



Centre de Calcul de l'Institut National de Physique Nucléaire et de Physique des Particules

# France status report

ALICE T1/2 Workshop

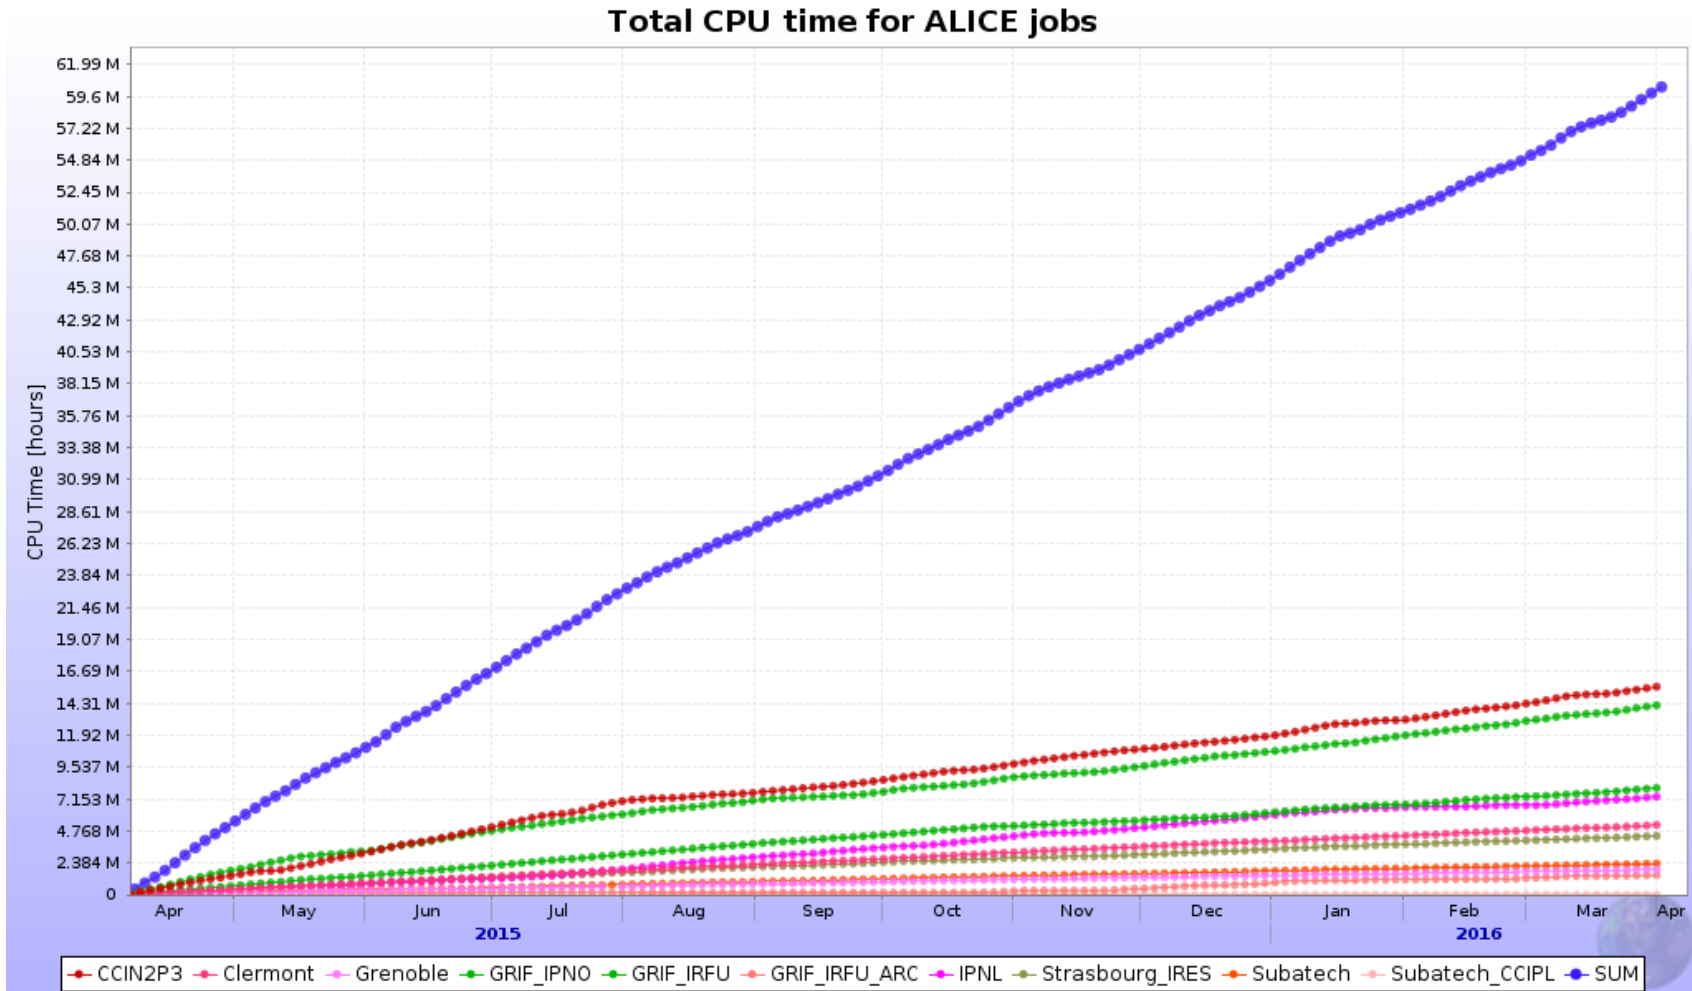
Bergen, Apr. 2016

Renaud Vernet

- Overview
- News from sites
- Questions & suggestions

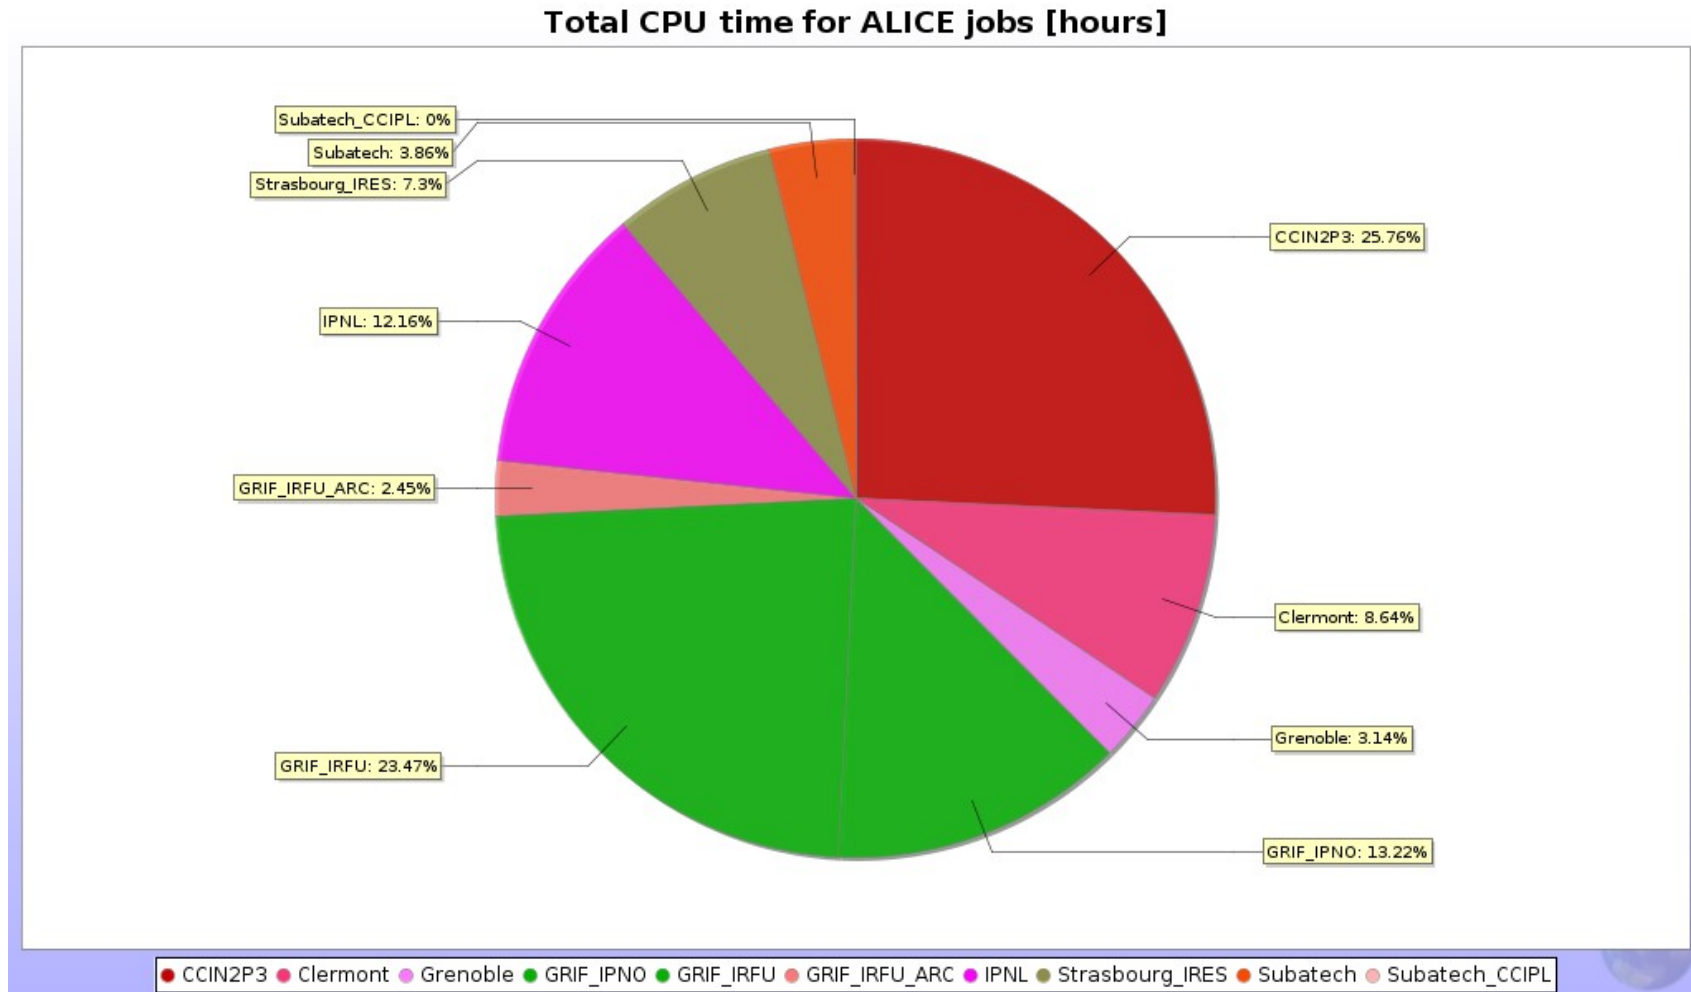
- 6 cities
  - Paris, Lyon, Nantes, Strasbourg, Grenoble, Clermont






- 60 Mhours delivered
  - 13 % of the 450 Mhours used by ALICE last year





- T1 = 25 %
- T2 = 75 %
  - Whole GRIF : 40 % of total

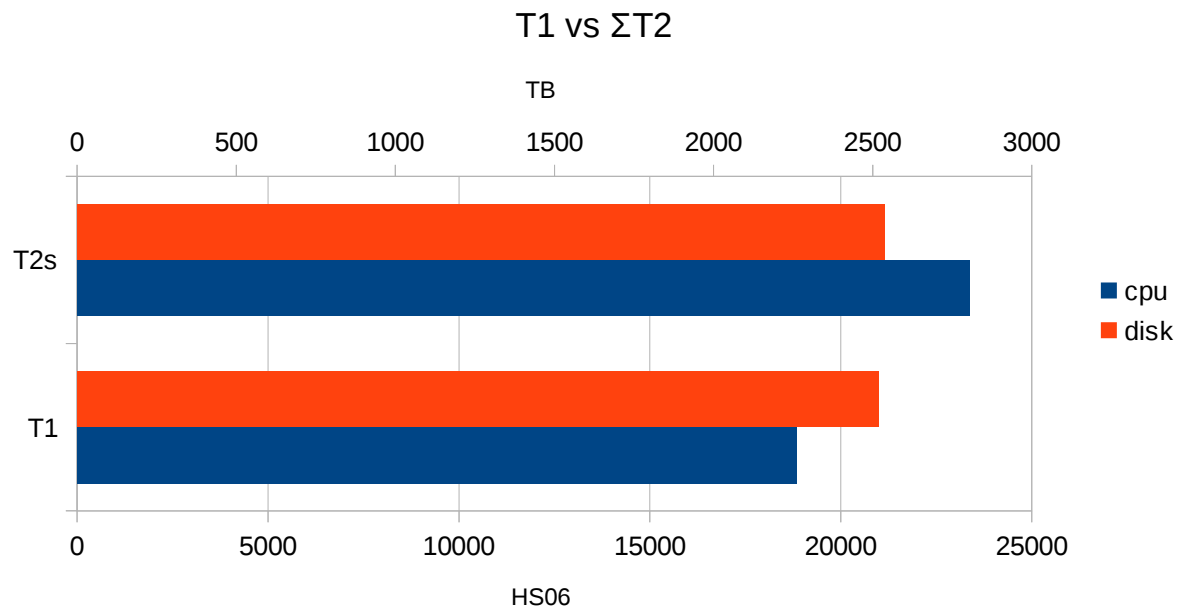
## AliEn SEs availability for reading



Statistics						
Link name	Data		Individual results of reading tests			Overall
	Starts	Ends	Successful	Failed	Success ratio	Availability
CCIN2P3::SE	06 Apr 2015 16:03	05 Apr 2016 14:03	4294	90	97.95%	97.99%
CCIN2P3::TAPE	06 Apr 2015 16:04	05 Apr 2016 14:03	4315	69	98.43%	98.49%
Clermont::SE	06 Apr 2015 16:06	05 Apr 2016 14:06	4350	34	99.22%	99.28%
Grenoble::SE	06 Apr 2015 16:10	05 Apr 2016 14:09	4367	17	99.61%	99.65%
GRIF_IPNO::SE	06 Apr 2015 16:10	05 Apr 2016 14:10	4363	21	99.52%	99.56%
GRIF_IRFU::DPM	06 Apr 2015 16:11	05 Apr 2016 14:10	4247	137	96.88%	96.89%
IPNL::SE	06 Apr 2015 16:13	05 Apr 2016 14:12	4363	21	99.52%	99.58%
Strasbourg_IRES::SE	06 Apr 2015 16:30	05 Apr 2016 14:28	4331	54	98.77%	98.81%
Subatech::EOS	06 Apr 2015 16:30	05 Apr 2016 14:30	4209	177	95.96%	96.02%

- Storage Availability > 96 % in all sites
  - Over the last year
  - Target OK

Growth 2015 → 2016			% requirements	
	T1	Σ T2	T1	Σ T2
<b>CPU</b>	<b>28%</b>	<b>19%</b>	<b>12%</b>	<b>10%</b>
<b>Disk</b>	<b>16%</b>	<b>51%</b>		
<b>Tape</b>	<b>32%</b>			



Global contribution ~ 10 %  
 ~ 50 % T1 / 50 % T2s

## LAN traffic

Traffic OUT

	Series	Last value	Min	Avg	Max	Total
1.	■ CCIN2P3	394.4 MB/s	0 B/s	525.2 MB/s	8.084 GB/s	15.39 PB
2.	■ Grenoble	17.02 MB/s	0 B/s	25.42 MB/s	579.7 MB/s	762.9 TB
3.	■ GRIF_IPNO	260.3 MB/s	0 B/s	109.4 MB/s	3.322 GB/s	3.205 PB
4.	■ IPNL	27.2 KB/s	0 B/s	0.175 MB/s	117.9 MB/s	5.245 TB
5.	■ Strasbourg_IRES	0.199 MB/s	0 B/s	1.918 MB/s	886.8 MB/s	57.56 TB
6.	■ Subatech	133.8 MB/s	0 B/s	58.42 MB/s	1.166 GB/s	1.712 PB
	<b>Total</b>	<b>805.7 MB/s</b>		<b>720.5 MB/s</b>		<b>21.11 PB</b>

## WAN traffic

Traffic OUT

	Series	Last value	Min	Avg	Max	Total
1.	■ CCIN2P3	14.14 MB/s	0 B/s	42.41 MB/s	6.048 GB/s	1.243 PB
2.	■ Grenoble	0.535 MB/s	0 B/s	1.717 MB/s	406.4 MB/s	51.54 TB
3.	■ GRIF_IPNO	0 B/s	0 B/s	0.931 MB/s	564.8 MB/s	27.95 TB
4.	■ IPNL	13.27 KB/s	0 B/s	0.153 MB/s	25.46 MB/s	4.603 TB
5.	■ Strasbourg_IRES	47.24 KB/s	0 B/s	0.224 MB/s	49.88 MB/s	6.71 TB
6.	■ Subatech	6.211 MB/s	0 B/s	3.283 MB/s	338 MB/s	98.54 TB
	<b>Total</b>	<b>20.95 MB/s</b>		<b>48.72 MB/s</b>		<b>1.428 PB</b>

- WAN is kept < 10 % LAN
  - Over the last year



- RENATER
  - Starting to upgrade backbone to 100 Gbps
- LHCOPN link CERN ↔ CC-IN2P3
  - 10 Gbps → 2x10 Gbps
- IPv6
  - Doesn't seem as urgent as before
    - Less pressure from CERN
  - Some services have it nevertheless (dCache)
  - In the todo list though

- CentOS 7
  - Waiting for green light from WLCG & experiments
- Pledges delivered
- Warranty extension of data servers
  - 5 → 7 years
  - Wait and see
  - Possible savings
- Puppetization of storage services ongoing
  - A few problems...

- Pledges
  - Already there
- Operations
  - Nothing special to report
- Memory consumption
  - ALICE jobs colliding with ATLAS analysis jobs
  - Needed to raise  $> 4$  GB/slot

- Pledges
  - Some delay, should be installed ~ end April
- Operations
  - OK, manpower just enough
  - Several data servers replaced
    - Some effort with rsync
- Flat profile of # jobs with time

- Pledges
  - OK
- Nothing special on operations, smooth
- High-mem jobs killing WNs
  - Much less frequent than before (still happens sometimes)
- New HPC farm (project)
  - 600 potential additional slots
    - with 3-4 GB/slot RAM
  - ALICE would have ~0 "priority"
    - ALICE jobs run when free, ALICE jobs killed when busy (local activities)
    - Or wait some time before killing (how, how long before, which signal?)
    - Any interest ?

- Pledges
  - Delays
    - Wrong setup of CPU servers delivered (Dell)
- Manpower
  - Lacks
  - Vobox operations done by Latchezar & Maarten essentially
  - Staying up-to-date (system-wise) not easy
- Xrootd drain last month
  - Huge number of very small files → 2 weeks draining
- Smooth workflow otherwise



- Pledges
  - N/A (Tier-3)
- RSS
  - ~2.8 GB per **aliproduct** job (last month)
  - Some workers lack memory → limit # slots ?
- ALICE uses efficiently free resources when available

- Pledges
  - No problem
- Smooth utilization of resources
- Operations easy
  
- HEPiX benchmark WG
  - Jerome Pansanel

- Pledges
  - Regional funds, take time
  - Disk delivery expected later this year (+400 TB)
- Smooth workflow
  
- New CCIPL resources
  - Pledged ? Not pledged ?

- Any interest to use (at sites) ? :
  - Cloud
  - Multicore
  - GPU and cie.
- EOS
  - Nice to have a "collaboration feedback" (small sites)
  - Should that be preferred to staying with native xrootd ?
- ML tests included in availability metric ?

- Communication and support model
  - Should we continue with emails ?
  - ggus tickets?
- Improve documentation
- Improve deployment model
- I/O tuning on data servers
  - Find the best kernel params for I/O
  - Any suggestions from ALICE (like with TCP/IP in the past) ?
  - Some web page where we can share experience ?

- **ALICE::CCIN2P3::TAPE**
  - Move to xrootd 4.2.3 uneasy
  - Unable to do 3rd party copy
    - Under investigation with Andy H.
- **ALICE::CCIN2P3::SE**
  - Today : mix of old and new servers (3.3.4, 4.2.3)
- Introduction of balanced common redirector
  - For both::TAPE (and::SE soon)
- Now puppetizing both clusters



- Puppetization
  - Need RPMs for monitoring
    - MLSensor
    - ApMon
  
- Monitoring
  - Not easy to understand what does what
    - → Hard to know where problems come from
  - A single monitoring tool would help

- Big xrd.sh that does everything
  - → split into smaller pieces ?
  - → make configurable servMon.sh
  
- → Simplify the whole thing, ideally
  - 1 set of RPMs
  - 1 set of config files

- Good French contribution to ALICE computing
- Smooth operations
  - memory problems sometimes
- Good support from CERN
  
- Suggestions to improve
  - Sharing of information
  - Storage service deployment