



Lyon Analysis Facility

- status & evolution -

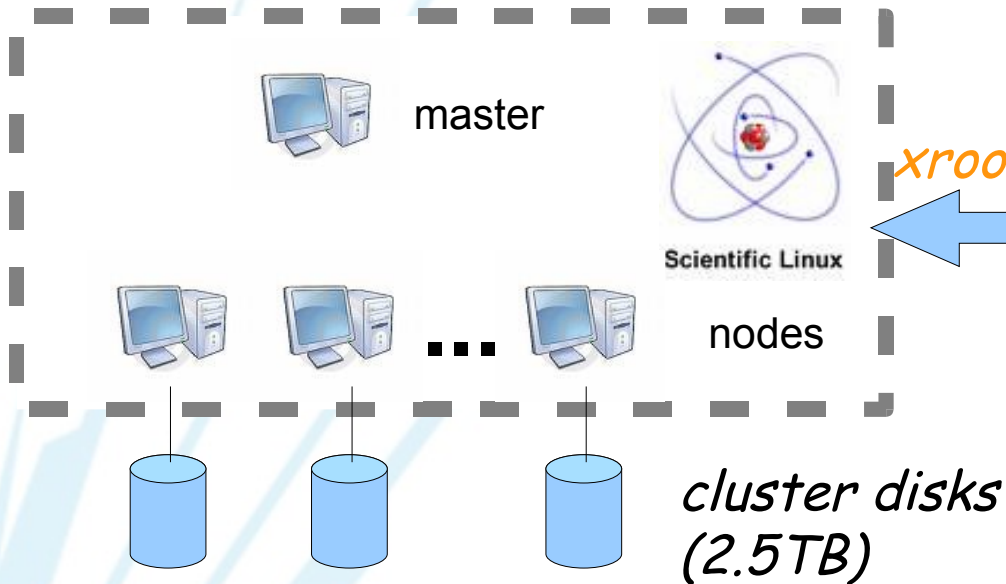
What is PROOF?




Parallel ROOT Facility

- Designed to be an alternative to GRID for data analysis
 - CPU power can be large + results are obtained quickly
- Code is submitted by the master and balances the load between the slaves (available CPUs)

PROOF



xrootd SE



(this is our current setup)

PROOF and LAF



■ Pros:

- ◆ ROOT = analysis software common to all HEP experiments (C++ based)
- ◆ Few requirements on OS and system
- ◆ Experiment-specific libraries are compiled on proof using PAR files (tar archives)
- ◆ ⇒ can be shared between several experiments/VOs
- ◆ Priorities can be defined

■ Cons:

- ◆ Main constraint is on analysis code design : code must inherit from ROOT's **TSelector**
- ◆ Not fully “interactive” ... but monitoring possible on each worker
- ◆ Can communicate with only 1 datafile manager : **xrootd**

■ Our current setup

- ◆ all nodes = Intel(R) Xeon(R) CPU / 8 cores [2.66GHz] / 16GB RAM / SL4
- ◆ 1 **master** (ccapl0001) + 20 **slaves** (ccapl00[02-21]) = 160 workers
- ◆ SE (xrootd) : 17TB
- ◆ Identification to PROOF based on certificate

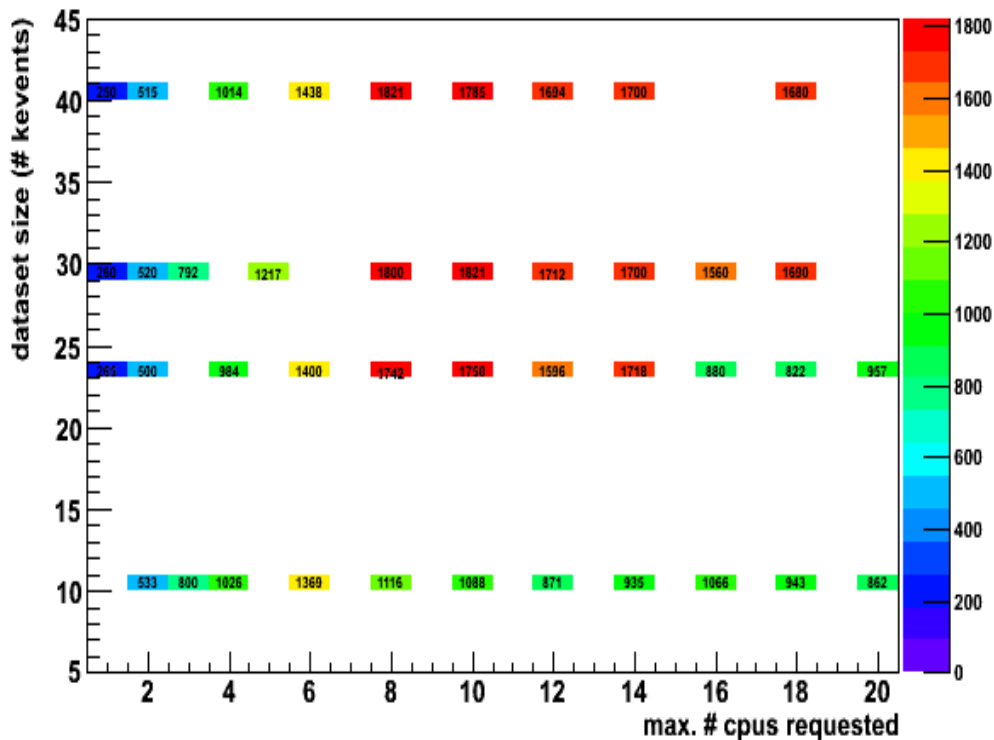


Live demo...

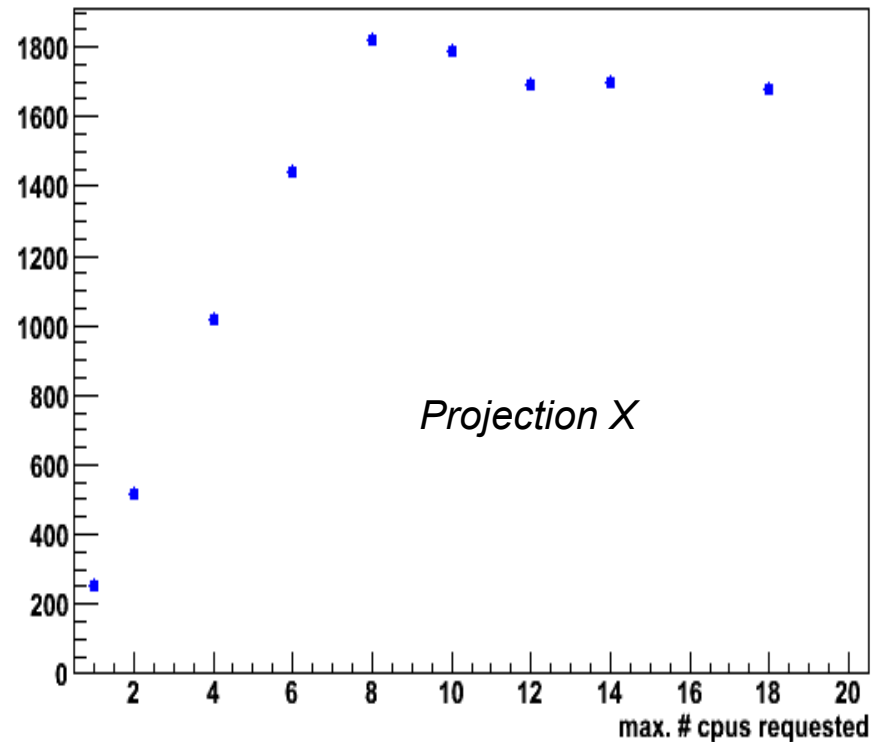
A few results (with ALICE software)



overall number of events per second



overall number of events per second



Saturation a 1800 events/s for typically 8-10 workers
 Data rate ~ 130 MB/s ~ bandwidth between data and workers

- only limiting factor ? No other effect behind ?
- answer will be given with a larger bandwidth

Situation for experiments :

ALICE



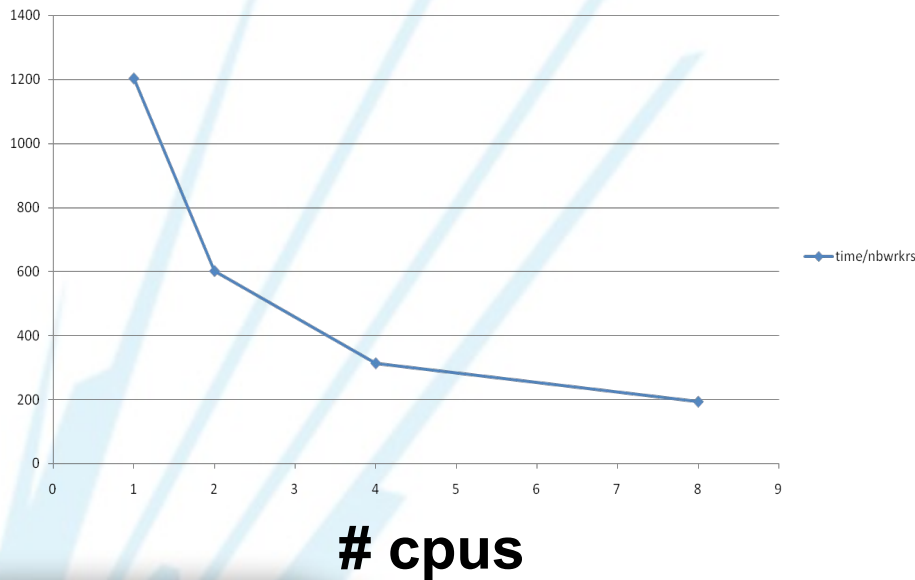
- PROOF already part of the computing model
- Extensively tested and used for several years
- Analysis tasks belonging to AliRoot run on both GRID and PROOF
- Unique issue :
 - ◆ Direct access to data on GRID and staging
 - AliEn services must run on Solaris, which is not the case yet
 - Work ongoing, hopefully completed soon
 - ◆ Other experiments may face that problem too (?)

Situation for experiments : ATLAS

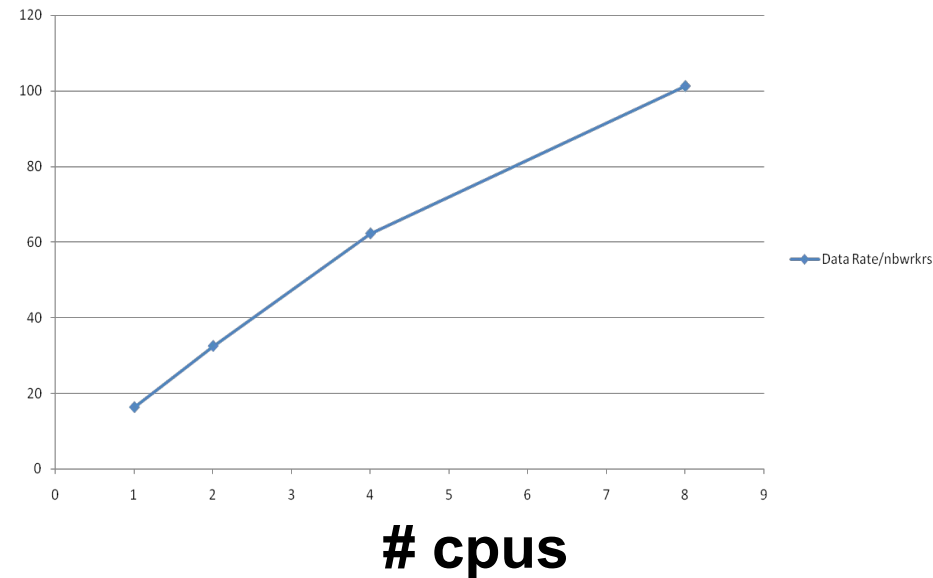


- Contact person: M. Airaj @ CEA saclay
- Tests successful
 - ◆ Library compilation OK
 - ◆ Access to data OK
- Analysis code run on 500k events stored on the xrootdSE

Processing time (s)



Data rate (MB/s)





Situation for experiments : CMS

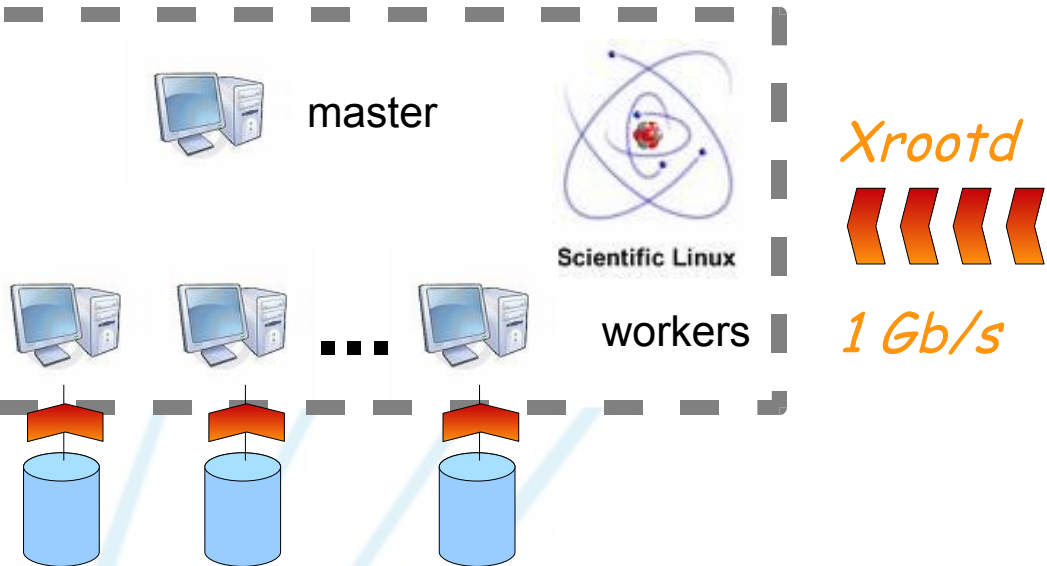


- Contact persons: M. Bluj, A. Kalinowski @ LLR
- No real tests performed yet
 - ◆ Not all CMS libraries can be loaded
 - ◆ Probable mismatch between ROOT CMS flavour and bare ROOT used by PROOF ... under investigation


Evaluation of different storage possibilities



PROOF



xrootd SE



- Read data from
 - ▶ “External” storage : dedicated xrootd SE
 - ▶ “Internal” storage : hard disks of the cluster nodes

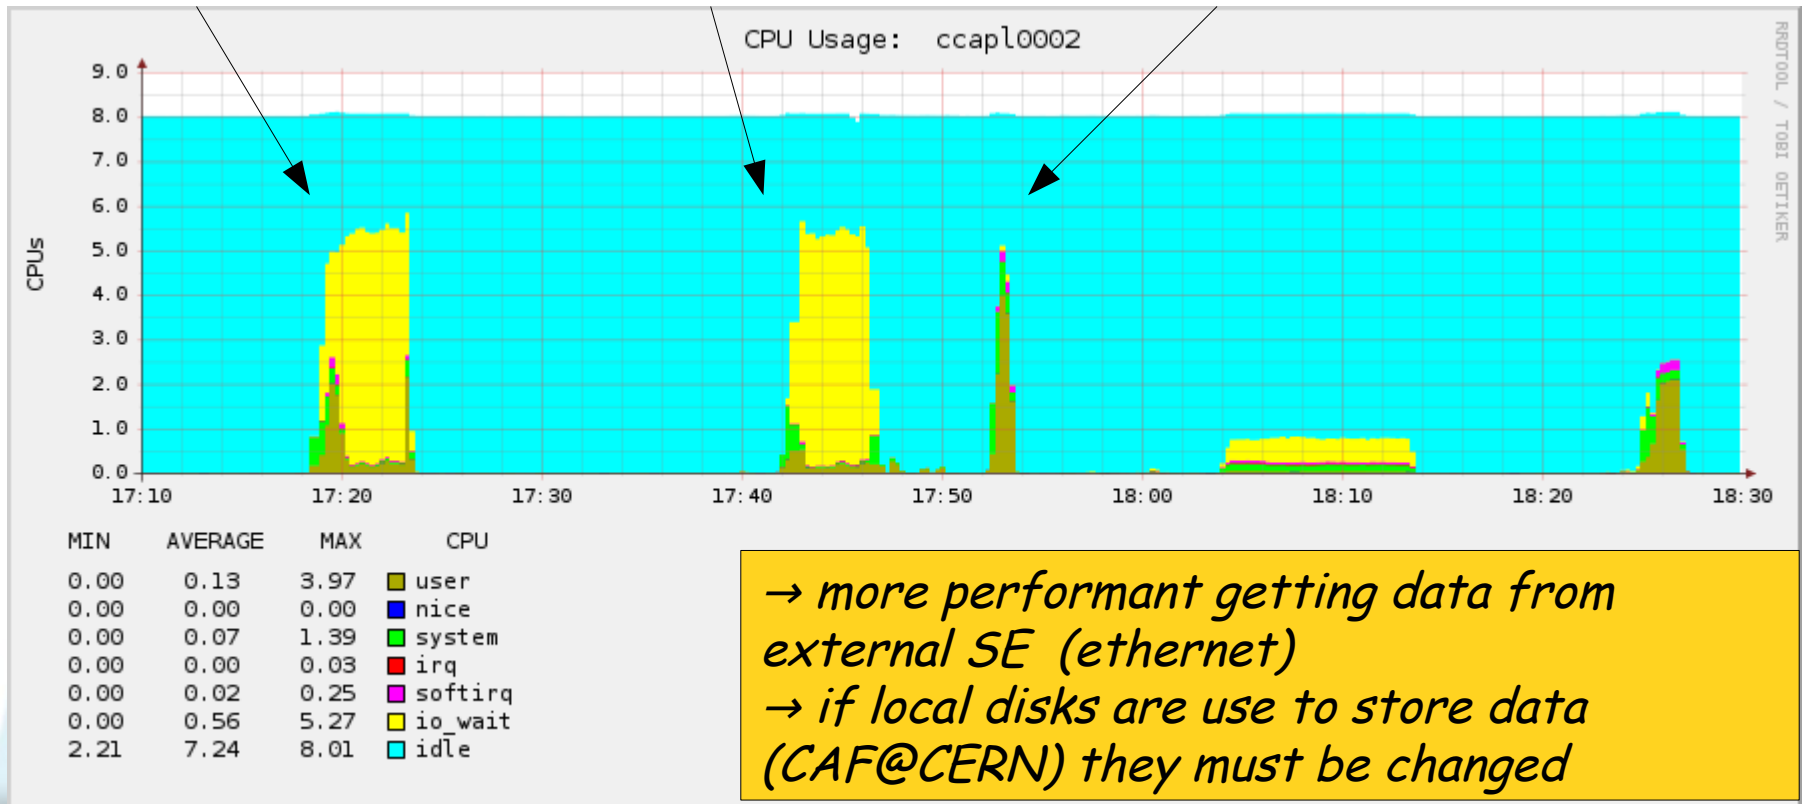
Performance benchmark (1 node only)



Analysis on cluster disks :
Worker reads data on another worker's disk

Analysis on cluster disks :
Worker reads data on its own disk

Analysis on Xrootd storage :
Worker reads data on SE





Summary



- LAF is working ...
 - ◆ Successful tests from ATLAS, ALICE
 - ◆ CMS feedback necessary
- ... but can & must be improved
 - ◆ bandwidth between SE and nodes
 - ◆ interface with GRID services (data staging)
- Tests from several simultaneous users needed
 - ◆ → stress the system !
 - ◆ results will help defining future improvements
- Possibility of using SSD disks for fast access (cache)
 - ◆ can that be supported by PROOF/xrootd ?
 - ◆ to be investigated
- No constraint on design, open to all good solutions