

French sites - status report

ALICE T1/2 Workshop

Strasbourg, May 2017

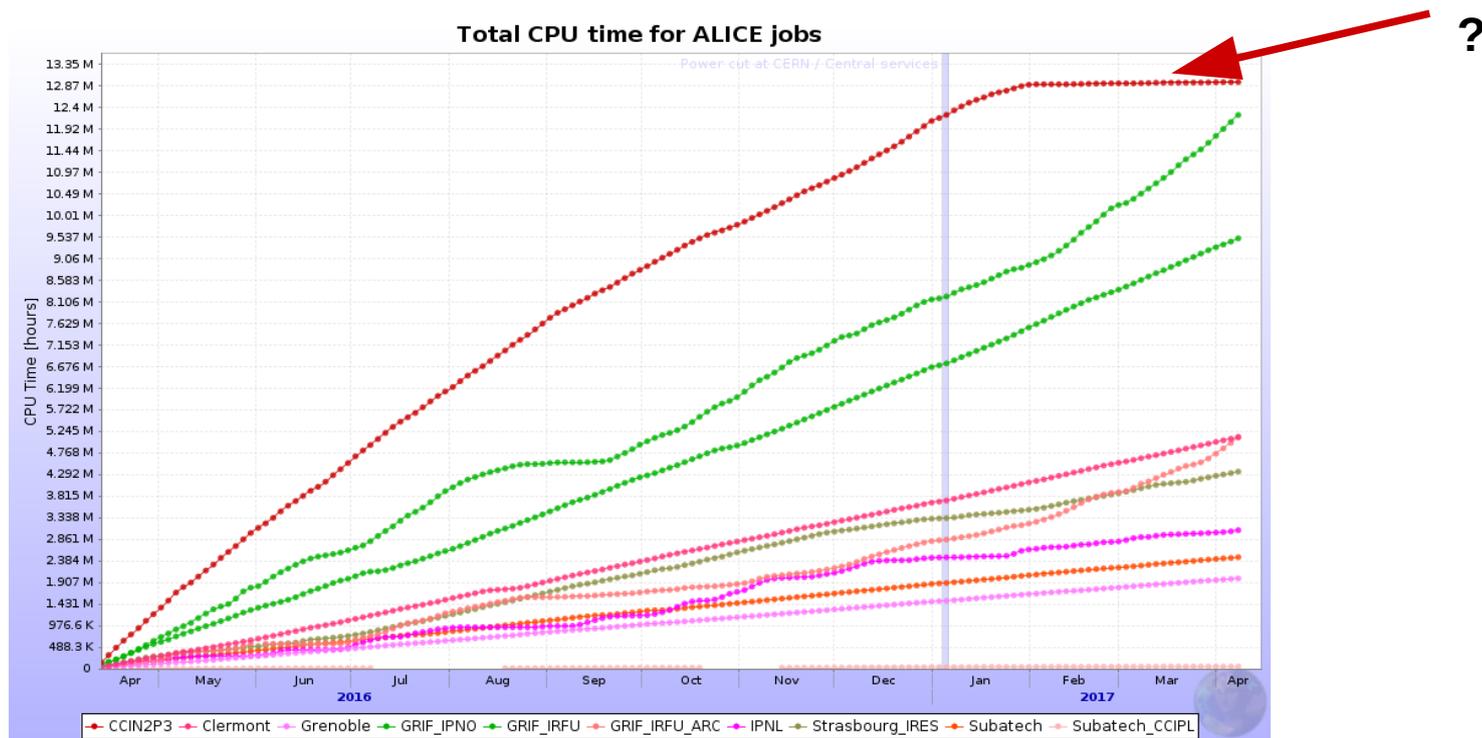
R. Vernet, for the French 'cloud'

ALICE computing in France

- 6 cities, 10 sites
 - Clermont, Grenoble, Lyon, Nantes, Paris, Strasbourg

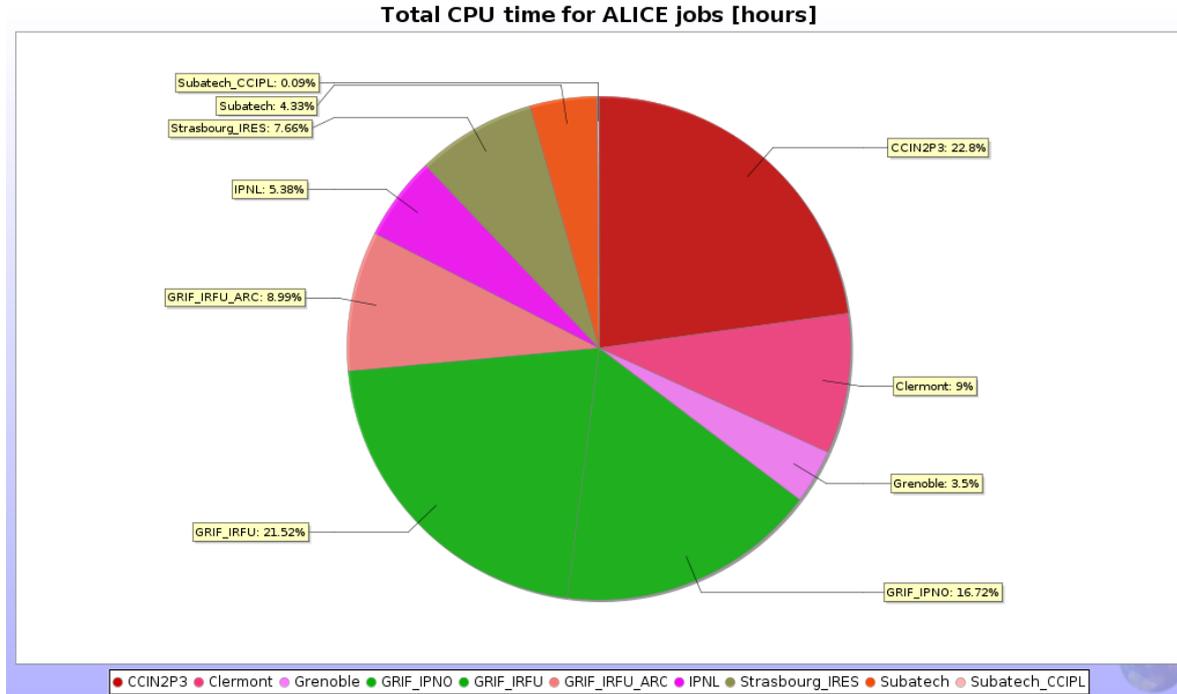


CPU time delivered



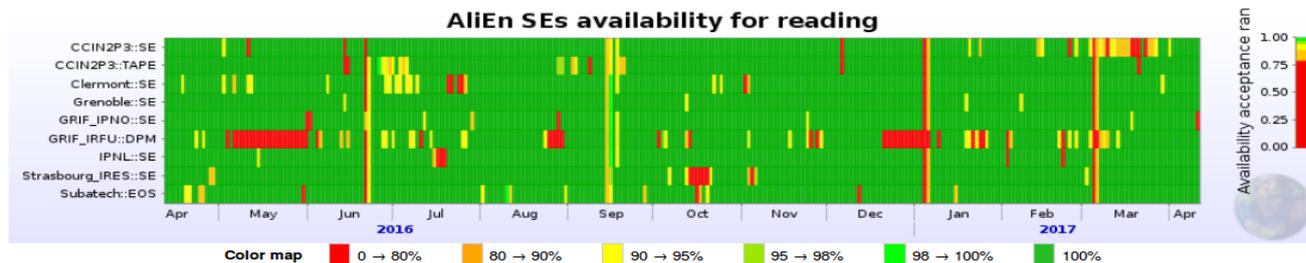
- ~ 60 Mhours delivered (steady)
 - 10 % of the 600 Mhours used by ALICE last year
 - Compared to 13 % last year

CPU breakdown in French sites



- T1 = 23 %
- T2 = 78 %
 - Whole GRIF : 40 % of total

Storage availability

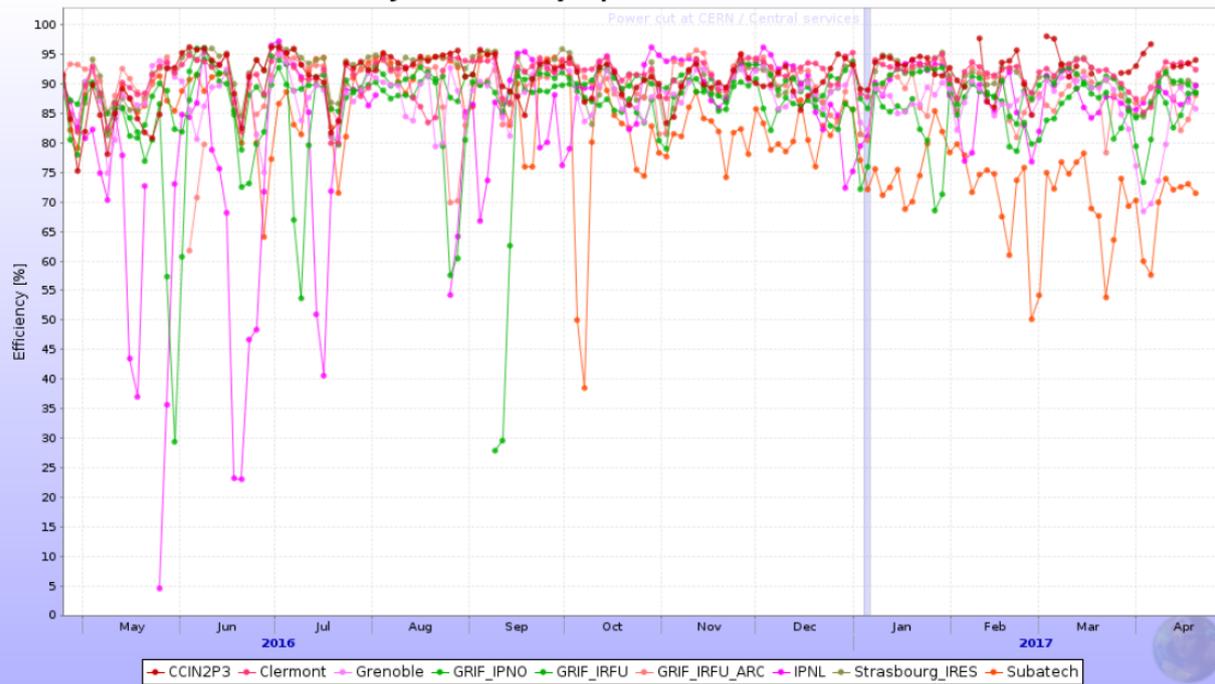


Statistics						
Link name	Data		Individual results of reading tests			Overall Availability
	Starts	Ends	Successful	Failed	Success ratio	
⚠ CCIN2P3::SE	11 Apr 2016 16:03	11 Apr 2017 16:03	4292	83	98.10%	98.10%
⚠ CCIN2P3::TAPE	11 Apr 2016 16:03	11 Apr 2017 16:03	4310	65	98.51%	98.51%
⚠ Clermont::SE	11 Apr 2016 16:06	11 Apr 2017 16:06	4307	67	98.47%	98.47%
⚠ Grenoble::SE	11 Apr 2016 16:09	11 Apr 2017 16:09	4345	29	99.34%	99.32%
⚠ GRIF_IPNO::SE	11 Apr 2016 16:10	11 Apr 2017 16:10	4326	48	98.90%	98.89%
⚠ GRIF_IRFU::DPM	11 Apr 2016 16:11	11 Apr 2017 16:10	3721	653	85.07%	85.01%
⚠ IPNL::SE	11 Apr 2016 16:12	11 Apr 2017 16:12	4303	70	98.40%	98.39%
⚠ Strasbourg_IRES::SE	11 Apr 2016 16:29	11 Apr 2017 16:29	4319	55	98.74%	98.75%
⚠ Subatech::EOS	11 Apr 2016 16:29	11 Apr 2017 16:30	4321	53	98.79%	98.79%

- Storage availability OK in most sites
- Problems @ GRIF-IRFU
 - Hard to investigate; DPM xrootd door solicited ; shared with FAX & AAA
 - No alerts sent to admin → fixed

CPU efficiency

Jobs efficiency (cpu time / wall time)



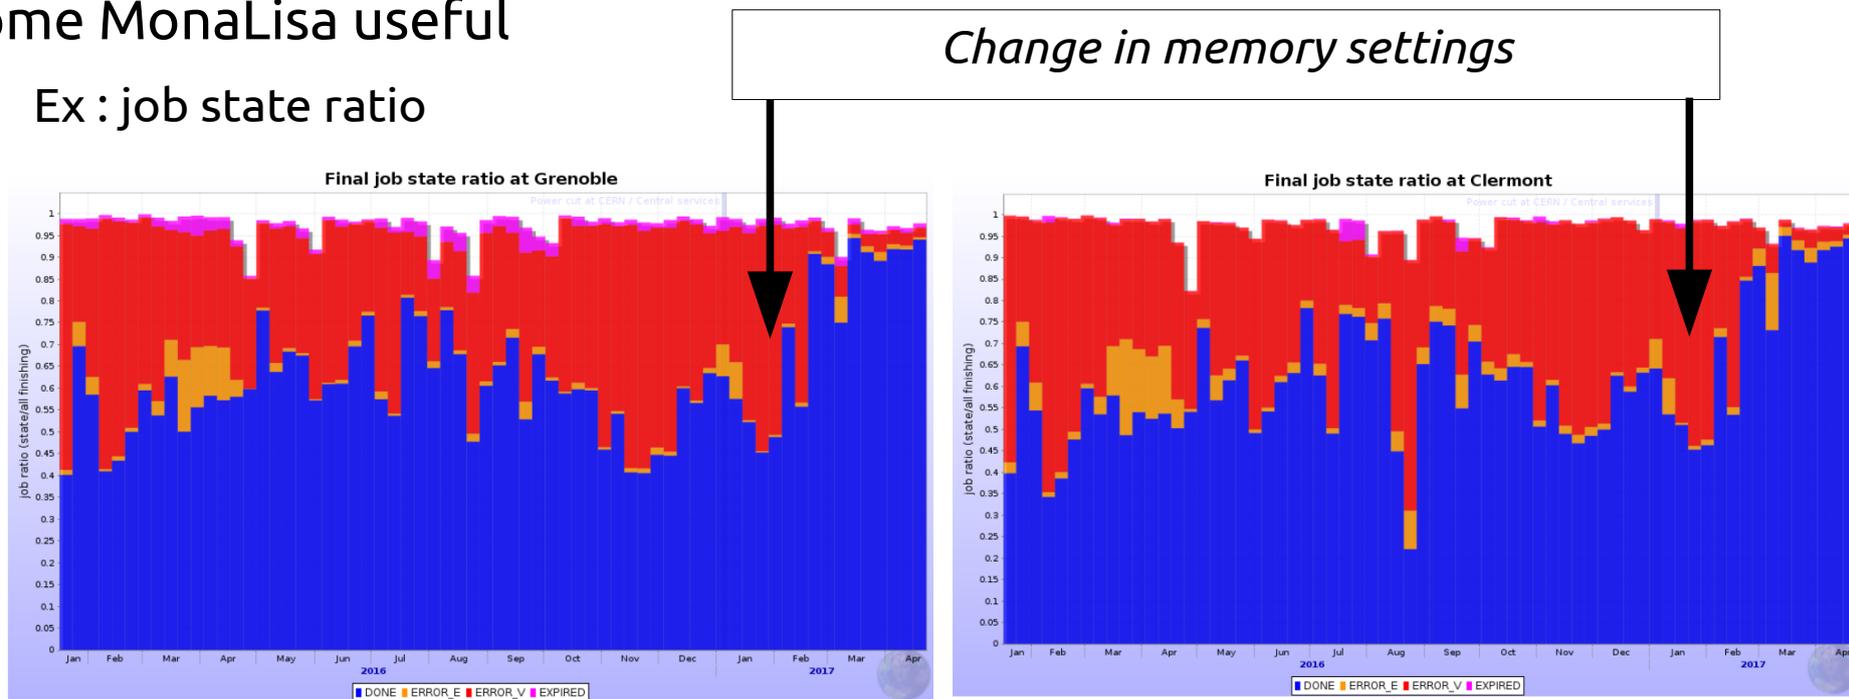
Jobs efficiency (cpu time / wall time)

	Series	Last value	Min	Avg	Max
1.	CCIN2P3	94.07	0	91.22	100
2.	Clermont	92.41	0	91.41	100
3.	Grenoble	85.78	0	87.61	100
4.	GRIF_IPNO	88.45	0	88.95	100
5.	GRIF_IRFU	88.56	0	84.24	100
6.	GRIF_IRFU_ARC	88.06	0	89.19	100
7.	IPNL	89.71	0	82.93	100
8.	Strasbourg_IRES	89.75	0	91.14	100
9.	Subatech	71.49	0	81.19	100
Total		87.59		87.54	

- Temporary problematic behaviours
- $85\% < \text{cpu_eff} < 95\%$ in most sites

Monitoring tools

- Some MonaLisa useful
 - Ex : job state ratio



- Can we have those visible in the main ML menu ? (and publicized)
 - Useful for site admins that don't have such tools locally

Growth 2016 → 2017

	T1	Σ T2
CPU	+60 %	+35 %
Disk	+18 %	+30 %
Tape	+93 %	

<CPU/Disk> ~ OK

- Budget 10 % higher than usual

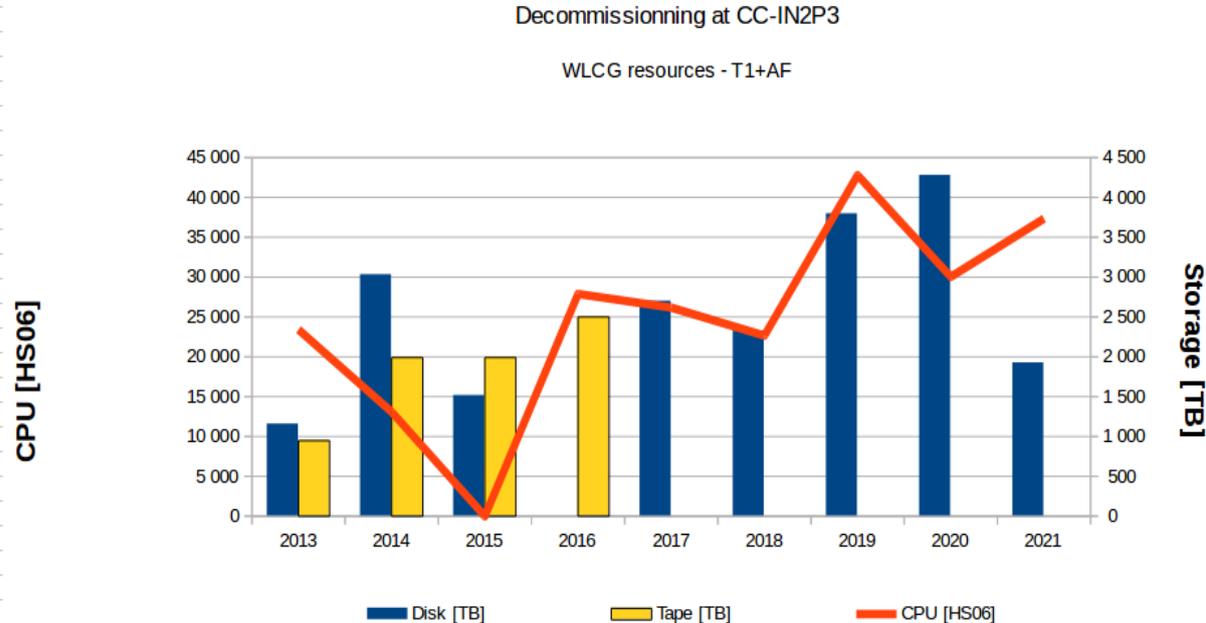
T1 = 12 % of ALICE requirements

T2 = 10 % of ALICE requirements

- Global yearly budget for « LCG France »
 - Stable over the last years
- Covers :
 - T1 : hardware replacement & growth (fully)
 - T2 : hardware replacement (partial)
- For T2s :
 - Budget comes from the labs
 - Sometimes contribution from universities
- Special extra budget to promote T1 growth in **2017**
 - Fully spent

Next years : previsions

- Replacement capacities :
- 2019 – 2020 will be tough
- No indication about future hardware prices
- CPU
 - Cost reduction ~ 10 % / year (*recently*)
- Disk
 - Cost reduction ~ 20 %/year (*recently*)



(Similar in T2s globally)

- Oracle stopping development of new drives & tapes
 - T10000 (we use that)
- Only one solution for the near future : LTO -8
 - Not before 2019
- Price evolution ?

- IPv6 migration
 - follows the plan (IPv6 working group)
 - T1 : dual stack storage OK
 - T2 : routing at sites ongoing
- Backbones for Run 3
 - CCIN2P3, GRIF, Strasbourg, Subatech : OK
 - Grenoble, Clermont, IPNL : may have to be upgraded
- LHC-ONE : OK
- Good work with RENATER (NREN)

- Operations smooth
- No big effort needed in average
- CERN team reactive

- EOS
 - Subatech & Strasbourg : migrated/ing
 - IPNL Lyon considering
 - → would like to move both VOs, possible ?
 - Other French sites
 - Not convinced : need & benefit ?

- **GRIF-IPNO :**

Replace Torque&Maui with HTCondor

- New vobox + cream to test : ~mid-May
- Parallel cluster to be opened for ALICE : work ongoing

- **GRIF-IRFU :**

- Connectivity upgrade 10 → 20 Gbps to be done 'soon'
- Full IPv6 dual stack
- 600 TiB storage (DPM+Xrd)
- ~3k threads used by ALICE

- Joint application for state/region funding successful
 - 300 k€ for Subatech
 - 150 k€ spent
 - 1 PB storage last October
 - 50 k€ for CPU in CCIPL
 - Remaining 150 k€ in 2019-2020
- All storage under EOS
 - with dual-managers
- Working on making use of our share in CCIPL
 - following the proof of concept already working at Subatech

- New CCIPL cluster (last October)
 - SLURM, Docker-ready, Intel-Omnipath
 - ALICE share is ~250 / 1600 cores
- Opened to users (March)
- Vobox installed and connected to Subatech Network (March)

- To be done
 - cvmfs on WNs
 - network config (gateway via vobox 10Gbits/s)
 - WLCG accounting ?
 - MoU (including security) ?

- 400 Gbps backbone
- CentOS 7
 - Several workers migrated
 - The rest will move smoothly
- Warranty extension of data servers : some experience now
 - Saved money
 - Spent time changing disks
- Plan for AFS decommissioning : 2019

- Not increase the number of services at sites
 - Can/should we reduce it ?
- Single entry point (storage & compute)
 - But keep resources distributed (like now)
 - Location-aware software

- Smooth operations
- Good performance of sites
 - GRIF-IRFU storage should get better now
- EOS
 - Some sites migrated/ing or considering
 - Some sites not-very-much convinced (yet)
- Run 3 : a forum ?