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

CCIN2P3 status report

ALICE T1/2 Workshop

Torino, 23-25 Feb. 2015

Renaud Vernet

- A national effort since 2004
  - Address LHC computing needs
  - T1,2,3 for all 4 experiments
  - Involvment of physicists
- An organisation
  - National coordination
  - Experiment representatives
  - Site representatives
  - Close collaboration with French NGI and NREN
- A community
  - Expertise, exchange, common actions

CCIN2P3 Computing rooms

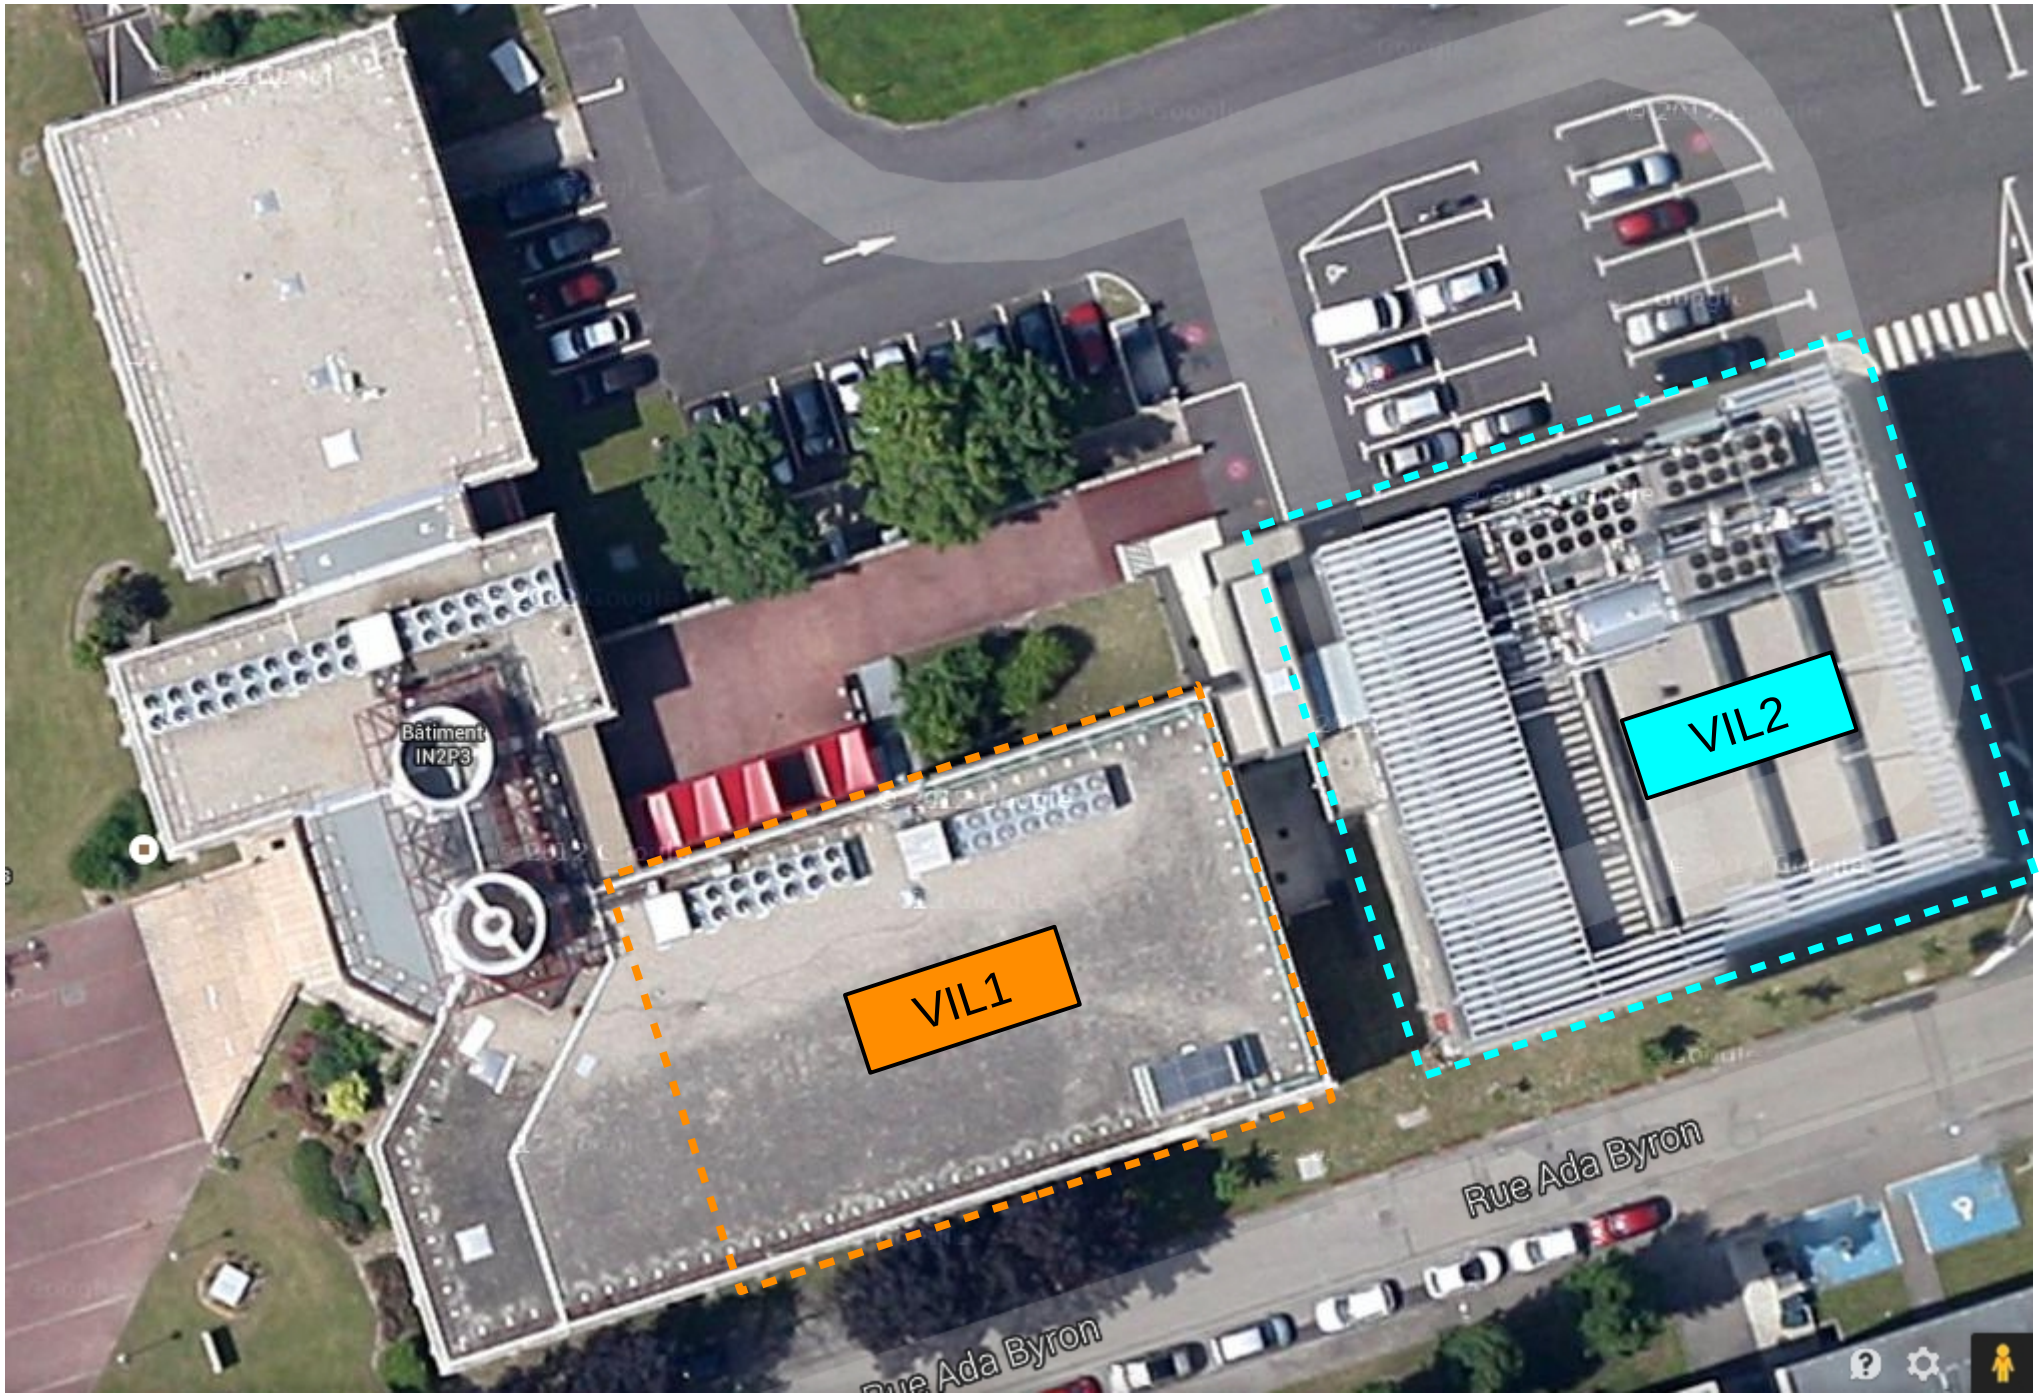

Network news

- Link Vil1 ↔ Vil2
  - upgraded to 2 x 100 Gbps
- Global simplification planned
  - Remove intermediate routers
  - One central router in Vil1 will stay
    - Connectivity with both internal and external adresses
- Improve connectivity with French NREN (Renater) in ~2016
    - Up to 100 Gpbs
- Interfaces
  - New storage to be in Vil2 : 2 x 10 Gbps links
    - (currently 2x1 Gpbs)

Network news (2) : IPv6

- 2 subnets in IPv6
  - Services
  - dCache
- Move to EfficientIP
  - Better management of migrations to IPv6 (esp. DNS)
  - In principle before June

- System
  - The Foreman + Puppet
  - Available test instances in SL7
    - CentOS is a serious possible alternative though
  - Continuous migration of services to virtualized hosts
    - Currently half of the services are virtualized
      - Openstack + VMWare

- Batch
  - Univa Grid Engine
  - Regular upgrades
  - Provides multicore
  - Still complex configuration
    - many users and experiments, with different needs

- Disk servers
  - Since last year, 4 TB disks
  - 1 server ~ 80 TB
  - Expected increase in bandwidth, saturations to monitor closely
  - Starting to deploy in Vil2
    - Benefit from better PUE
- HPSS
  - New technology T10K-D (StorageTek/Oracle)
  - Transition ongoing
    - Mix of several technos (T10K-B, -C, -D)
    - Currently 30 "D" drives in production
  - Expected growth of ~7 PB this year (total LHC)

- 2100+ running jobs over the last year *(vs 1400 in tsukuba)*
  - ~twice the pledged value
  - Move to special queue last summer benefitful (less zombie)

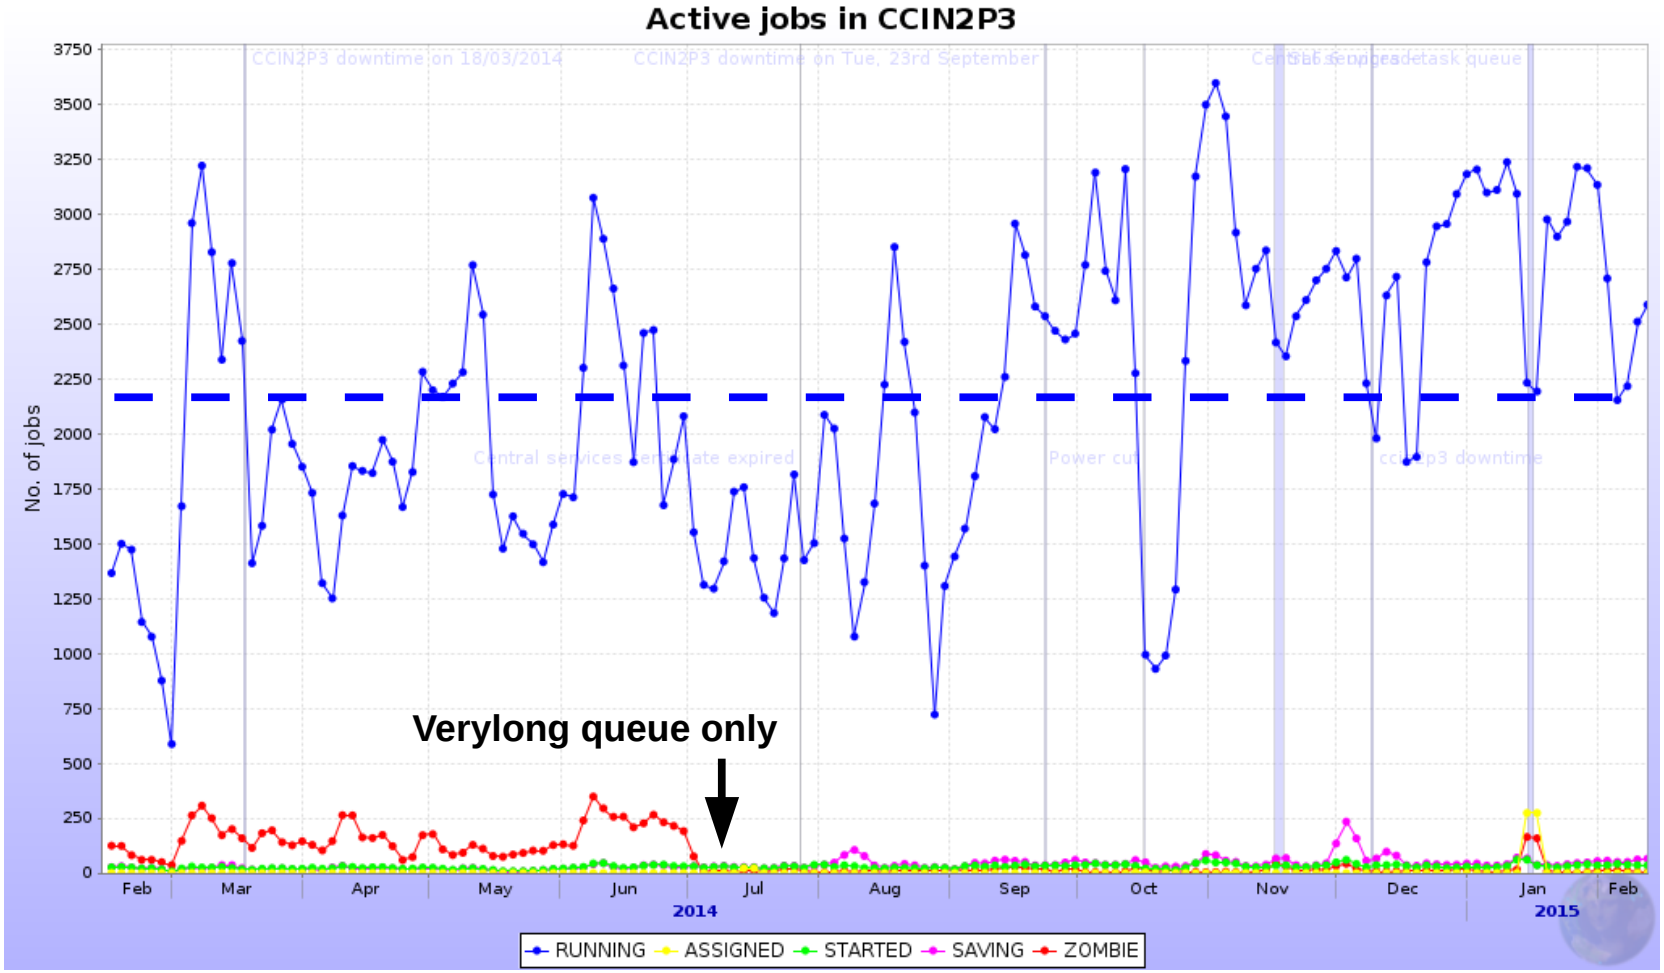

Availability / reliability

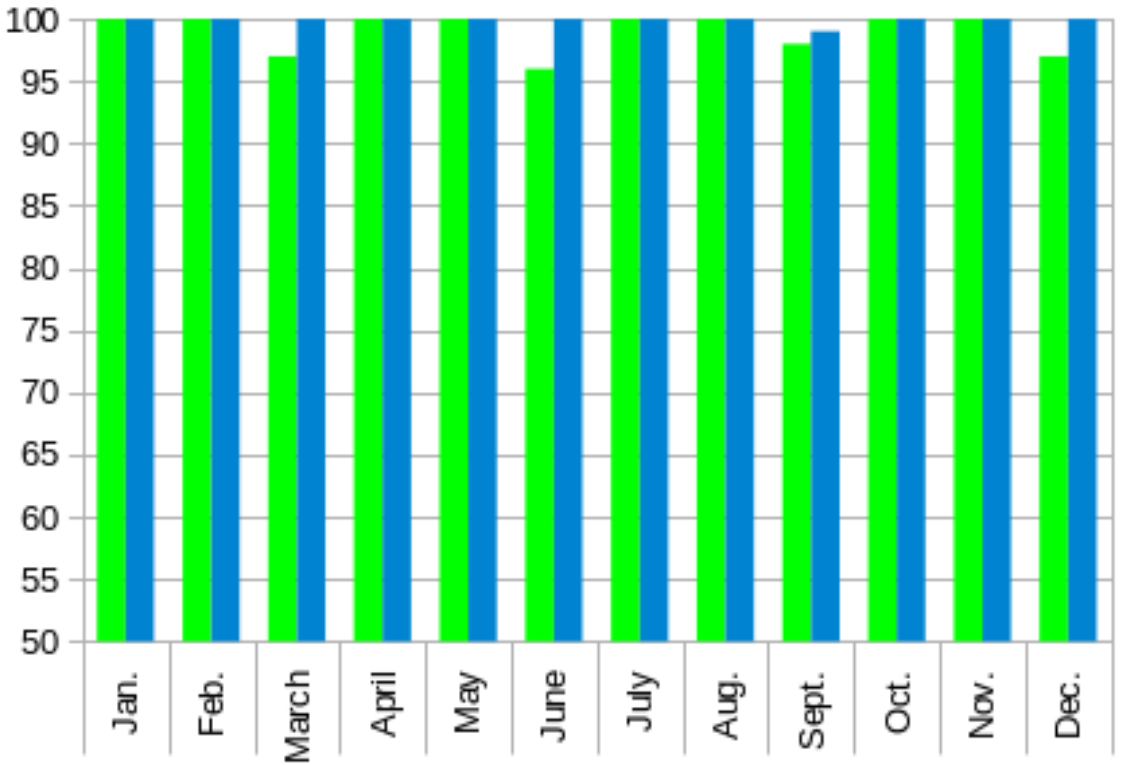

~99 %
~100 %

- No major show-stopper

Over the last year

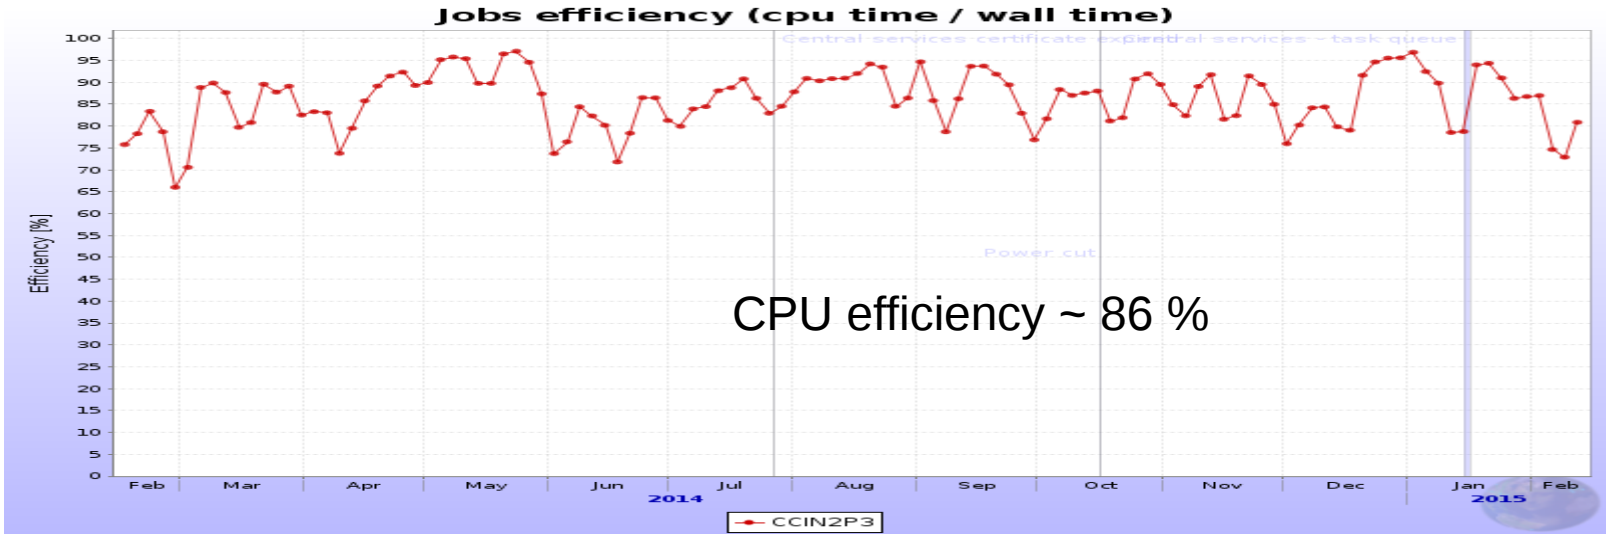

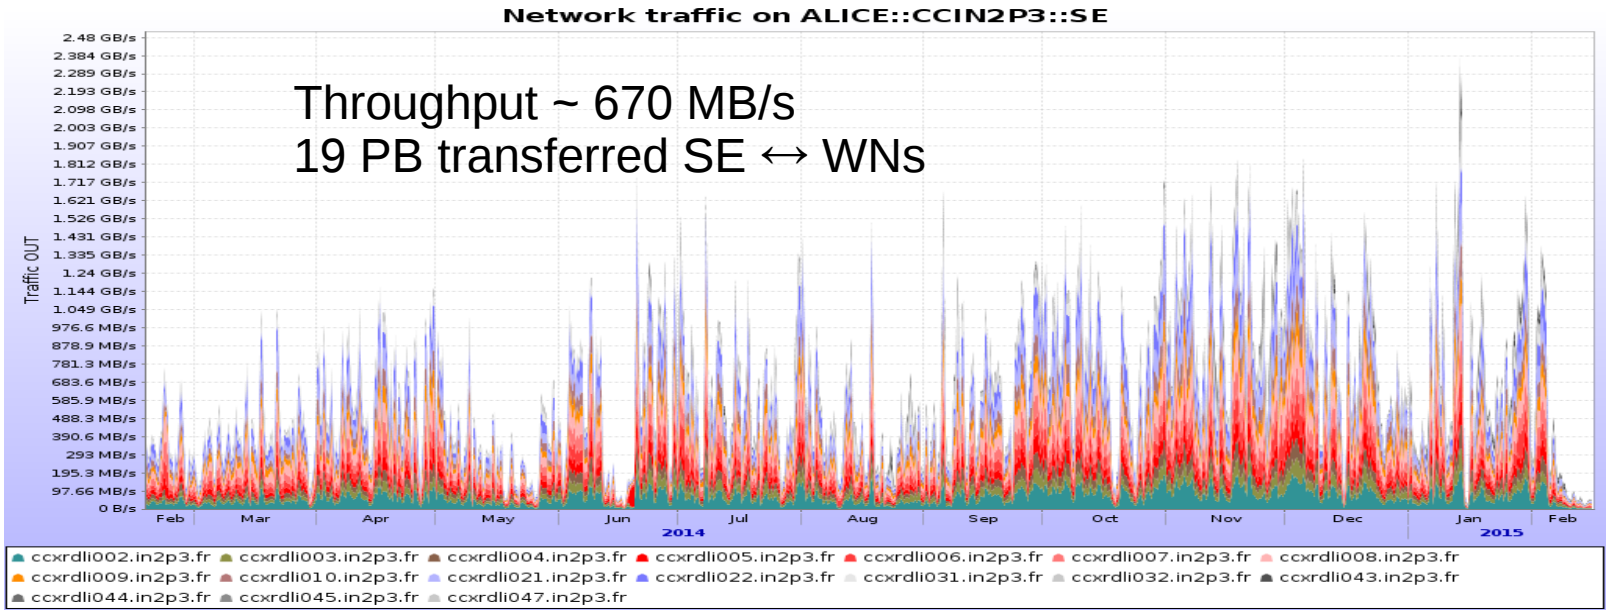

CPU efficiency vs server load (more detail)

- Only one server with heavy load can be enough to drop seriously the CPU usage
- We have to check periodically the storage status
- Restarting the daemon on the server usually helps

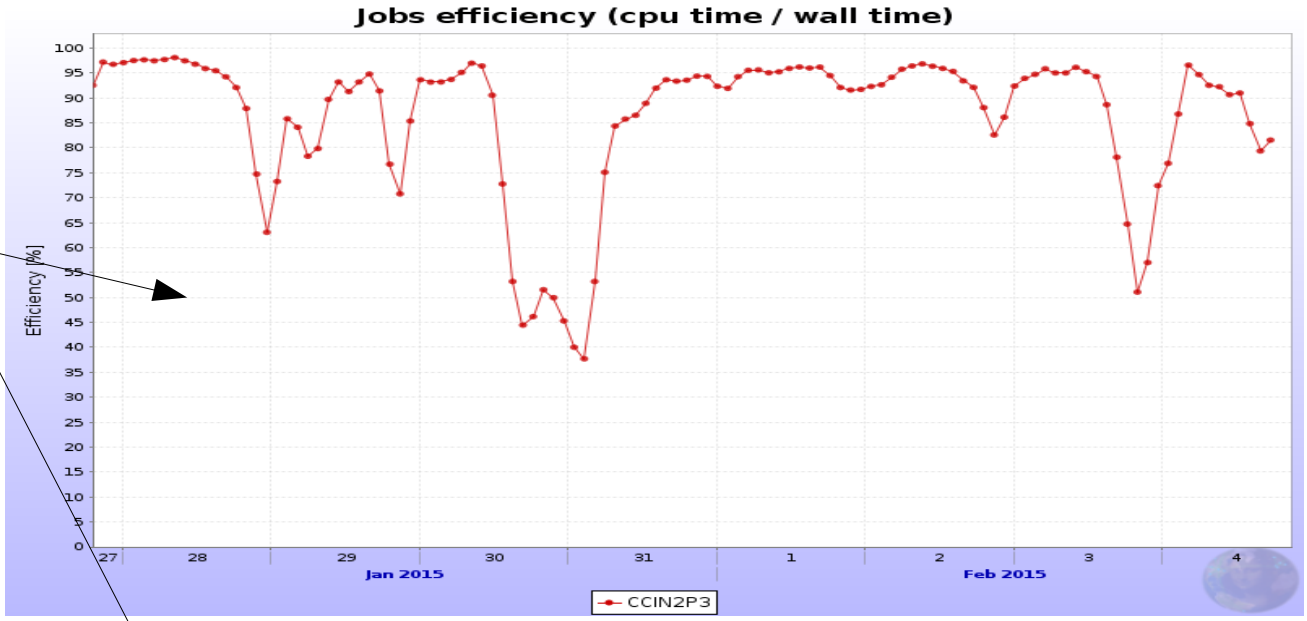

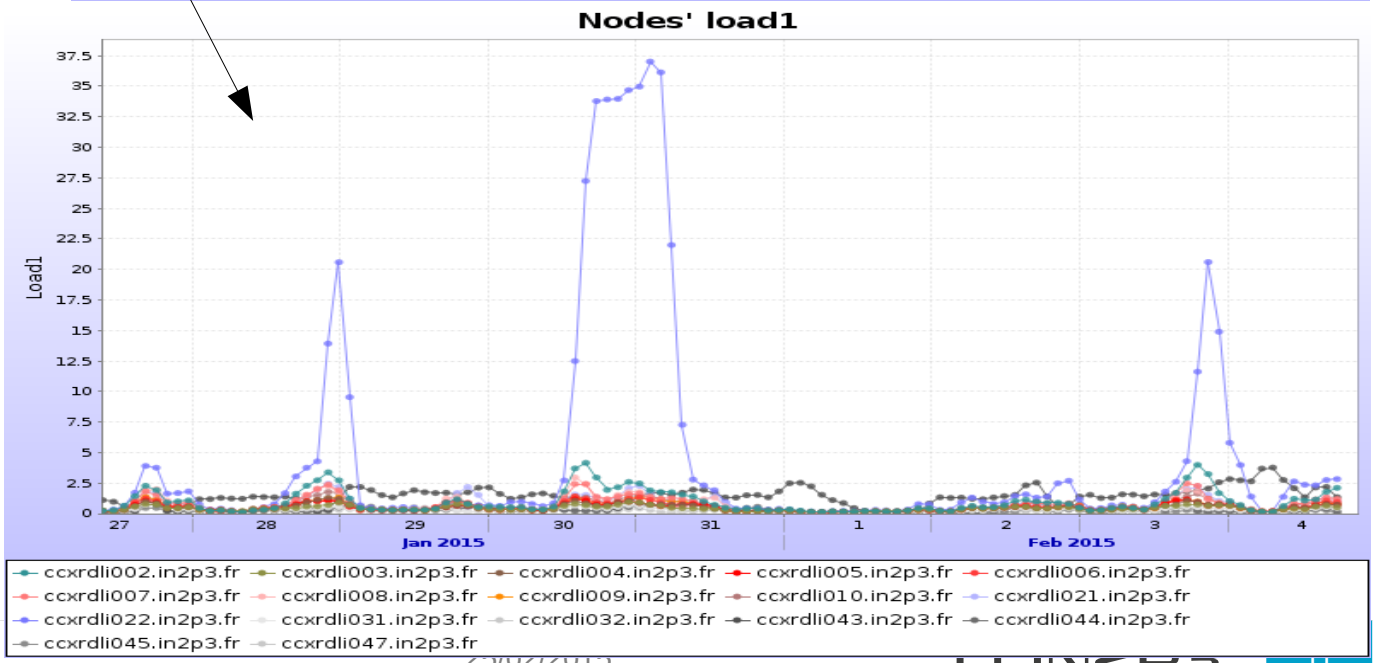

SE availability

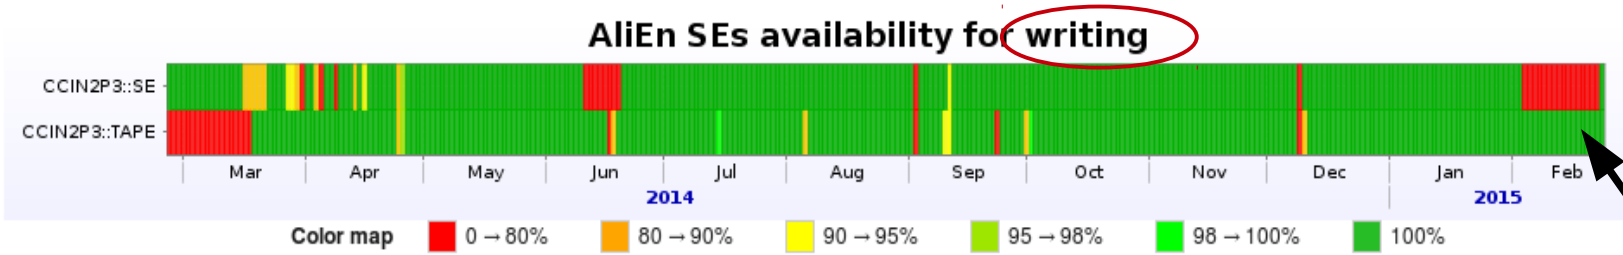

Statistics

Link name	Data		Individual results of writing tests			Overall
	Starts	Ends	Successful	Failed	Success ratio	Availability
CCIN2P3::SE	24 Feb 2014 14:03	24 Feb 2015 14:03	4002	376	91.41%	91.42%
CCIN2P3::TAPE	24 Feb 2014 14:04	24 Feb 2015 14:04	4099	283	93.54%	93.63%

Full storage

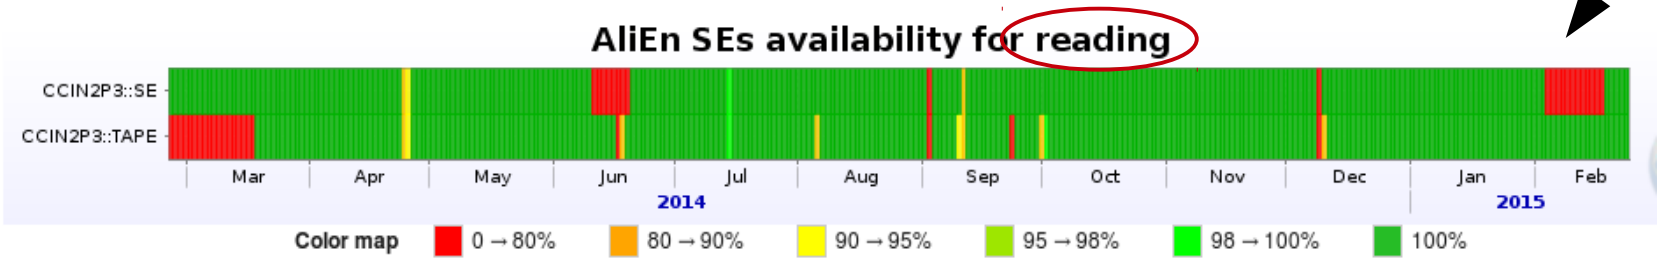

Statistics

Link name	Data		Individual results of reading tests			Overall
	Starts	Ends	Successful	Failed	Success ratio	Availability
CCIN2P3::SE	24 Feb 2014 14:03	24 Feb 2015 14:03	4083	296	93.24%	93.27%
CCIN2P3::TAPE	24 Feb 2014 14:04	24 Feb 2015 14:04	4098	284	93.52%	93.62%

ALICE resource growth at CCIN2P3

http://wlcg-rebus.cern.ch/apps/pledges/resources/

	2014	2015	Growth	% T1 requ.
CPU (HS06)	11500	14700	+ 28 %	12 %
Disk (TB)	1260	2180	+ 73 %	12 %
Tape (TB)	1050	1420	+ 35 %	14 %

8-10 % last year

- Nice resource growth for 2015 (LHC globally)
  - Bigger growth than ALICE requirement
  - Not much hardware to renew + budget OK
- Network backbone OK
- Tape buffer
  - Currently 300 TB
  - Should be resized for this year ?
    - Expected raw data in 2015 = 4.2 PB
    - => Tape buffer should be ~450 TB... OK ?

Conclusions

- Smooth operations in general
- Good support from the CERN team
  - Thanks
- Very opportunistic use of resources
- ALICE operations is not time consuming !
  - (at least for me...)

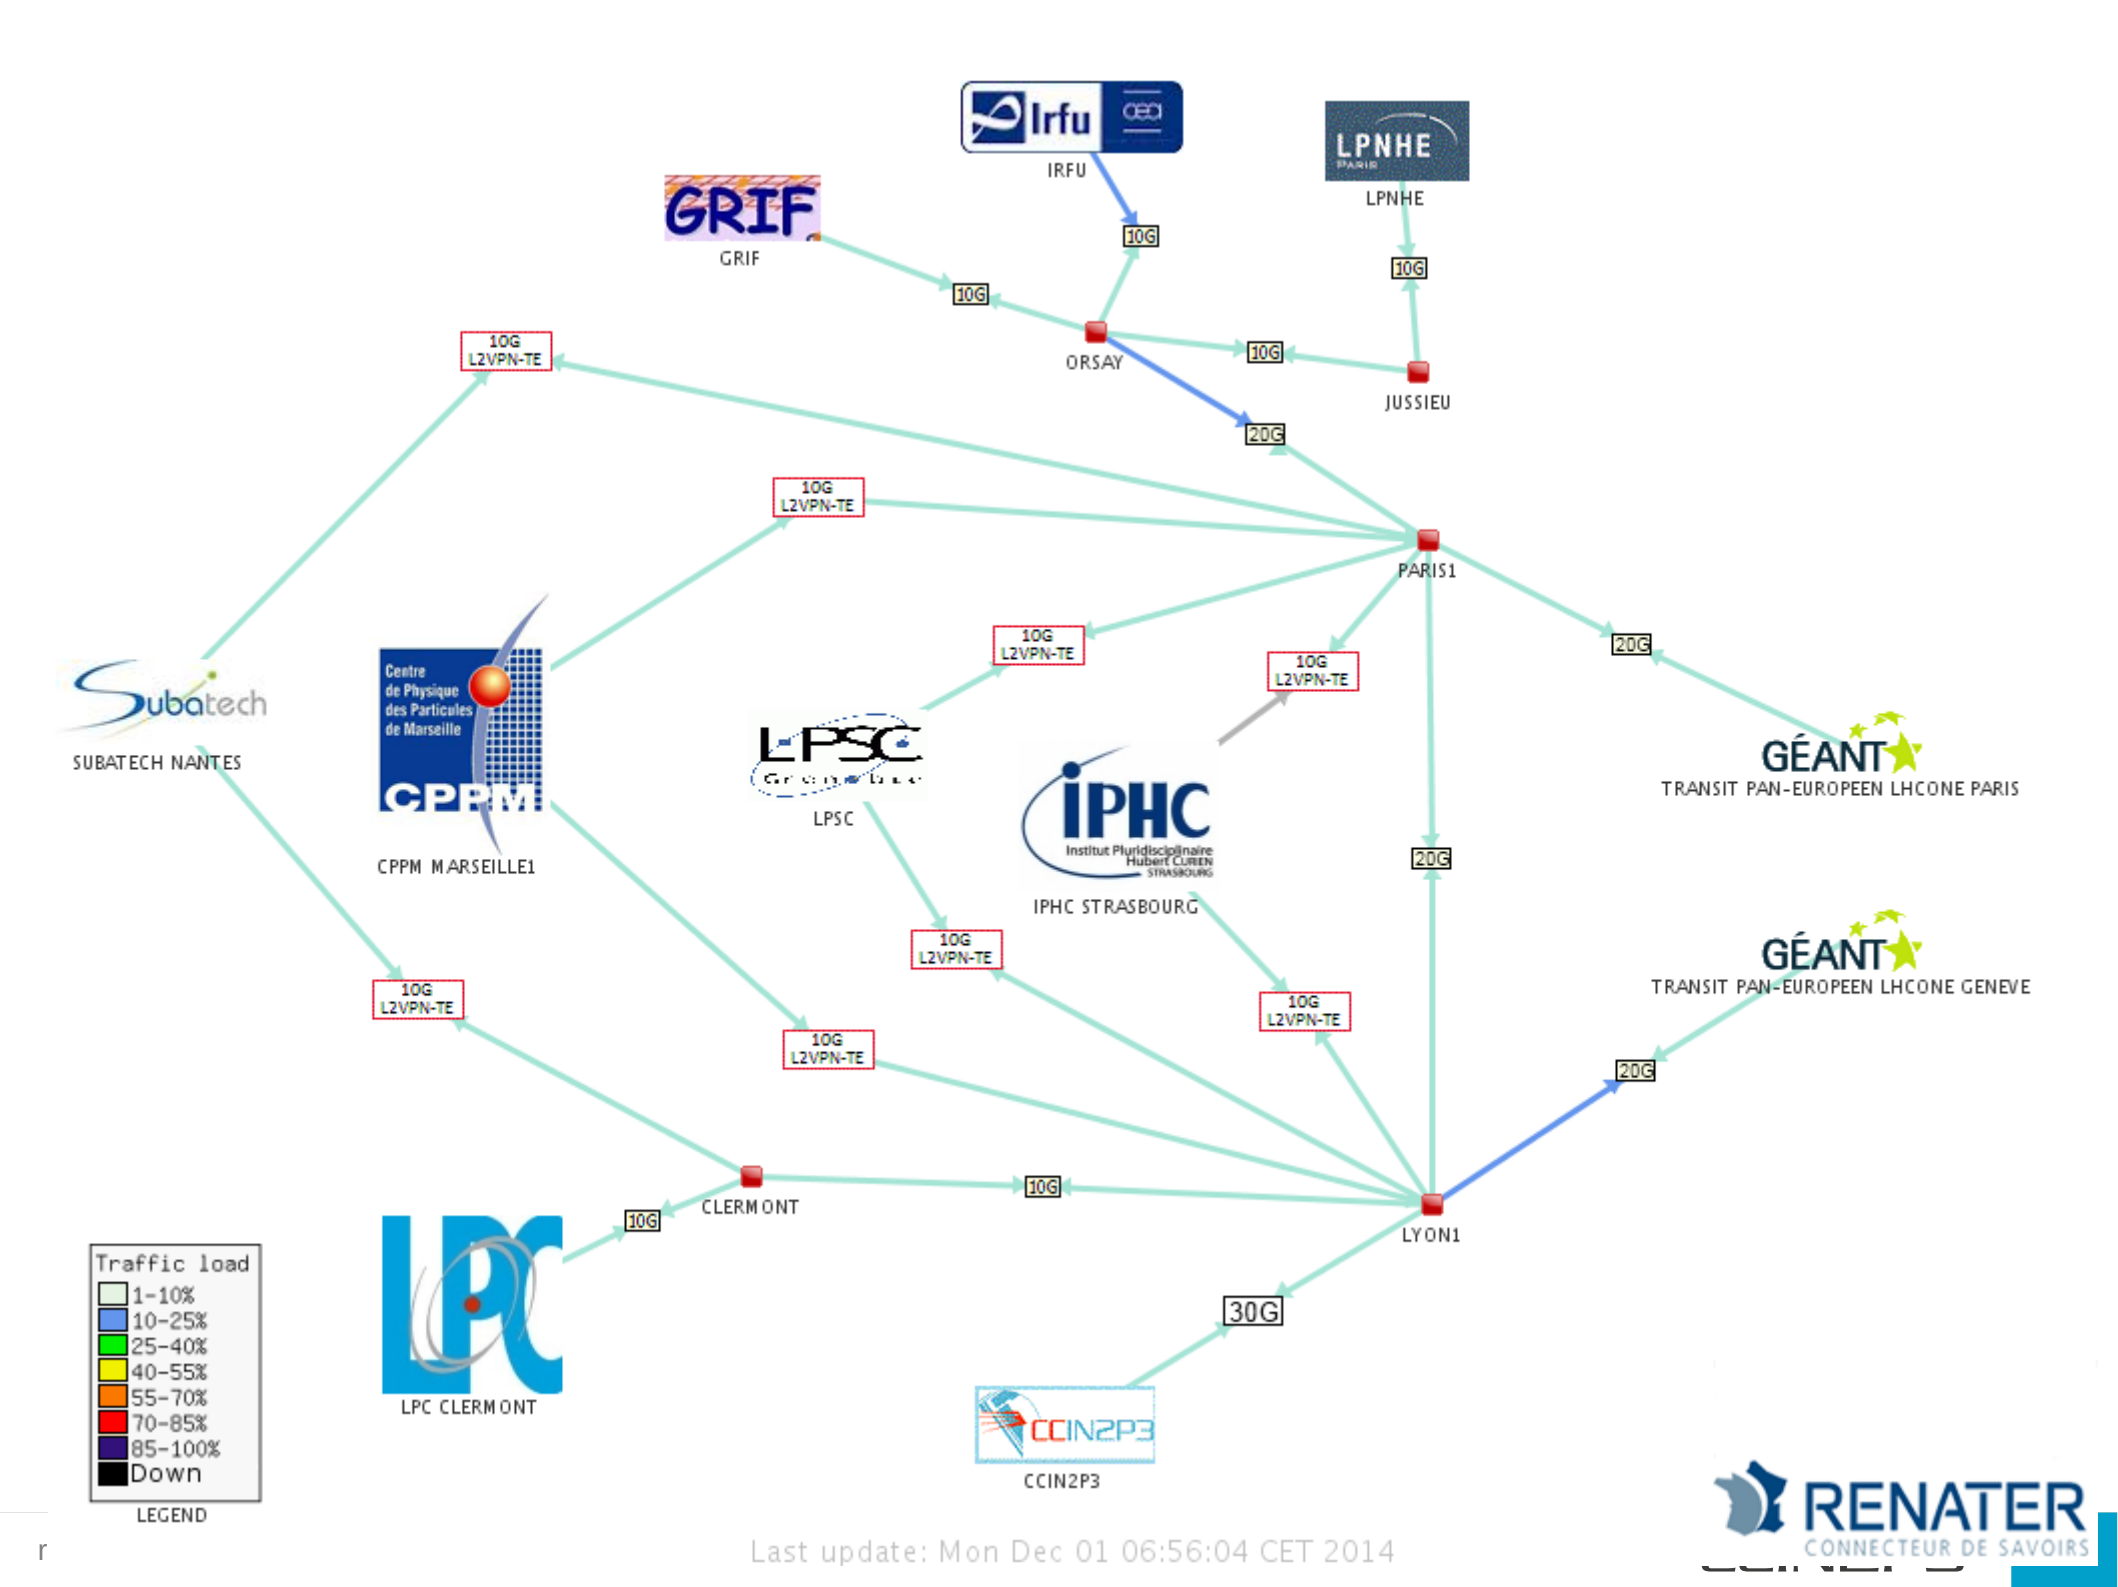
IRFU
LPNHE
GRIF
ORSAY
JUSSIEU
PARIS1
SUBATECH NANTES
CPPM MARSEILLE1
LPSC
IPHC STRASBOURG
TRANSIT PAN-EUROPEEN LHCONE PARIS
TRANSIT PAN-EUROPEEN LHCONE GENEVE
CLERMONT
LYON1
LPC CLERMONT
CCIN2P3
10G
20G
30G
10G L2VPN-TE
Traffic load
1-10%
10-25%
25-40%
40-55%
55-70%
70-85%
85-100%
Down
LEGEND