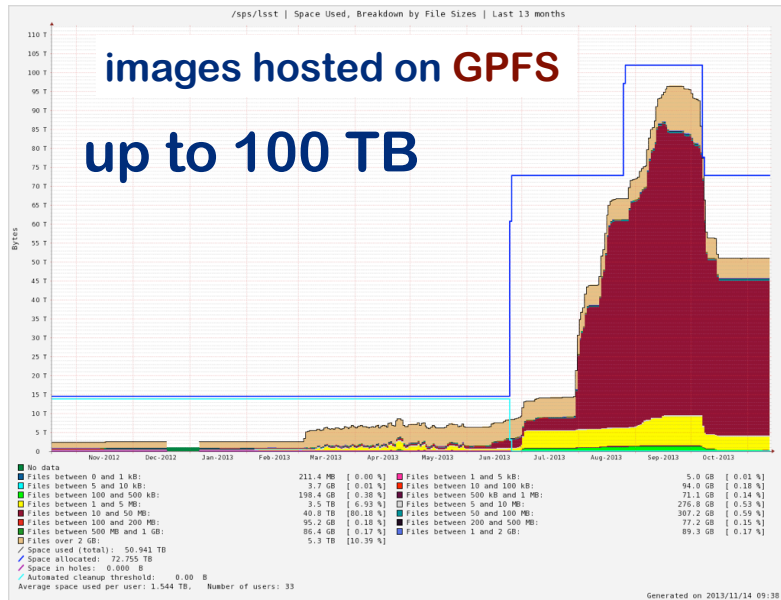
rachid.lemrani@cc.in2p3.fr

LSST meeting - December 18th, 2013

l r f u
—
c e a
—
s a c l a y



DC-2013 : Storage



Some additional storage temporarily provided, borrowed from other groups ...

➔ Better anticipate the needs in the future

Currently > 8 M files on sps and > 1 M files on AFS


➔ Reduce number of files / use only sps (not AFS) ?



DC-2013 : IRODS



IRODS used for data transfers, data management



Need to use tapes and reduce disk space :


➔ **IRODS/ HPSS interface**

Data management takes long :

➔ **(again) reduce number of files**



Number of jobs



← Up to 700 batch slots
accessing /sps/lst

all LSST jobs : 325 KHours
Data Challenge : 260 KHours

DC jobs submitted on dedicated
batch project P_Isst_prod
restricted to Isstprod

DC successful jobs :
215 KHours wallclock
92 KHours cpu


➡ **43% CPU efficiency**



Number of jobs per day



Integral 3.684e+04



restricted to “lsstprod” jobs
that didn’t fail : ~ 37 000 jobs

Numerous types of jobs
launched somewhat sequentially

➔ Significant amount of deadtime




CPU per week



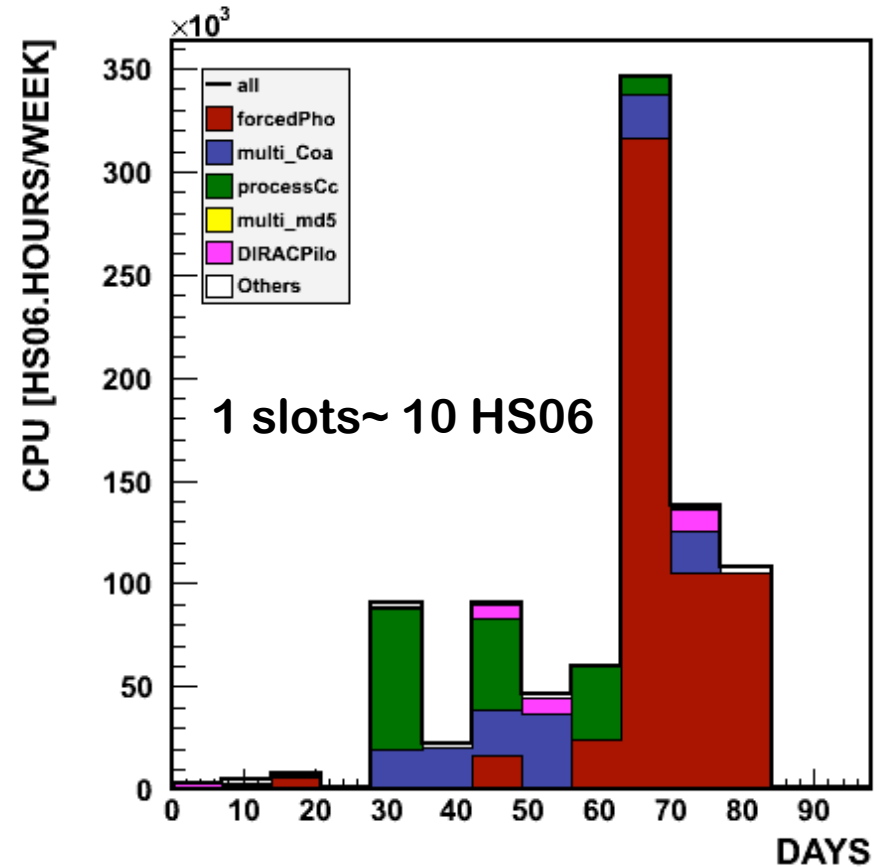
Walltime (Hours)

Integral 2.151e+05



CPU (HS06.Hours)

Integral 9.218e+05



➡ co-add and forced-photometry


➡ forced-photometry



Efficiency of jobs






Integral 3.66e+04



wallclock weighted

Integral 2.149e+05




-  - Forced-photometry better after upgrade of MYSQL server
-  - Coadd jobs very inefficient (high I/O) : partly bottleneck due to few SL6 workers at that time running several jobs at the same time



Memory usage



Integral 3.526e+04



➔ OK



LSST and Cloud




Test running LSST software using Openstack

Production of simulated images
using all available cores of the
virtual machines

Total production : 150 000 CPU hours
70 to 100 hours for one pair
of images (2 x 15 s exposition)
on a VM with 20 cores

1single CDD : 1/189 focal plan



credit : Dominique Boutigny



DC-2013 : Database



See Osman Aidel's talk :

**Mysql Server : bi-proc, bi-core 40 GB RAM, 9TB Disk
Space Usage of 4,3 TB : 770 GB (18%) indexes 3,6 GB (82%) Data,**

- Main issue to overcome :

- ENDLESS INDEXING ~ 2 billion lines for each the 5 filter tables

Trick : - Load tables without indexing

- Switch-on indexing with appropriate configuration

status : Indexing takes 15 hours for each table

-> The MySQL database will be available for physics analysis



Large scale test at CC-IN2P3



See Fabrice Jammes' talk

Tests “Very usefull uncovering unexpected issues“

Largest QSERV platform ever with 310 nodes

From July to Septembre 2013

IBM iDataplex 8 cores / 16 GB RAM / 160 GB disks (~130 Go for data)

- The machines are installed by CC-IN2P3 :
Automatic installation of the nodes using puppet
- QSERV service is administrated by CC-IN2P3 :
Automatic installation of the QSERV software,
Extensive adaptation and correction of original scripts
- Distribution of the data on the local disks of the nodes
from high-throughput GPFS disk space



Conclusions



- Very challenging DC with only 5TB of input RAW data
- Coadd very I/O intensive : low cpu efficiency
- Huge amount of files : ~ 20 M
- Huge Database : 2 B lines per table
- “Heavy” implication of CC-IN2P3 Staff : Database, Sys. Admins and Storage