



CC-IN2P3 Site Report

HEPIX Spring meeting 2013

Bologna

Sébastien Gadrat





Outline



- ✓ General information
- ✓ Staff changes
- ✓ Batch system & Cloud
- ✓ Storage
- ✓ Quality management
- ✓ Perspectives

General information about CC-IN2P3 - Lyon



- ✓ French National computing center of IN2P3 / CNRS
in association with IRFU (CEA – Commissariat à l'Énergie Atomique)


- ✓ Groups and Users (with their respective CPU consumption):
 - ✓ Tier-1 for LHC experiments (ATLAS, CMS, ALICE and LHCb) : 70%
 - ✓ Other HEP experiments (mostly D0): 8 %
 - ✓ Astroparticle experiments: 9 %
 - ✓ 94 groups in total (see next slide):
 - ✓ 69 groups in physics (HEP+astroparticle+nuclear physics): 98 %
 - ✓ 25 groups in multidisciplinary (biology, others): 2 %

- ✓ Dedicated support for the 4 LHC experiments, astroparticle and biology.

Diversity of scientific activities using the CC



Number of groups by scientific domain and activity in 2013





Staff changes



- ✓ **Pierre-Étienne Macchi** replaces Dominique Boutigny as director (since January)
- ✓ **Philippe Olivero** retired last month
- ✓ **David Bouvet** replaces Ghita Rahal as Support team leader (since January)
- ✓ Important people turnover foreseen this year (end of contract for several people)
 - ✓ Difficult to anticipate (rules for temporary contracts are getting tougher)
 - ✓ Budget issue as well



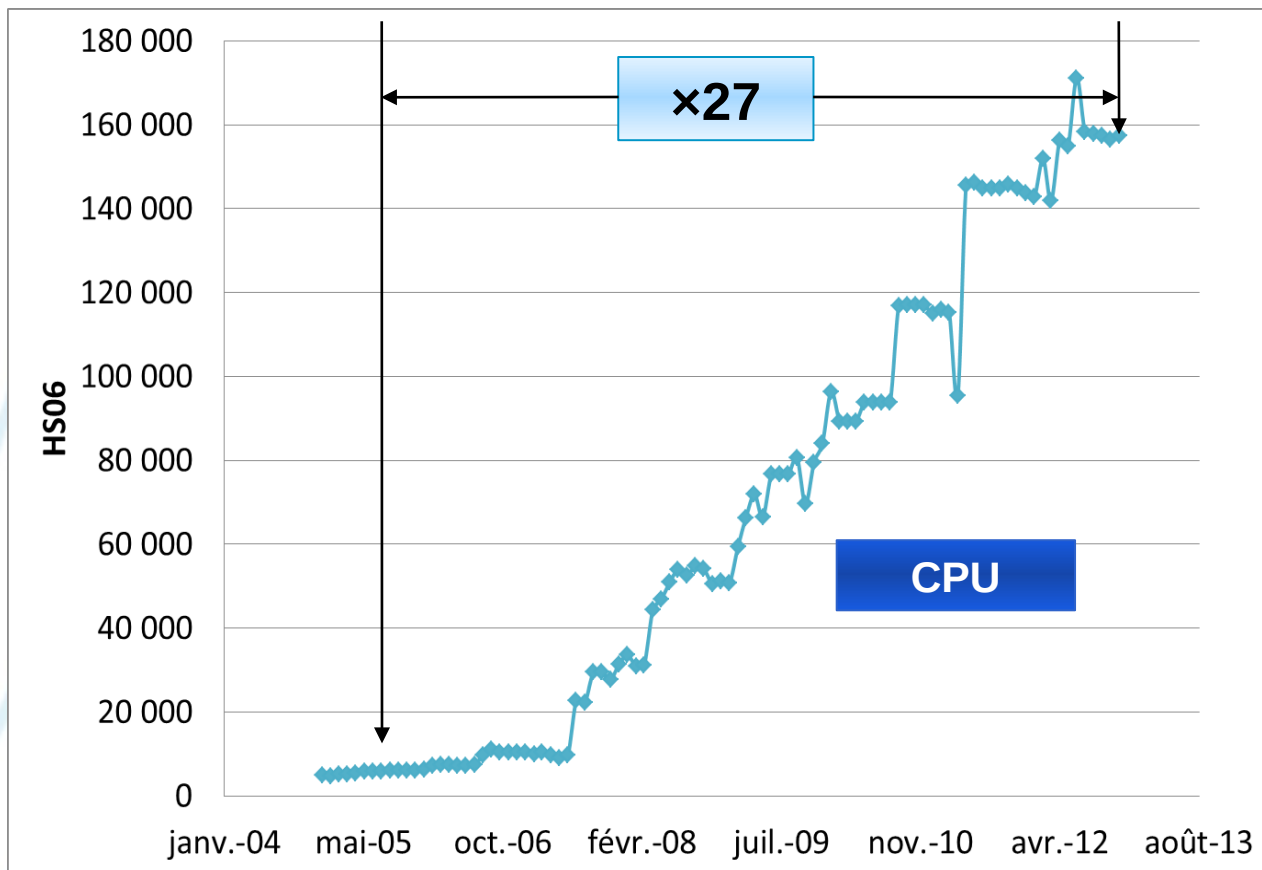
Batch and Cloud



- ✓ Univa Grid Engine (v8.1.4) will replace Oracle Grid Engine (v6.2update6_20)
 - ✓ Univa support began this April
 - ✓ Migration foreseen end of May, expected to be transparent for users

- ✓ Main motivations for this migration
 - ✓ Low reactivity from the Oracle support (bug fix)
 - ✓ Low visibility on the long term evolution of the product (required features were assigned low priority, when some were already in UGE)

Evolution of the computing power (HS06)



Status of the workers as of April 2013



Dell Power Edge / Nber of cores					Total					
					Number of workers	Number of cores	Worker power (HS06)	Power %		
G E	prod	INTER	Pwg 1950b	8	4	32	302	0 %		
			Total		4	32	302	0 %		
		SEQ	Pwg C6100	24	315	7 560	69 556	41 %		
			Pwg C6220	32	160	5 120	56 218	33 %		
			Pwg M610	16	149	2 384	22 076	13 %		
			Total		624	15 064	147 850	88 %		
		MULTICO RES	Pwg C6100	24	1	24	221	0 %		
			Pwg M610	16	25	400	3 704	2 %		
			Total		26	424	3 925	2 %		
		MULTISE Q	Pwg M610	16	50	800	7 408	4 %		
			Total		50	800	7 408	4 %		
		PARA	Pwg M610	16	64	1 024	9 482	6 %		
			Total		64	1 024	9 482	6 %		
		Total					768	17 344	168 967	100 %
		Total					768	17 344	168 967	100 %

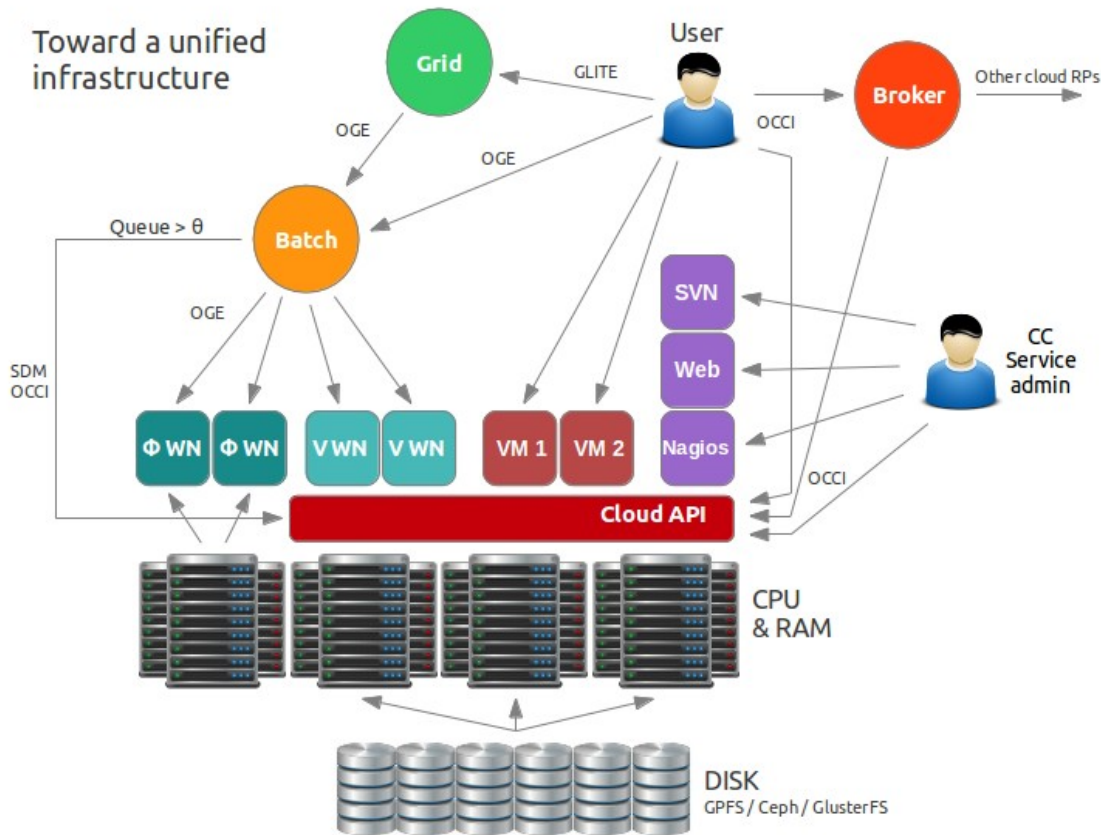
+7 % in HS06 compared to last April, 2012 (last Site Report).

Cloud@CC-IN2P3



- ✓ A Cloud project has been initiated in Spring 2012
 - ✓ Goal is to provide an OpenStack based IAAS
 - ✓ Federated with EGI FedCloud testbed

Toward a unified infrastructure



✓ Main use-cases

- ✓ LSST/EUCLID/Integral/LHC
- ✓ Academic cloud
- ✓ Private cloud for infrastructure services

**Presentation at next
Fall HEPiX meeting**




Storage

Storage news and evolution



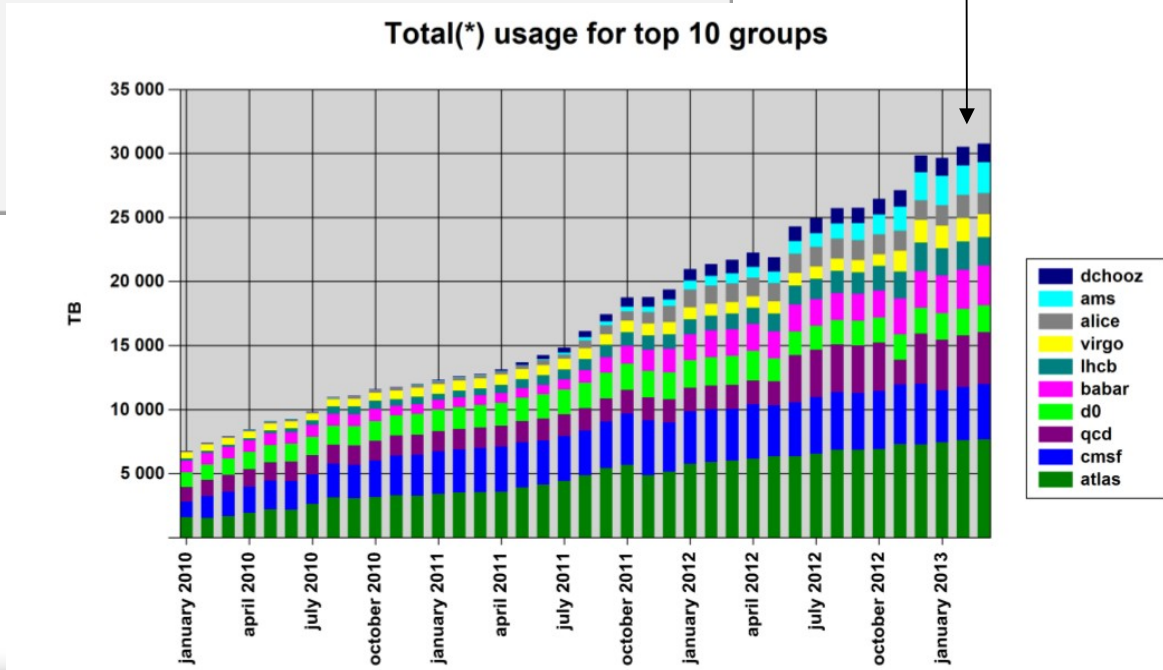
- ✓ No increase foreseen this year, except
 - ✓ Tapes (HPSS): +5.5 PB for 2012 (~-1 PB for LCG), +4 PB foreseen 2013 (~2 PB for LCG)
 - ✓ Disks (dCache + xrootd): +600 TB (all for LCG)
- ✓ Constraints
 - ✓ Lots of hardware to be replaced
 - ✓ Tight budget
- ✓ Perspectives (strongly depending on the budget)
 - ✓ More disks + disk renewal for xrootd (besides LCG) and GPFS
 - ✓ Robot for backup double copy in TSM

Storage usage



Total storage
(disks + tape)
> 30 PB

MINIMUM	AVERAGE	MAXIMUM	Files
26.029M	29.285M	34.197M	files-lcg-chimera/files-atlas
1.569M	1.692M	1.870M	files-lcg-chimera/files-cms
1.233M	1.513M	1.707M	files-lcg-chimera/files-lhcb



✓ Quite a big increase over the few last months (years)!

(*)Total for dcache, hps-bande, iRODS, SRB and xroot services.

Storage Summary



■ dCache (updated)

- 2 instances: LCG (v1.9.12-25) & EGEE (v2.2.8)
- +300 pools on 180 servers
(x4540 Dell R510/SL5, Solaris being retired)
- ~7,4 PiB
- 37.5 M files, 7.5 M directories

■ SRB

- v3.5
- 250 TiB disk + 3 PiB on HPSS
- 8 x4540 32 TB 10/ZFS + MCAT on Oracle

■ iRods

- v2.5
- 11 x4540 32/64 TiB + 1 Dell v510 54 TiB
- 520 TiB disk space, 1.7 PiB managed (incl. Tapes)
- ~22 M files

■ xrootd

- v3.0.5
- 20 x4540 32/64 TiB + 16 R510 54 TiB
- 1.5 PiB

■ HPSS (updated)

- 22 PB, 41 M files on tapes / 432 TB on disks
- v7.3.3.8 on IBM AIX 6.1
- Core server: IBM eServer p550
- Tape drives: 60 x T10K-B / 24 x T10K-A / 15 x T10K-C
- 11 x tape servers
 - 6 x IBM AIX (p520)
 - 5 x Dell PowerEdge R710/720 RHEL5
- 12 x Dell disk servers
 - PowerEdge R510 + PowerVault
 - MD1200 + 10Gbps Eth + RHEL5 = 36 TB
- ACSLS 8.2

■ OpenAFS (updated)

- v1.6.1 and v1.6.2 (servers), v1.6.1 (worker nodes)
- 42 GiB
- 42 servers, ~1400 clients

■ SPS (GPFS) (updated)

- v3.5.0.7
- 370 M files, 35 filesystems (67 groups)
- 1.2 PiB (7 x IBM DCS9550 + 2 x IBM DCS3700)
- 850 nodes, 35 servers

■ TSM (Last update Nov 2011)

- 4 servers (AIX 6, TSM 6) each w/ 4 TiB on DS8300
- 1 billion+ files, ~1 PiB
- 20 LTO4 drives, ~2000 LTO4 tapes
- 3 TiB / day

Quality management and perspectives

- ✓ CMDB project : see Emmanouil's talk on Monday afternoon

- ✓ New ticketing system OTRS in production since last October

- ✓ Second computing room fully operationnal
 - ✓ Electrical redundancy done last December (see next slide)
 - ✓ Should allow to cope with the LCG ressources increase for the future (until ~2020)
 - ✓ Will allow to provide resources to new experiments (astroparticle : LSST, CTA, EUCLID)
 - ✓ CC-IN2P3 will process 50 % of LSST data and store them all

Second computing room evolution



Currently

80 racks* (26 used so far)

* capacity installed

1.5 MW

Redundancy N+1

Redundant HV

2011

50 racks

600 kW

Redundancy N+1

Minimum electrical
autonomy

2015

120-160 racks*

1.5 MW

Redundancy N+1

Redundant HV power

2019

240 racks*

3.2 MW

Redundancy N+1 and 2N

Redundant HV Power

On top of the existing 1 MW computer room

- ✓ The electrical redundancy is provided by 2 independent HV power supplies
- ✓ Successful tests to run in degraded conditions

- ✓ Relatively small resource increase this year due to tight budget
 - ✓ Increase for LHC experiments only (mainly tapes, but disk too)
- ✓ Partnerships with the industrial world (new biomedical projects recently started: eTricks, IRT BIOASTER)
- ✓ The second computing room will allow
 - ✓ To cope with the increase of the LHC experiments needs
 - ✓ Support new experiments (LSST, EUCLID, CTA)



Thank you !



Questions ?