Operation in France

2013

ALICE T1/2 workshop
Tsukuba
05/03/2014

Renaud Vernet (CCIN2P3)

筑波大学
University of Tsukuba

LCG France

Outline

- LCG France network
- Resource status
- Sites reports

French sites

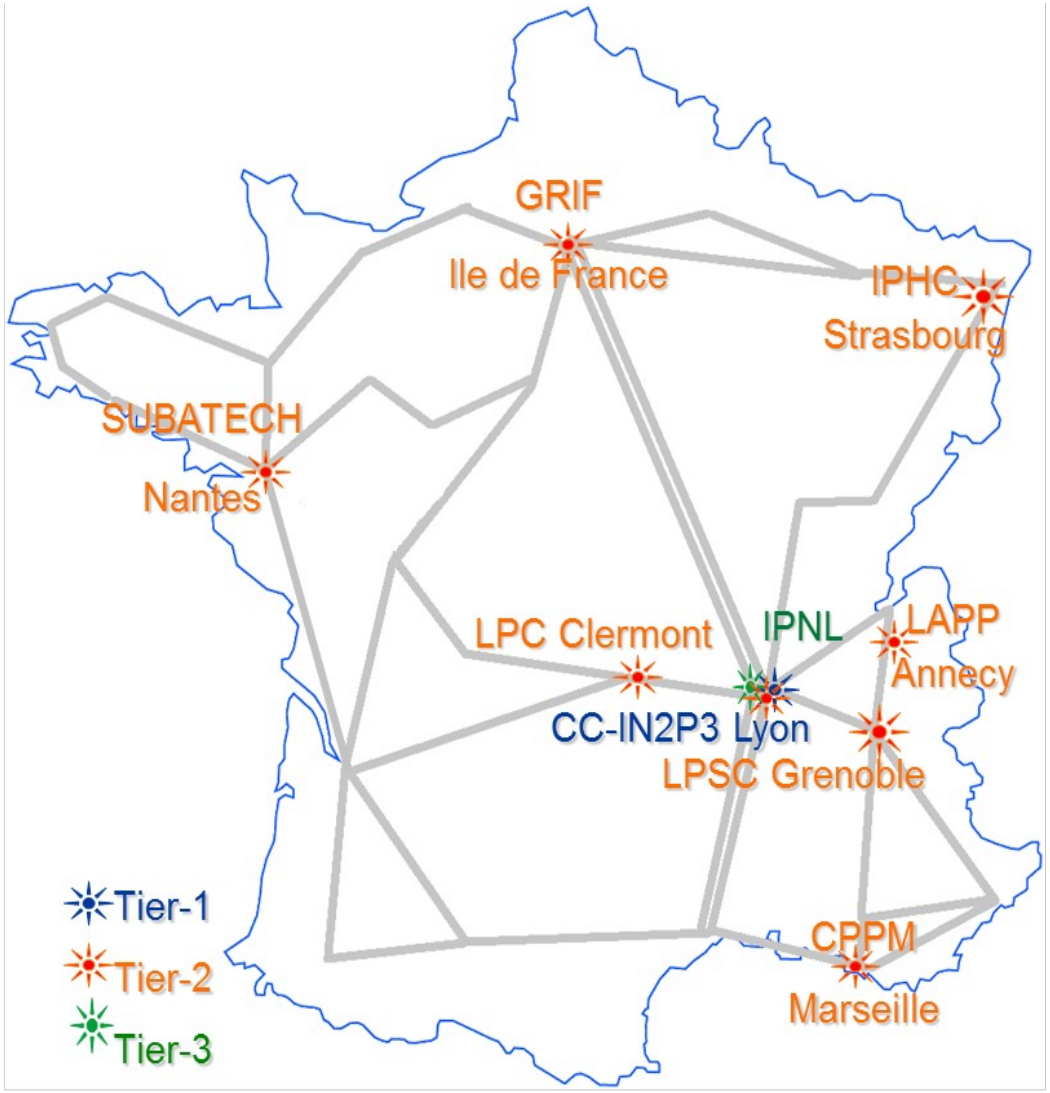

Role	Site	ALICE	ATLAS	CMS	LHCb
Tier-1	IN2P3-CC	✓	✓	✓	✓
Tier-2	IN2P3-CC-T2 (AF)	✓	✓	✓	✓
	IN2P3-CPPM		✓		✓
	GRIF	✓	✓	✓	✓
	IN2P3-LPC	✓	✓		✓
	IN2P3-IPHC	✓		✓	
	IN2P3-LAPP		✓		✓
	IN2P3-LPSC	✓	✓		
	IN2P3-SUBATECH	✓			
Tier-3	IN2P3-IPNL	✓		✓	

LHC connectivity

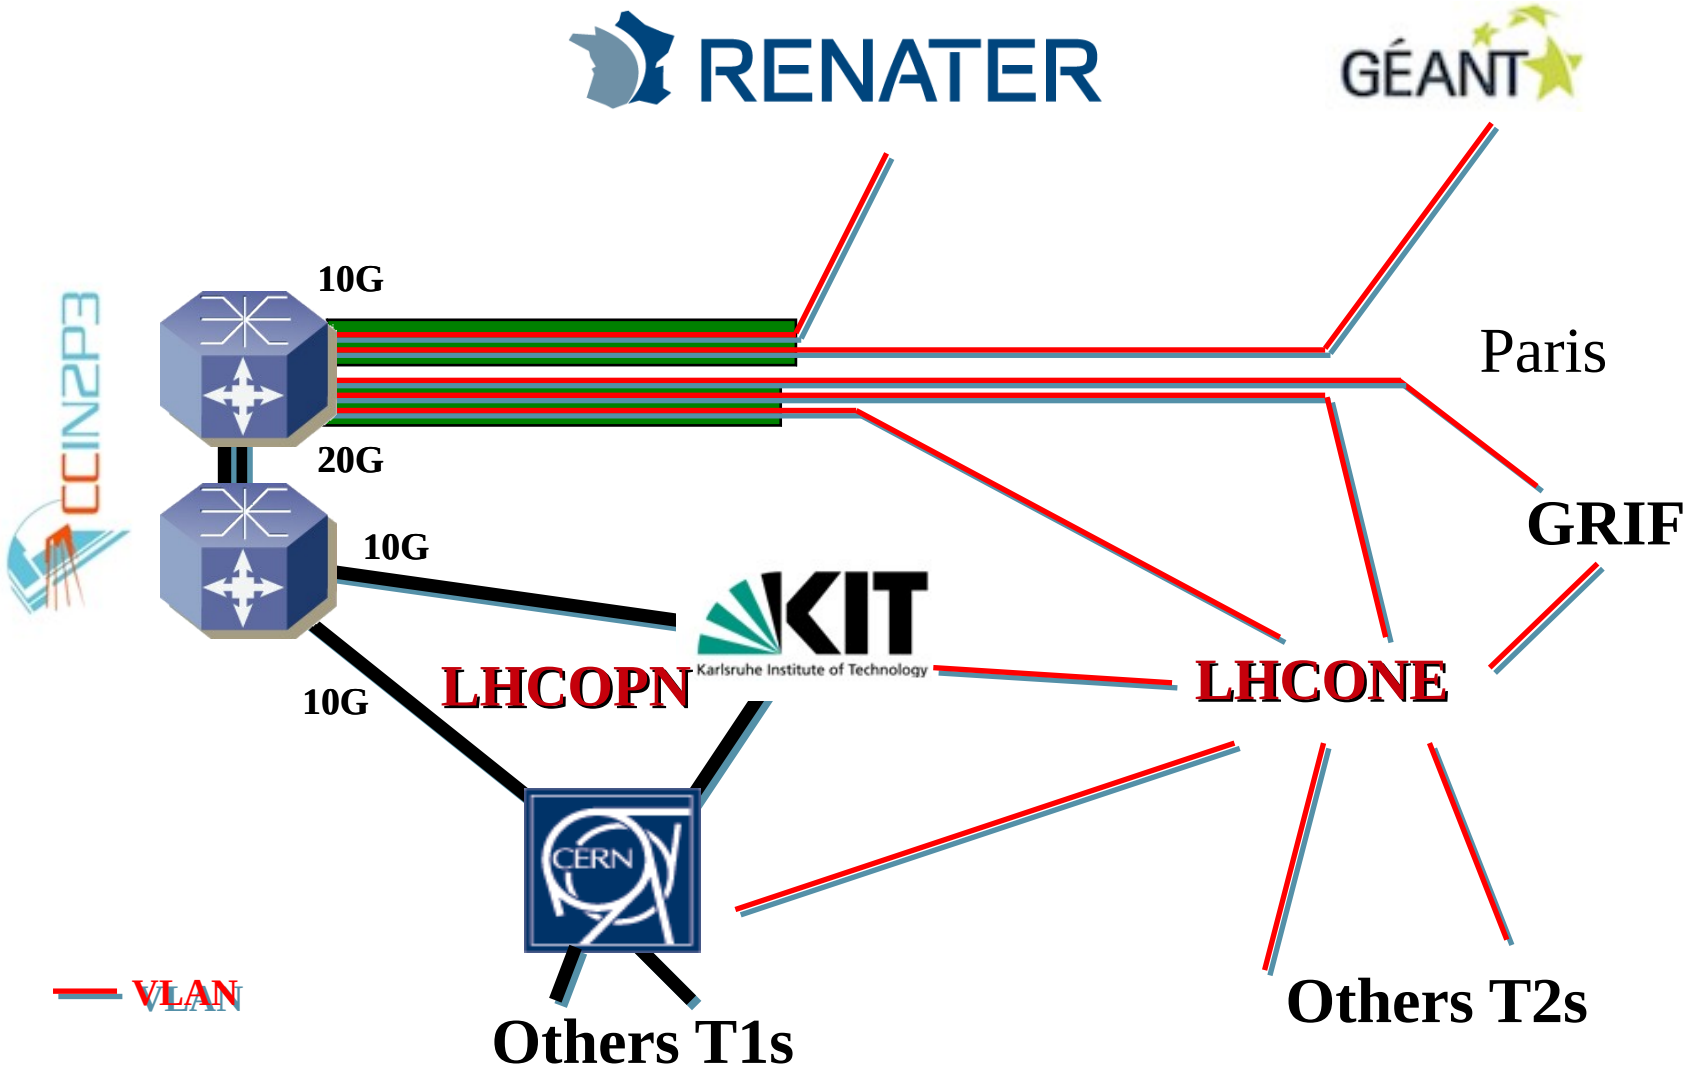

- All the French sites have a 10 Gbps connection to LHCONE
  - Depending on the size :
    - Shared with generic IP, L2VPN or dedicated lambda
- LHC-ONE is used more and more
- CCIN2P3 and GRIF have a 20 Gbps connection to LHCONE

Bandwidth (country view)

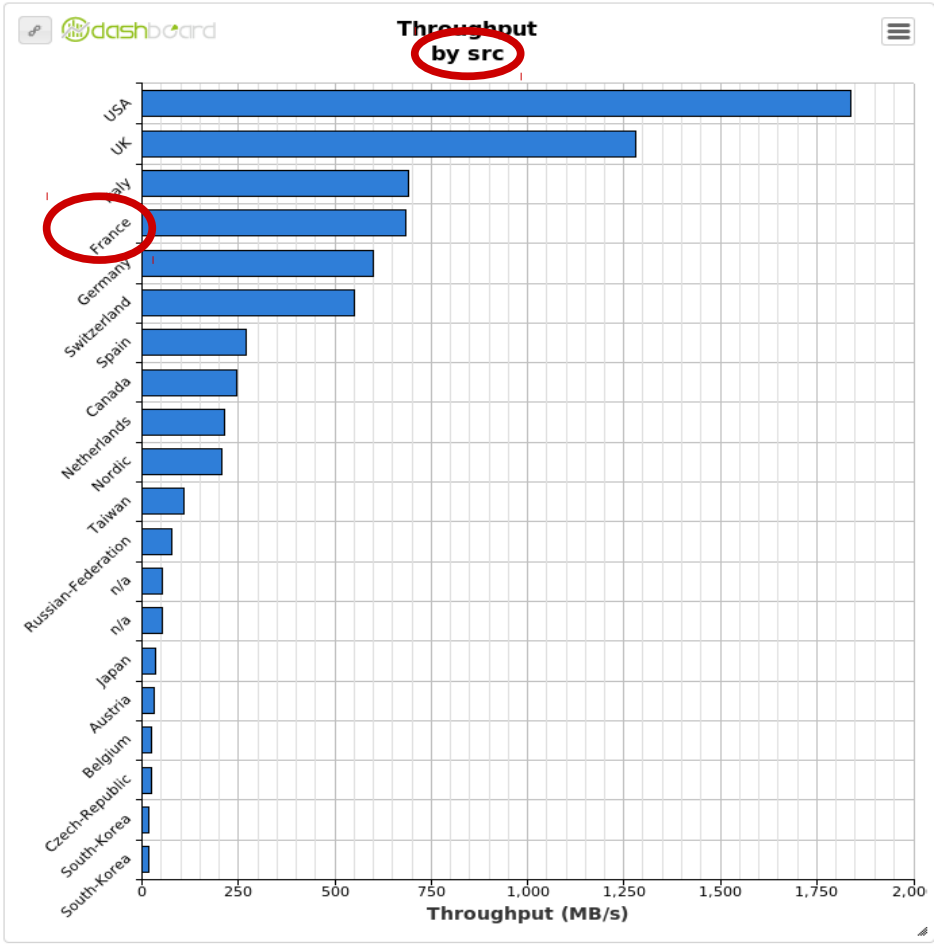

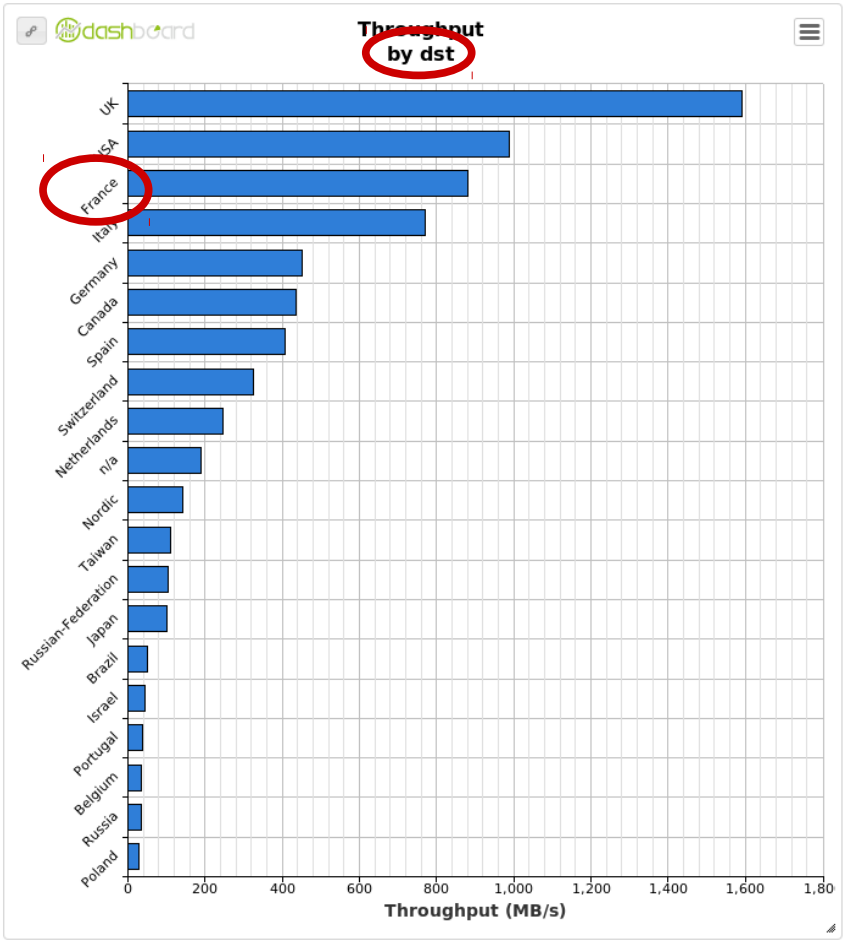

Does NOT include ALICE transfers (xrootd) but
Gives a good indication of where we stand

French contribution to LHC computing

Normalised CPU time (HEP-SPEC06)

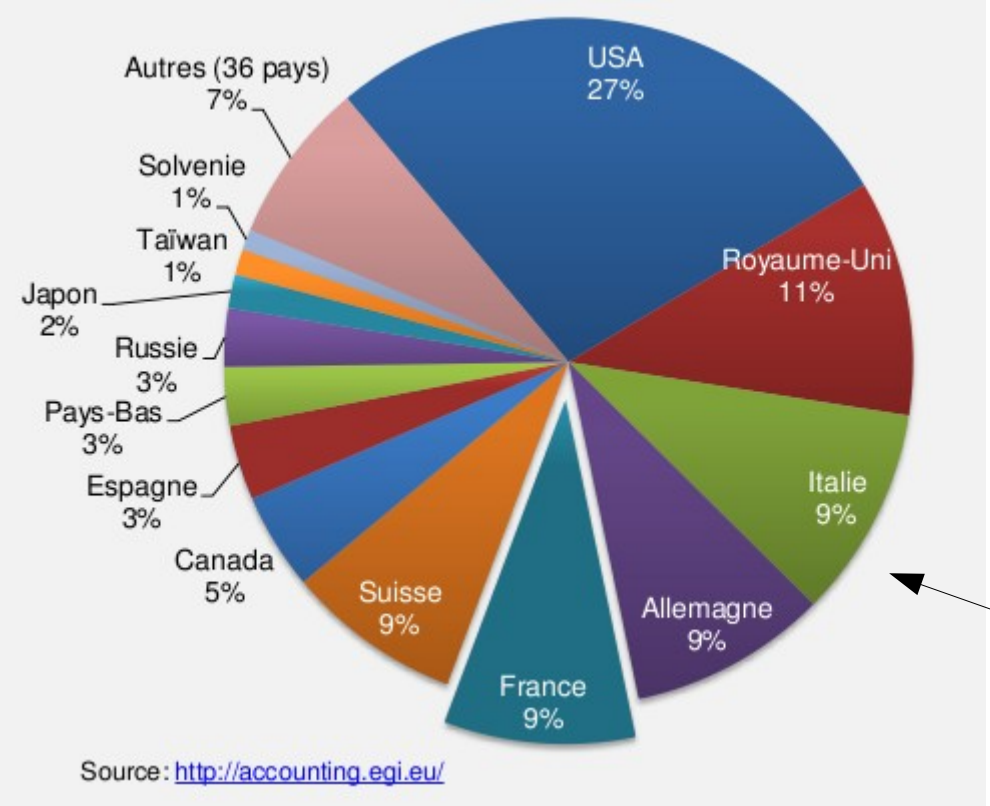

Evolution with time

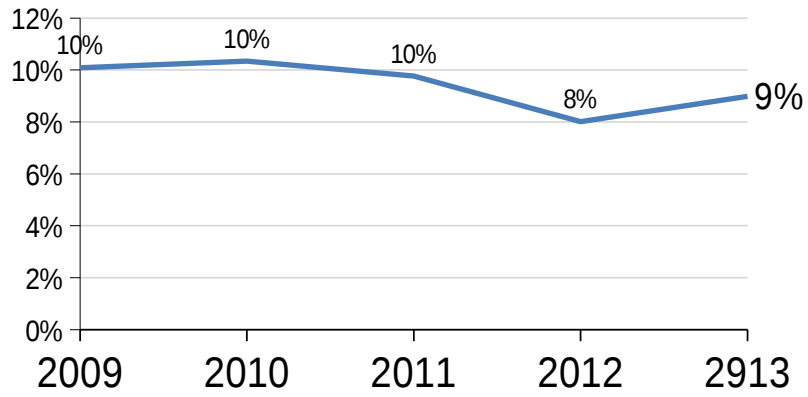

4 LHC V0s

T1 vs T2,3 resource usage

Delivered CPU

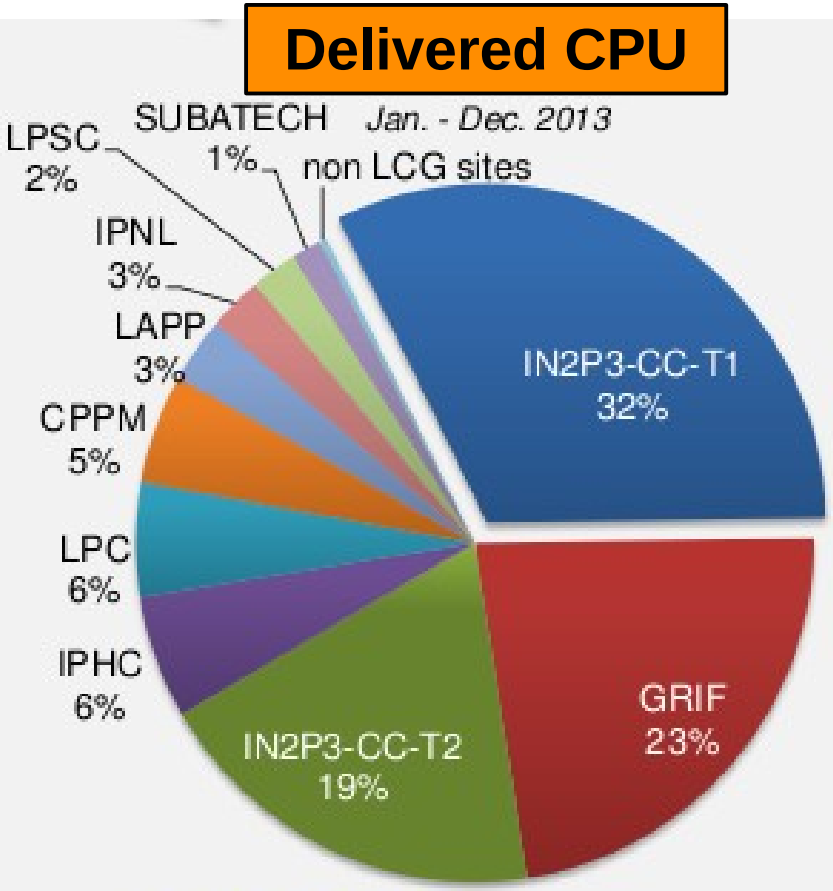

Allocated disk

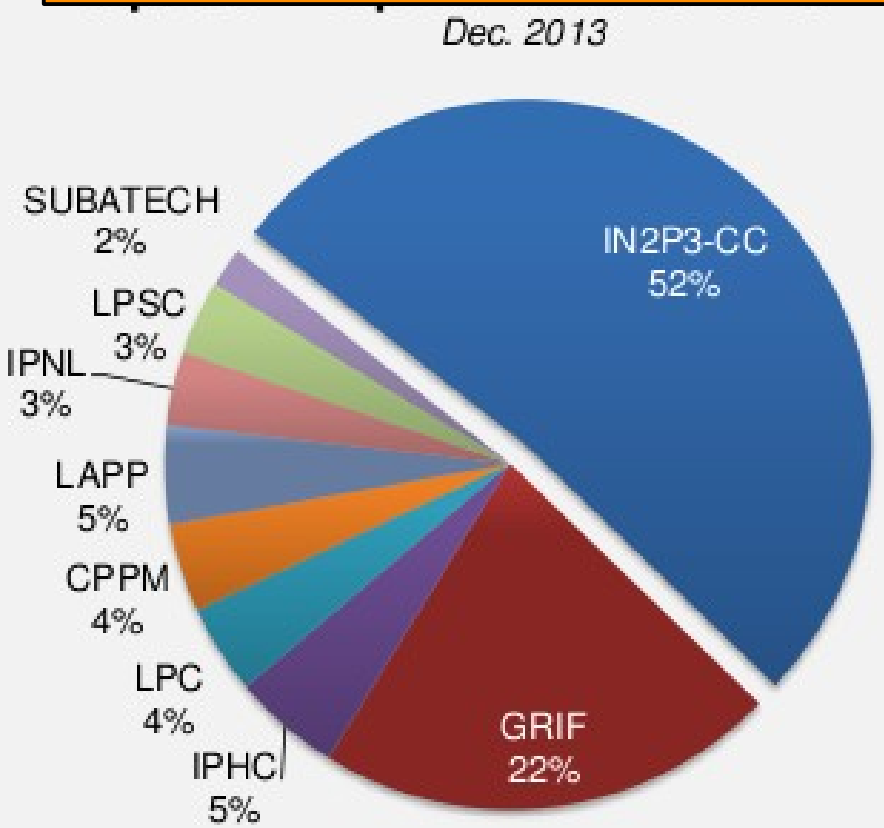

Sources : http://accounting.egi.eu/
http://gstat.egi.eu/gstat/summary/EGI_NGI/NGI_FRANCE/

4 LHC V0s

LHC computing resources in France

- End 2013 : LHC computing in France (all LHC VOs)
  - Pledged (T1+T2) + Non-pledged (T3)
- CPU : ~200 kHS06, Disk ~ 19 Po, Tape ~ 10 Po
  - T1 resources ~ T2+T3 resources

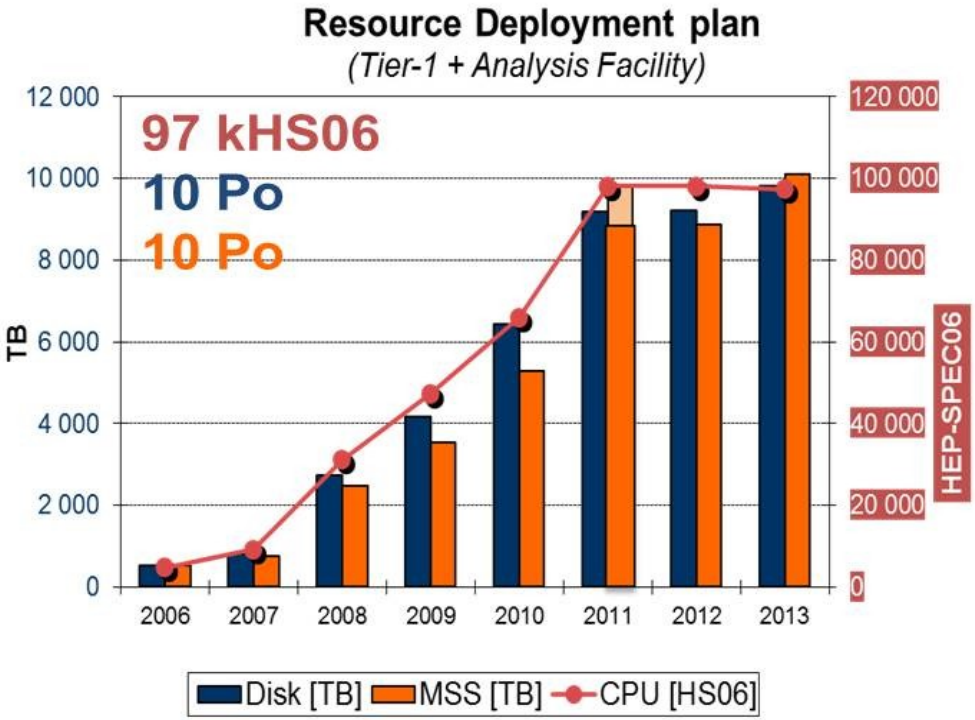

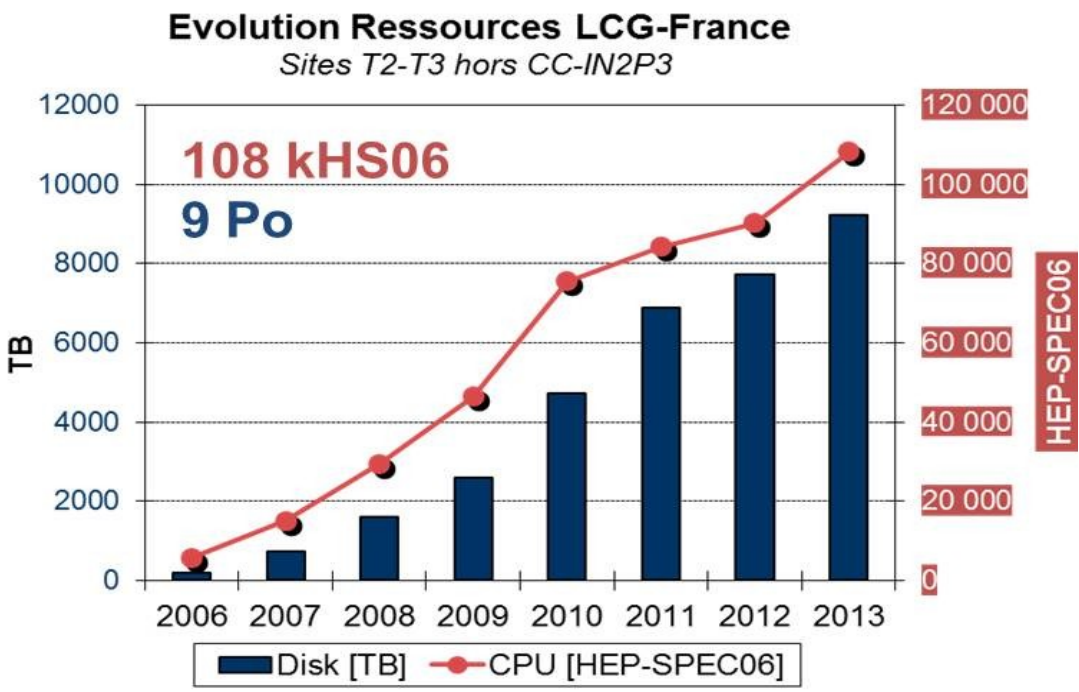

ALICE pledged resources in France

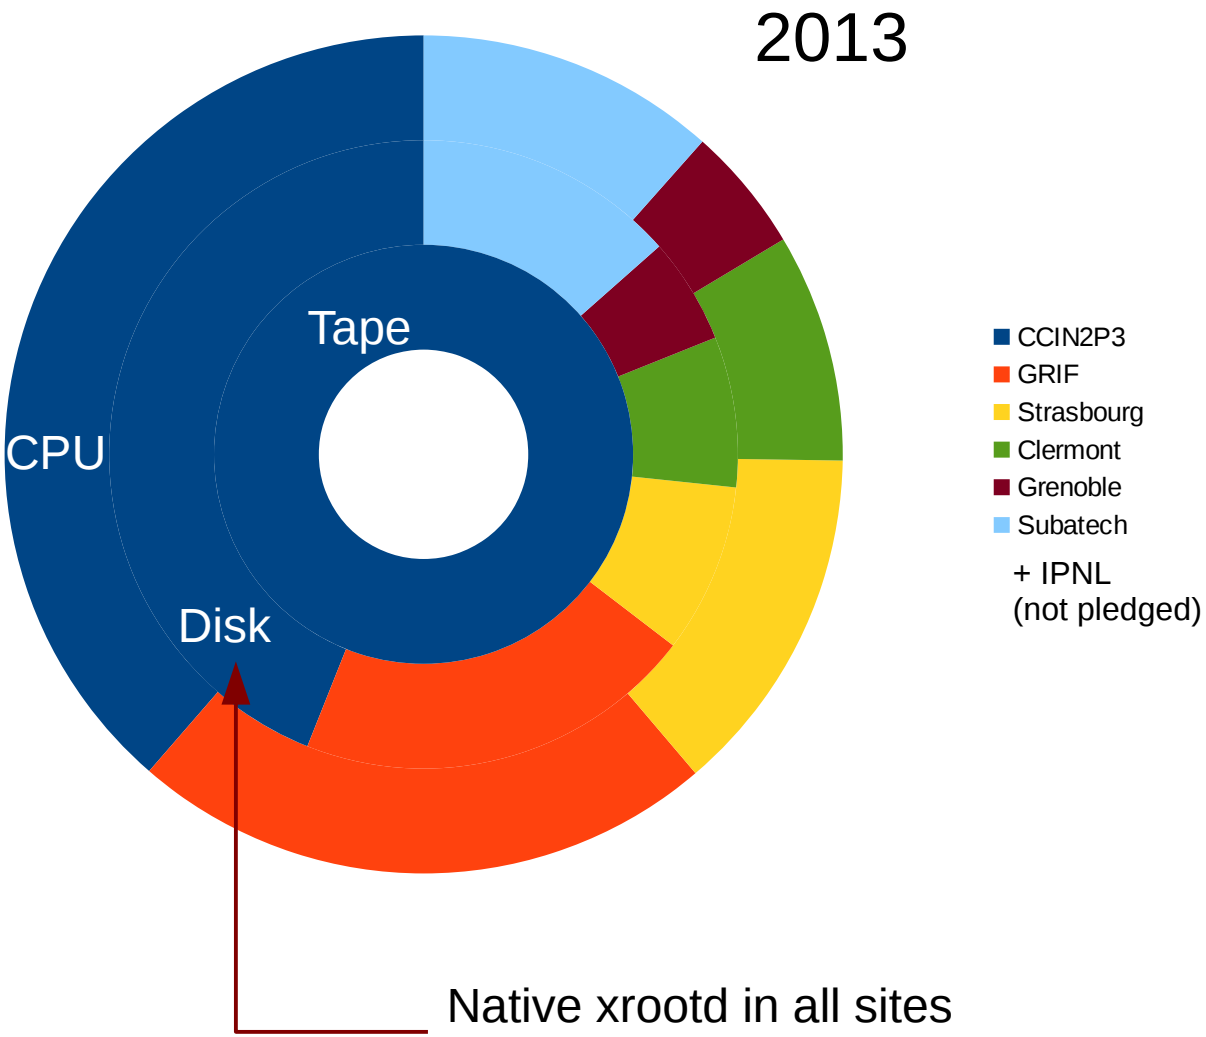

Native xrootd in all sites
but GRIF-IRFU

ALICE 2014 T1 resources forecast:
+15 % CPU , + 25 % disk , +0 % tapes

ALICE 2014 T2 resources forecast:
+15 % CPU , + 2 % disk

Year-by-year budget
Good budget in 2014
No idea about 2015

Availability

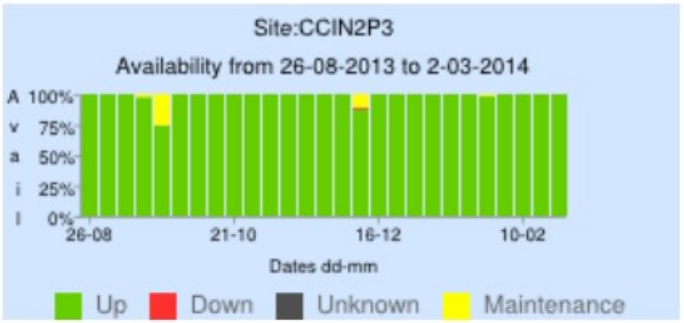

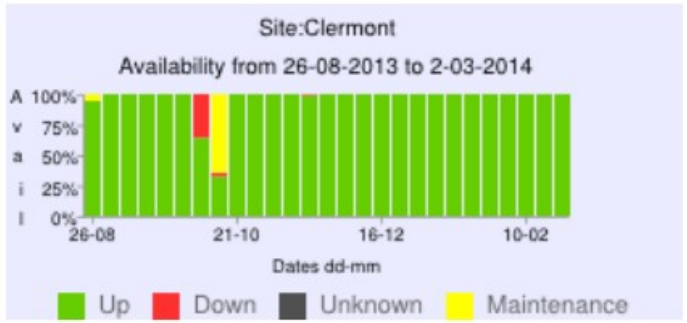

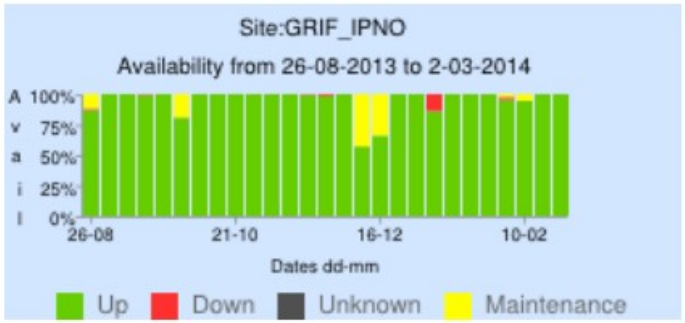

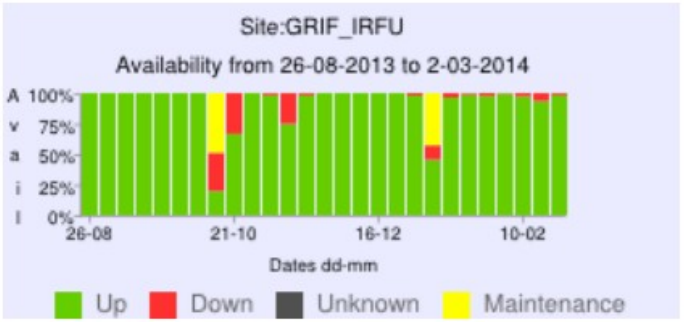

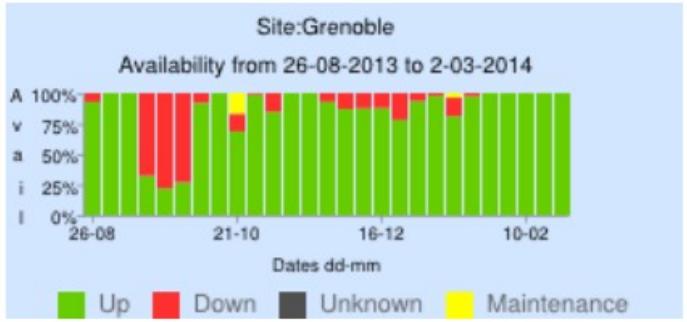

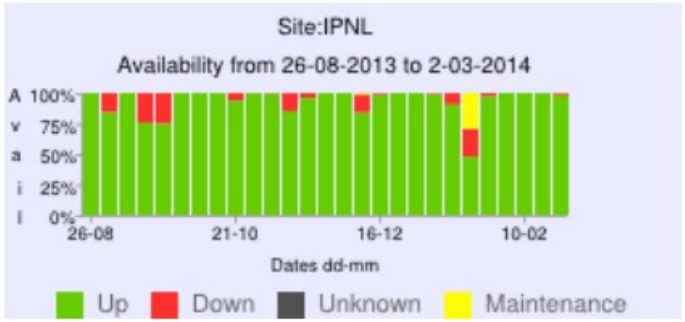

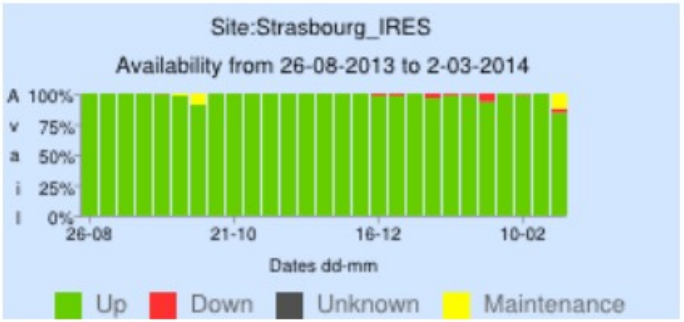

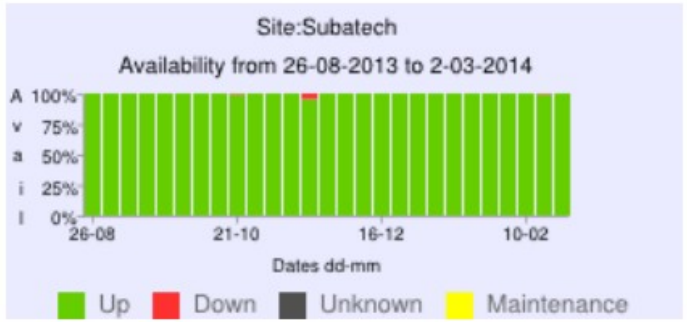

Reliability

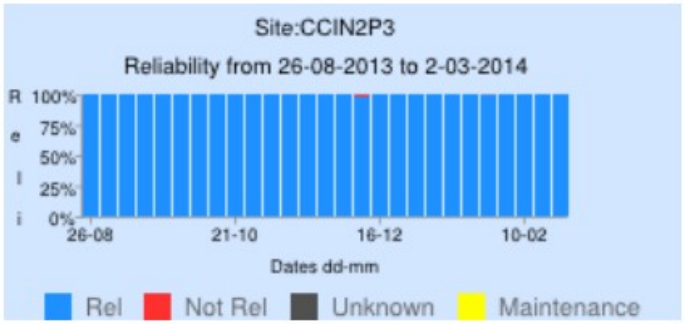

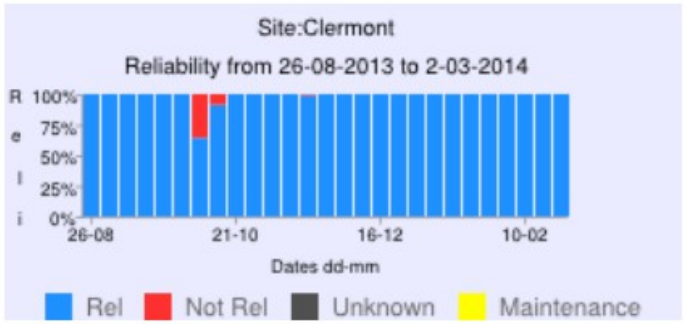

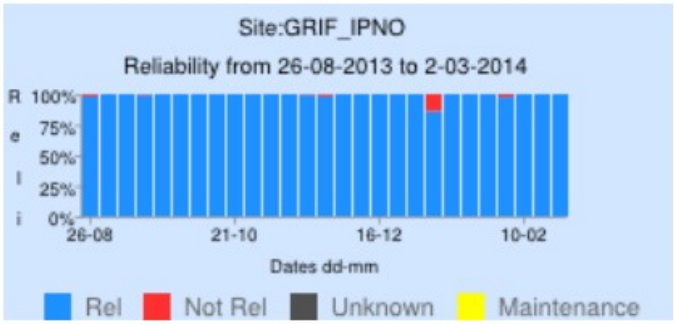

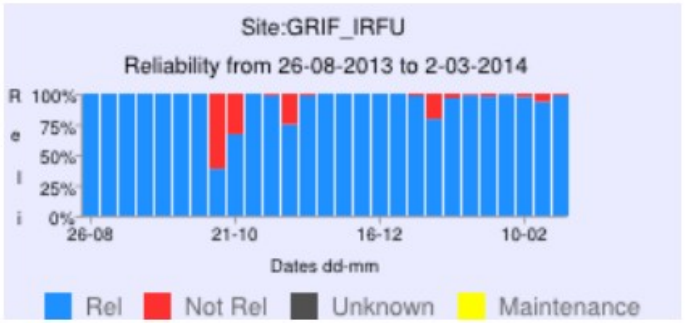

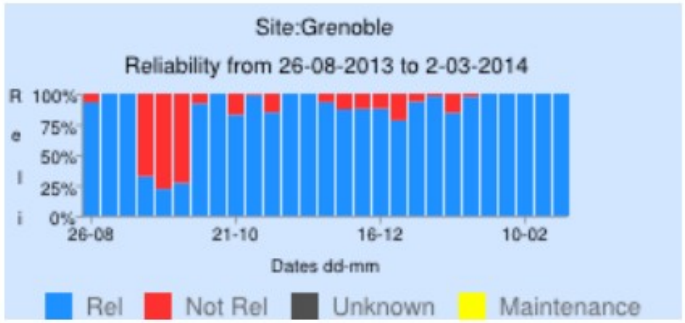

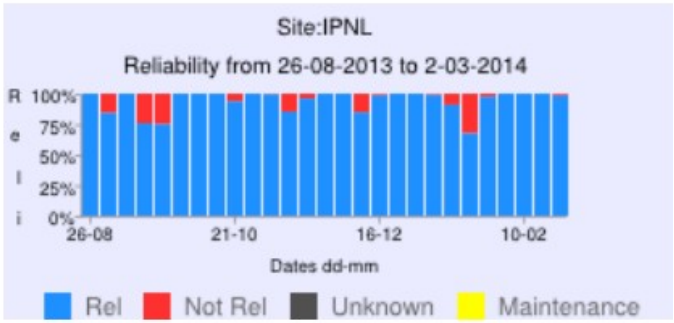

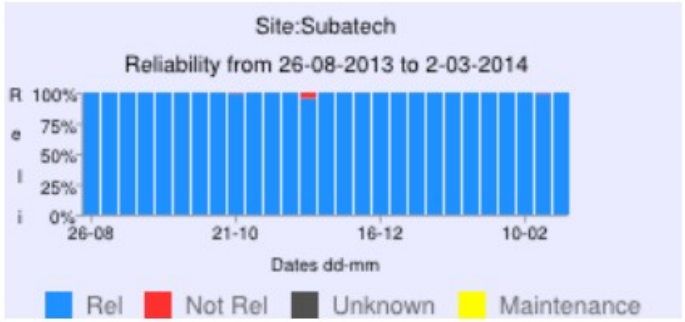

CPU efficiency

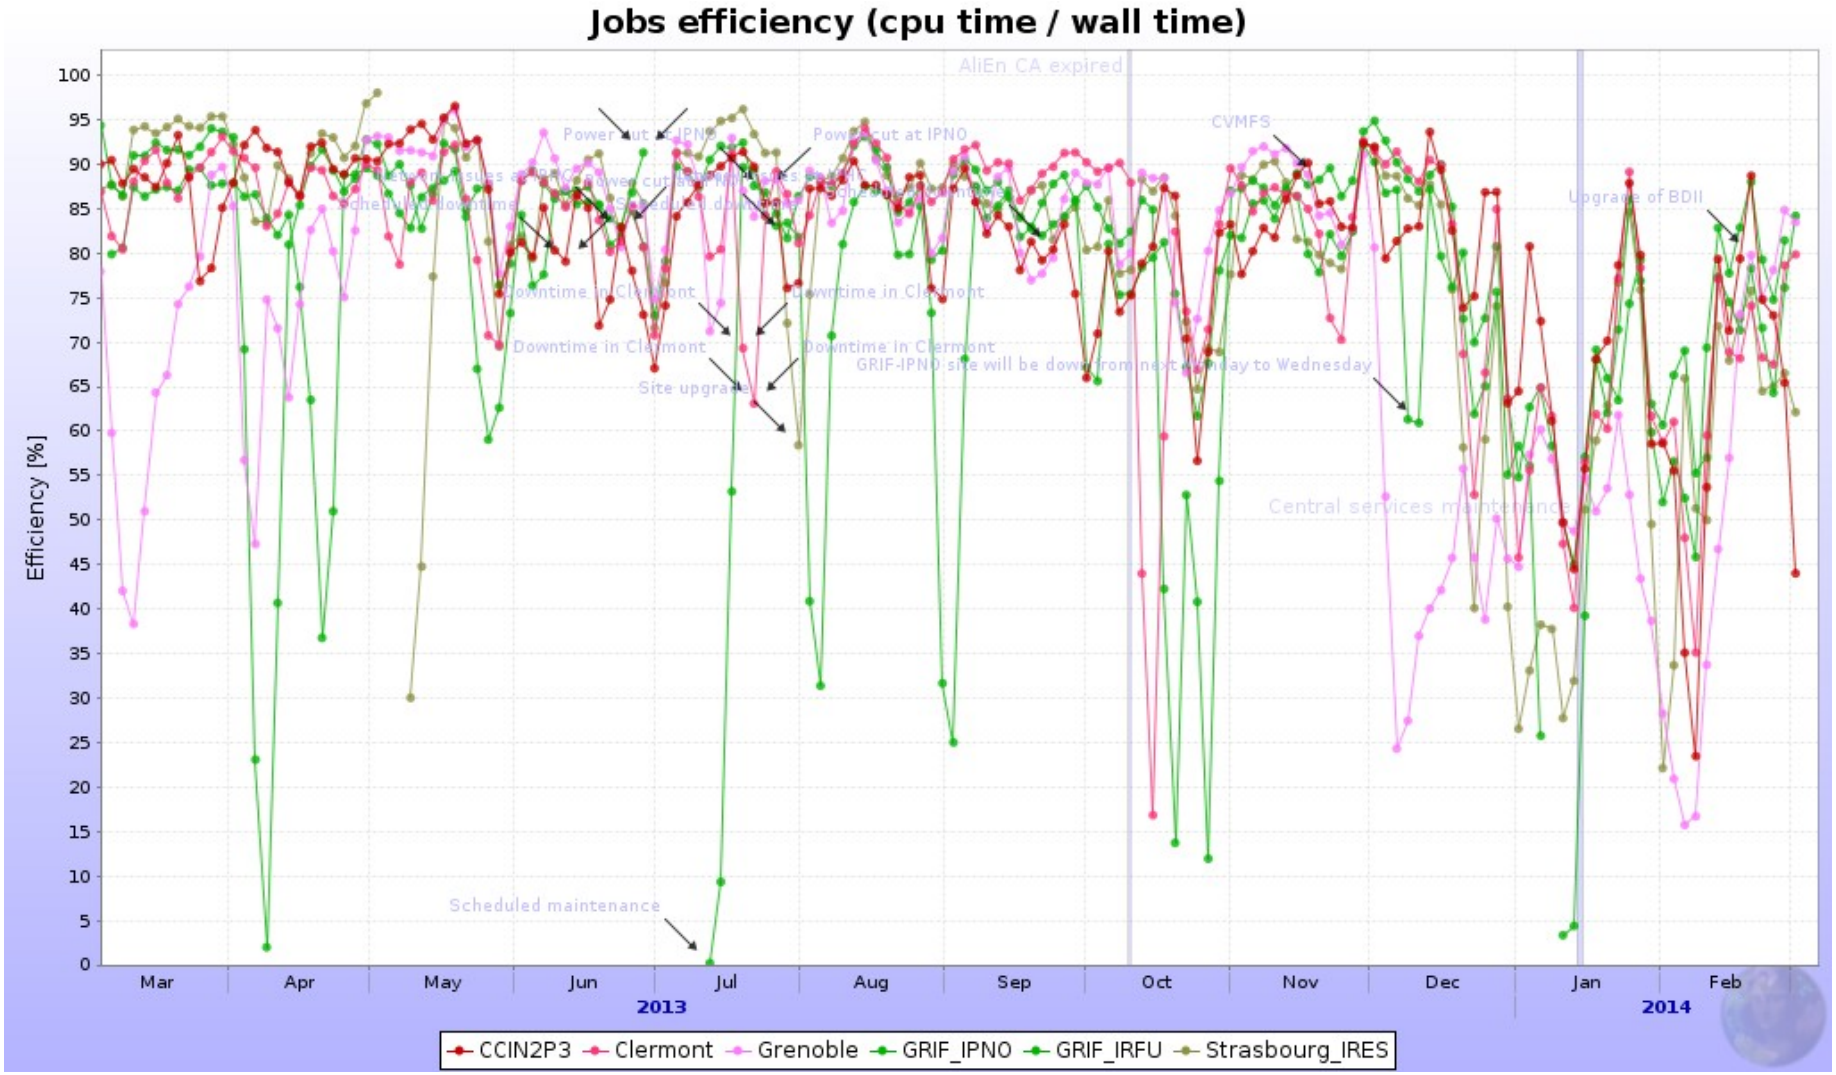

80% average (unweighted) including analysis, OK

Memory usage

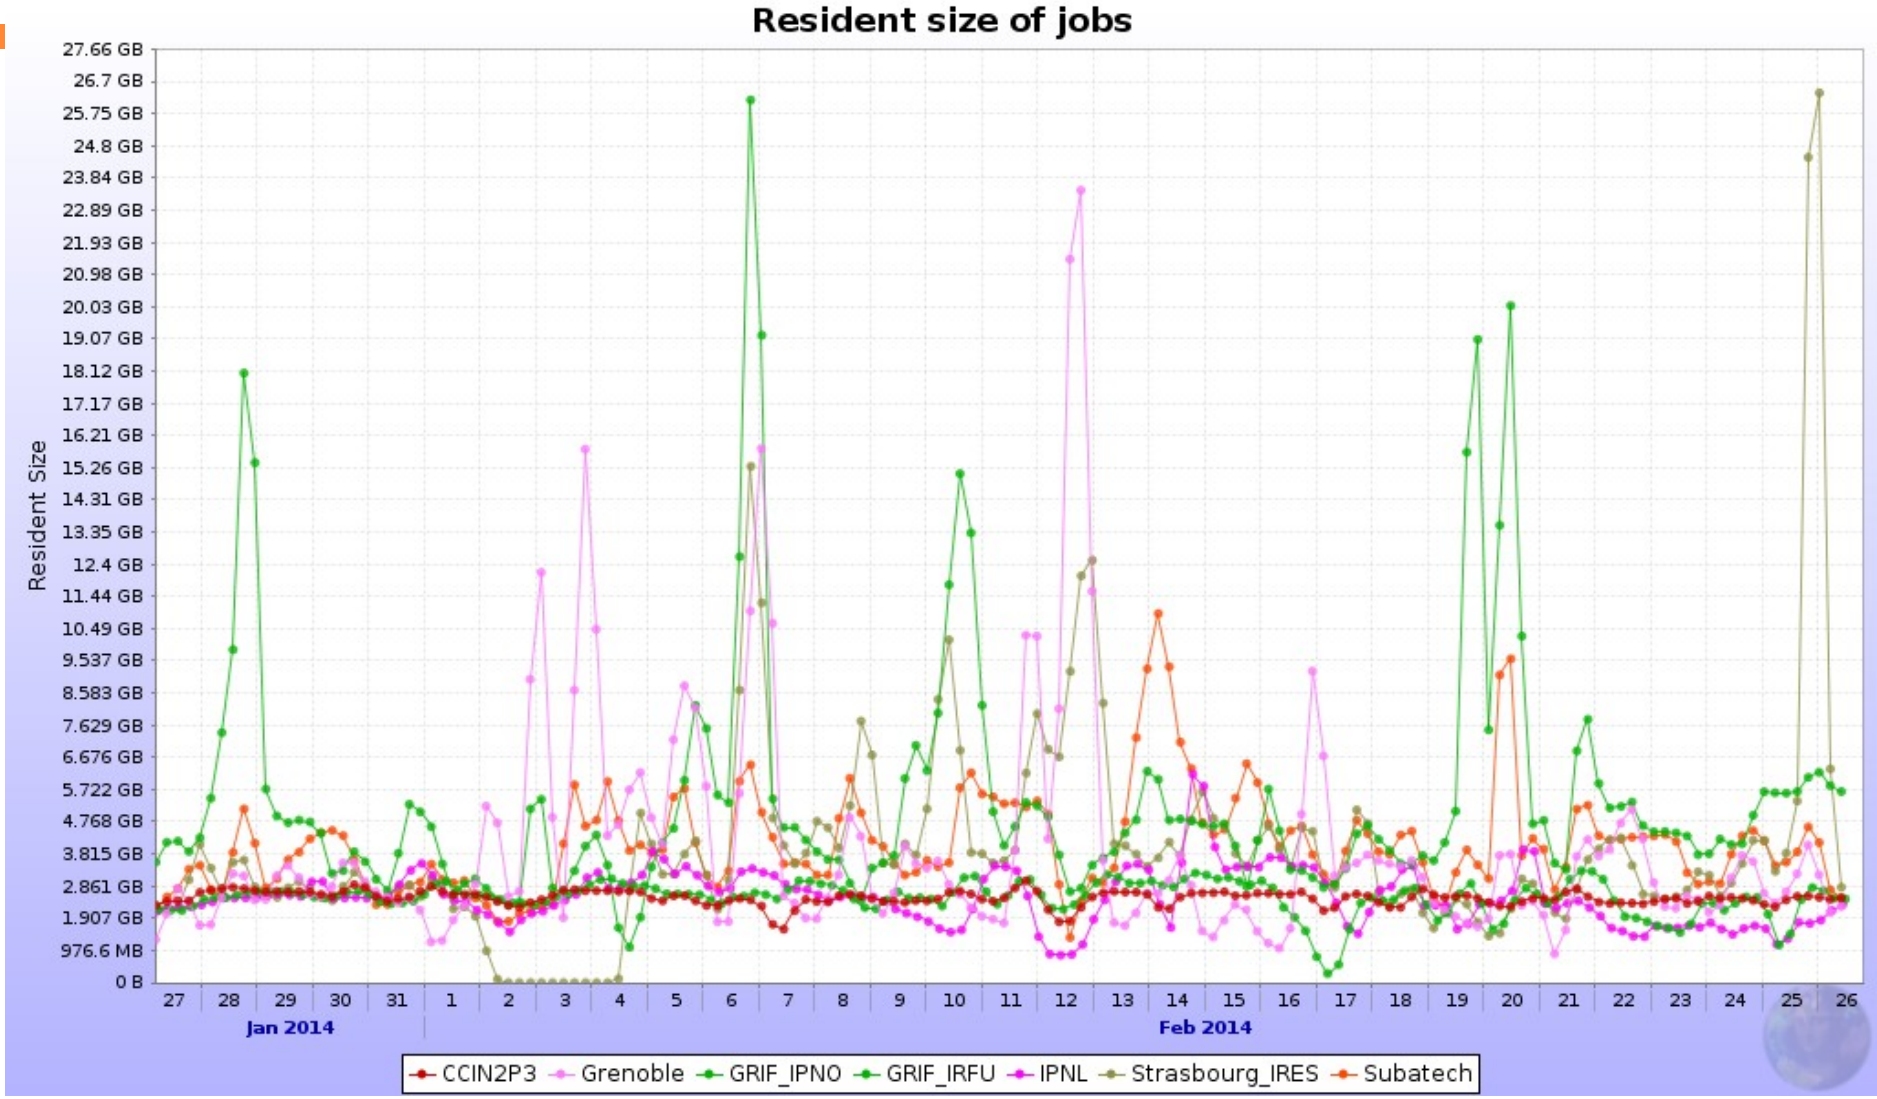

Situation improved, sometimes still problematic though...

CCIN2P3 report

- Availability > 97 % MoU
- Succesful transition from Oracle to Univa Grid Engine
  - Benefitful, especially for support concerns
- Transition bittorrent – CVMFS done
  - Not smooth, several days w/o jobs
- Xrootd 3.3.4
  - First attempts last summer failed (regular crashes)
    - Had to stick on 3.0.1 for several months some time
  - Now works fine for ::SE, still problems for ::Tape
    - Under investigation
- IPv6
  - Ongoing… goal is to be ready before the end 2014

Sites news/questions/issues

- LPSC
  - Infrastructure consolidation (virtualisation, 10 Gbps LHCONE)
- GRIF
  - New computing room
- IPNL
  - New backbone 10 Gbps
- Subatech
  - New room for workers
  - SAF used a lot
- No special news from other sites

Sites news/questions/issues

- CVMFS installed everywhere
  - Subatech -> CCIN2P3 -> stratum1
  - Should be done at other sites?
- EOS under test (Subatech)
  - Training needed
- Xrootd : data migration procedure ?
- Site feeling: few problems and good cooperation with the offline team ☺

ccin2p3 xrd connections

- Comparison of internal vs external connections to alice::ccin2p3::SE

In terms of # connections:
internal ~ external

Same shape

No direct traffic measurement of internal vs external

any suggestion ?

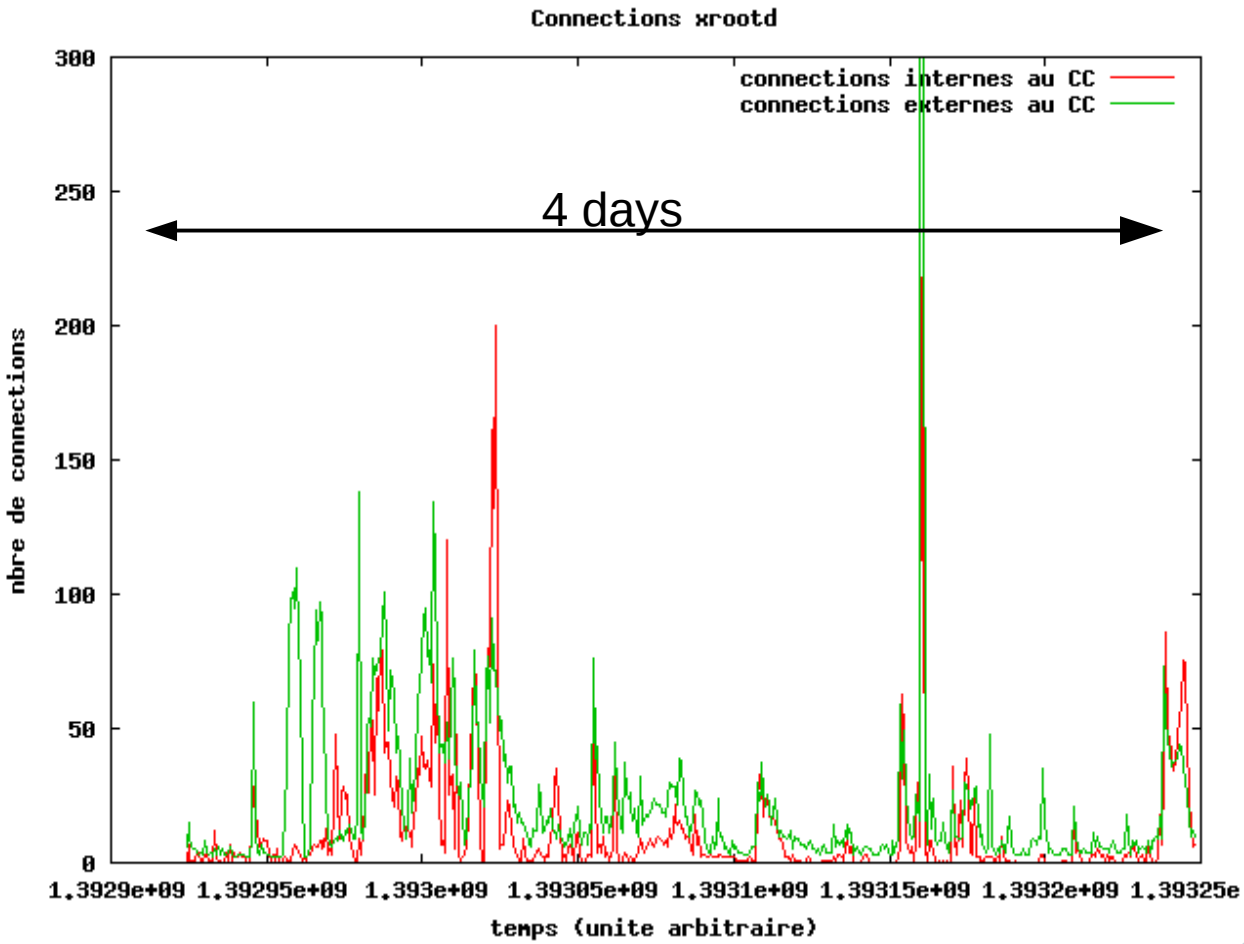

Connections vs traffic

connections

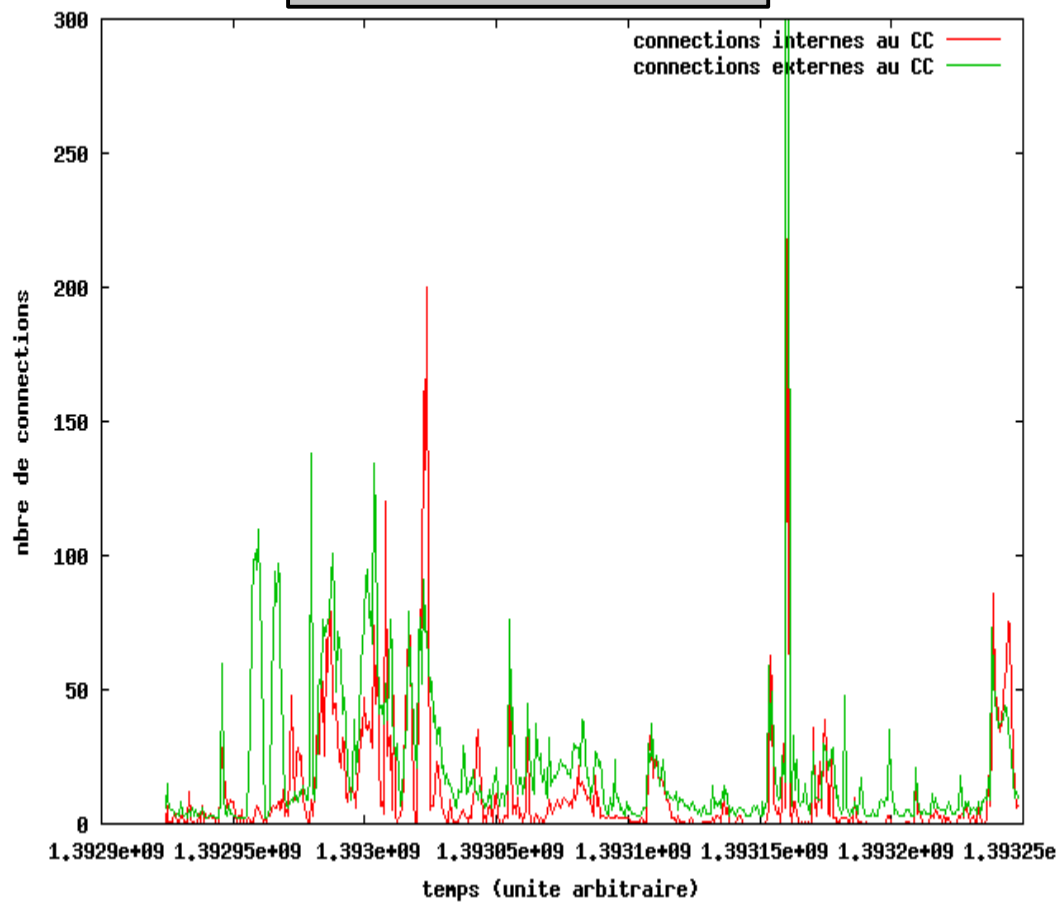

Traffic

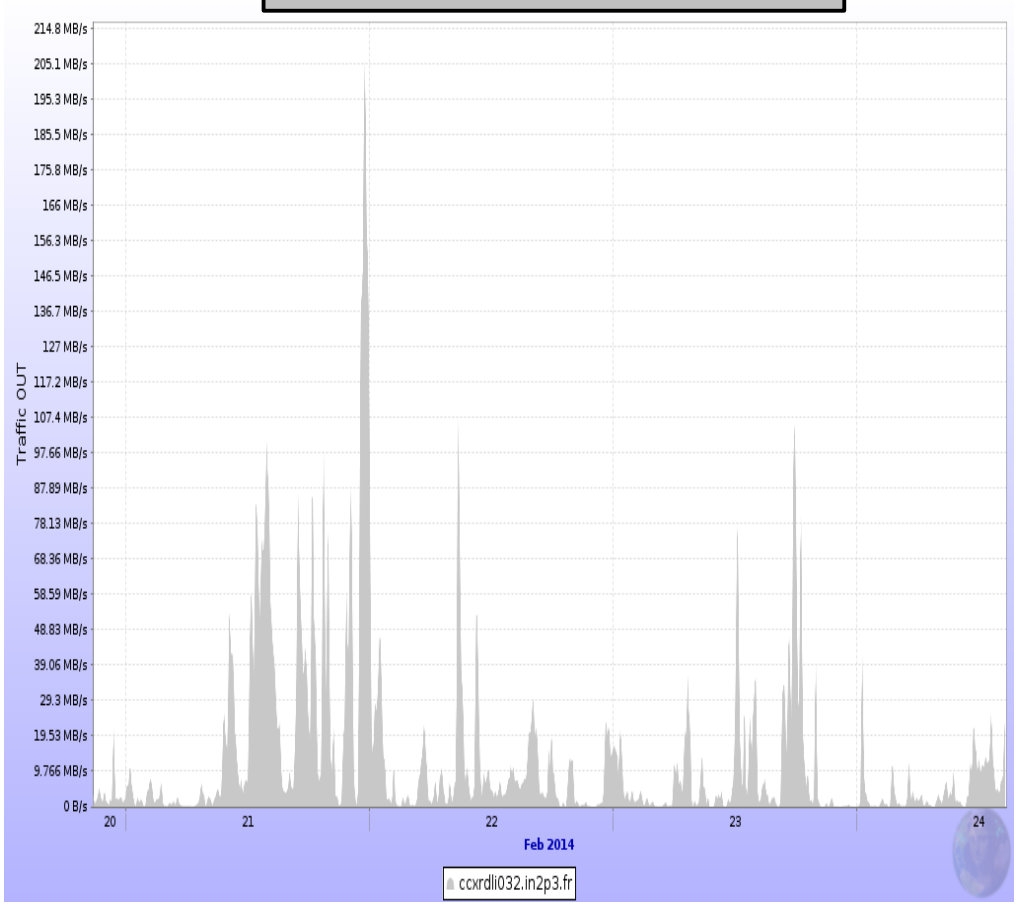

No trivial correlation
But good indication of the trend

Corollary

- # connections ~ scales with traffic
  - 50% of the connections come from the outside
  - ⇒ 50% of the traffic goes outside ccin2p3
- Concern…
  - especially since ALICE is no longer going to be the only VO exploiting WAN access to data
  - Last year, ALICE transfers were able to ~saturate Lyon's link to LHCONE for a few days

This probably needs
Work together with sites

Summary

- France provides ~10% of the WLCG resources
- ~60% of French CPU & Disk for ALICE provided by T2s and T3 (IPNL)
- Good availability/reliability of T1
- Network in the top5 in terms of bandwidth
- Good CPU efficiency (~80%) including analysis
- A few concerns for the near future
  - Remote data access strategy
  - Maintenance of AAF
  - Memory
- But globally very satisfactory operations and interaction with CERN core offline team

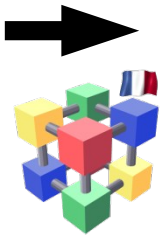

Ending questions

- AAF support ?
  - Problematic situation wrt CVMFS
- WN tarball maintained in EMI3 ?
- Multicore? Cloud?
  - We already provide that at T1
- ARC_CE ??