



Centre de Calcul

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

The CNRS logo, consisting of the letters 'CNRS' in a white, sans-serif font inside a dark blue circle. The background of the top right corner of the slide is a photograph of a server rack aisle with blue lighting.

cnrs

CC-IN2P3 report status

FJPPL 2024 – 30/01/2024

Preview



- **Missions**
- **IN2P3 Experiments**
- **Staff**
- **Computer rooms**
- **Equipex+ FITS**
- **Electricity cost evolution**
- **Computing platform evolutions**
- **Mass storage platform evolutions**
- **Infrastructure deployment methods**
- **Network connectivity evolution**
- **Summary**

Missions



Providing IT resources to the scientific experiments supported by IN2P3.

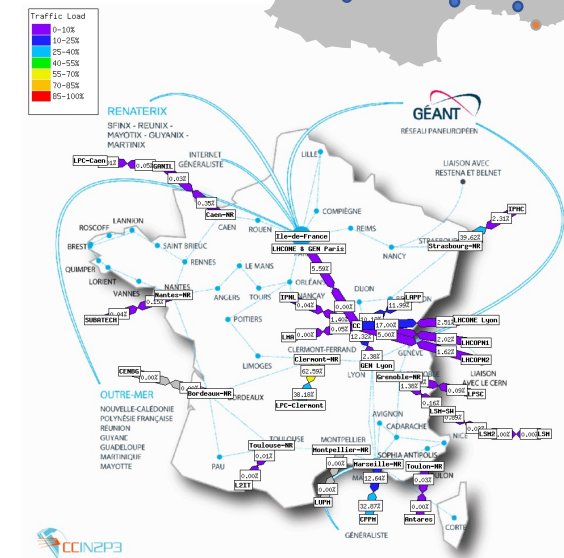

- Data storage, analysis and computing resources

Providing common institutional services

- Mail and messaging services, project management tools, documents management service, CAD services, etc.

Providing, managing and planning future evolutions of the network connectivity of IN2P3 sites

- IP management, routing, coordination of network operators deployments



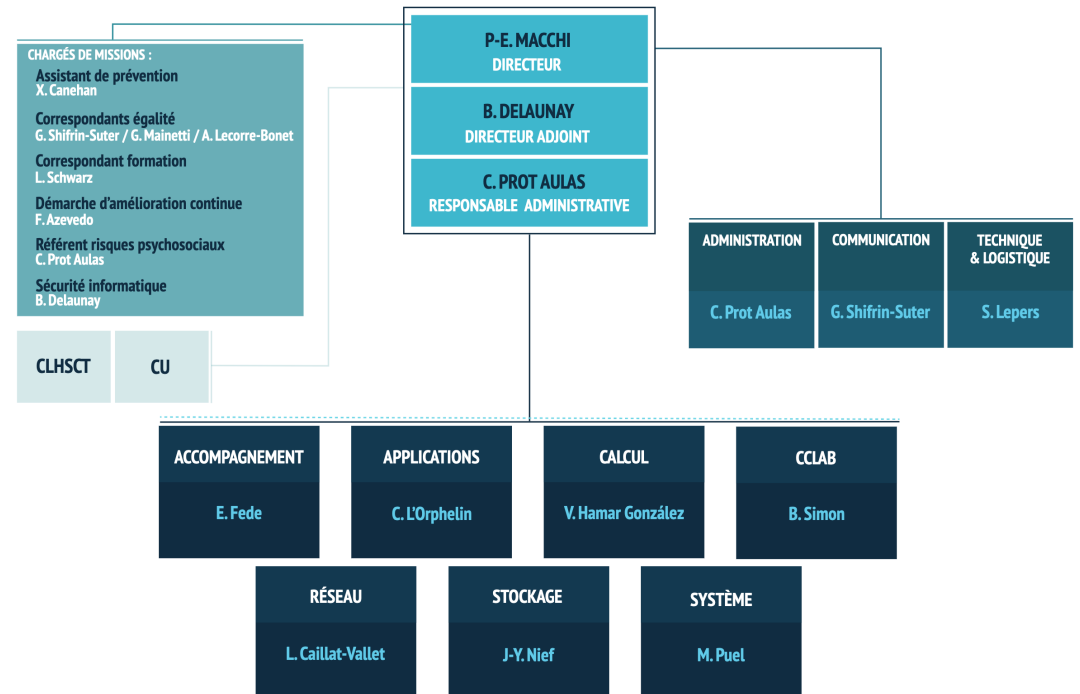
Staff



80 people, 85% on permanent positions

- 2 researchers
- 14 people non IT (Administration and financial service, Communication, Facility Service Management)
- 64 people for IT (Applications, Computing, Network, System, Storage, Support)

The number of people working at CC-IN2P3 is constant for years



Computer rooms



2 computer rooms VIL1 and VIL2

VIL1

- Built in 1986
- 850 m², 1MW IT, 150 racks
- Old cooling system in the raised floor (PUE=2)



VIL2

- Built in 2011
- 850 m², 3.6 MW IT, 240 racks
- Modern cooling system with hot corridors (PUE=1.4)



2 redundant power sources for the site and a generator for VIL1

Equipex+ FITS



Project involving the two computer centers IDRIS and CC-IN2P3 of the CNRS, aiming to provide IT facilities to French research infrastructures

4 use cases to start : LHC experiments, LSST / Vera Rubin Observatory, French Institute of Bioinformatics (IFB), Synchrotron Soleil

Two parts in this project :

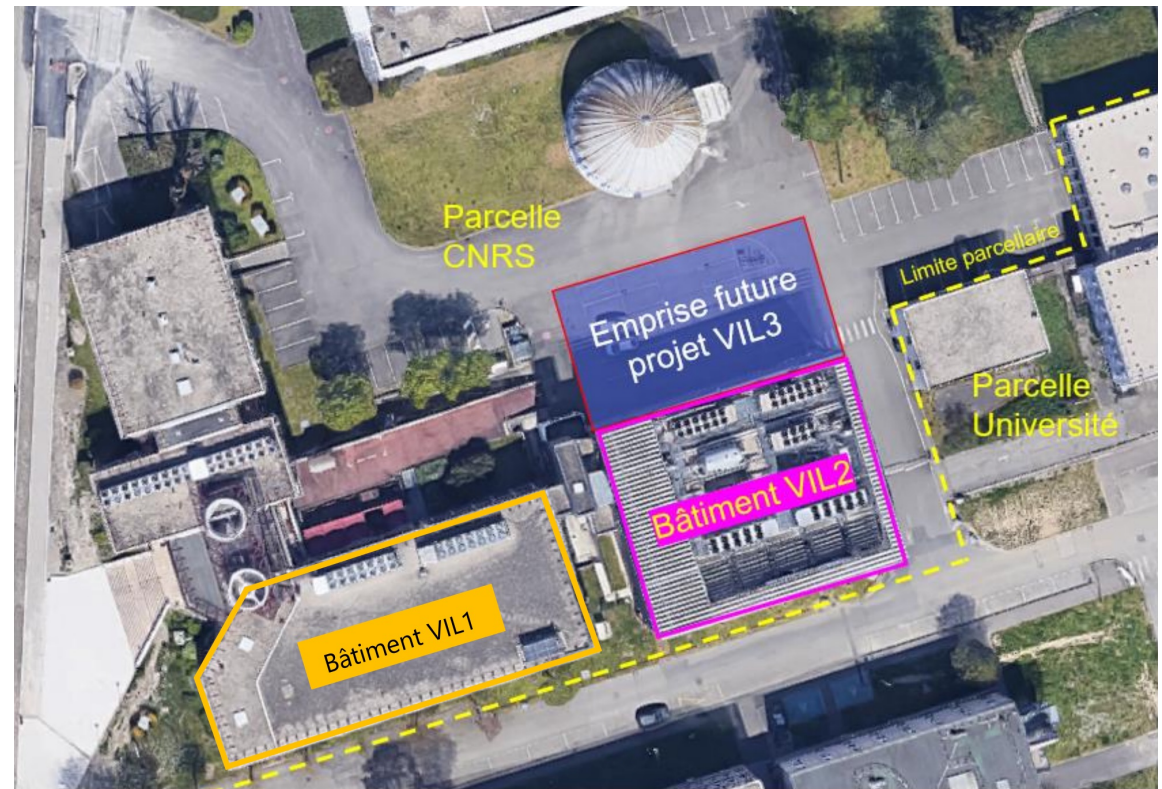

- Build a storage system distributed across IDRIS and CC-IN2P3, interconnected to identity federation systems and implementing different classes of services (double copies etc.)
- Renovation of a computer room at IDRIS and construction of a new room at CC-IN2P3

<https://www.fits.cnrs.fr/>

Equipex+ FITS



- New IT room at CC-IN2P3 : VIL3
132 racks, 550 m², 2 MW for IT
- The consortium of companies for the building has been selected, work will start on February 2024



Electricity cost evolution

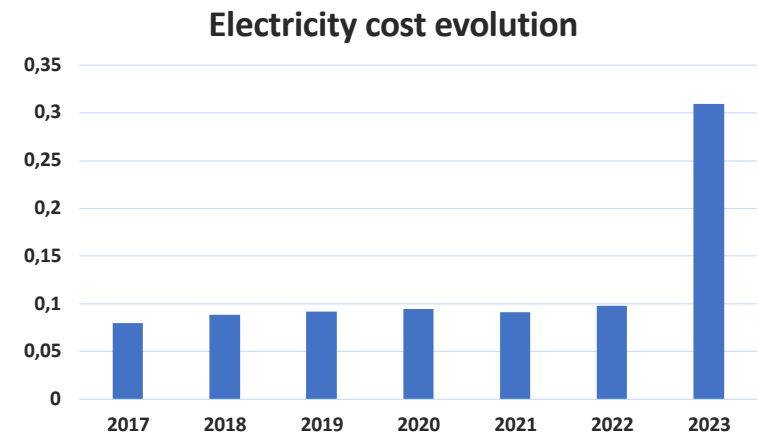


High increase of the electricity cost in 2023

- Less than 10 cents of € per kWh in 2022 to 31 cents of €

Electricity bill over years

- 1,09 M€ in 2020
- 1,10 M€ in 2021
- 1,20 M€ in 2022
- **3,80 M€ in 2023**



In comparison, the annual budget of the CC-IN2P3 is 7.3 M€ for CAPEX+OPEX

New public procurement in 2024, downward trend is expected

- Price should be lower than 20 cents of € per kWh

Computing evolutions



2019 : 1 computing platform driven by Univa Grid Engine

- HTC, HPC and GPU computing activities
- Grid and non grid jobs

2020 : 2 computing platforms

- HTCondor for HTC grid jobs
- Univa Grid Engine for non grid HTC, HPC and GPU jobs

2021 : purposes, status and plan presented at the Evaluation and Survey Committee

- Univa Grid Engine replacement scheduled on Q1 2022

Purposes :

- Resolving growing operational challenges
- Reducing financial costs associated with the escalating pricing of the UGE software license

Computing in 2023




2023 : 2 computing platforms in production

- Q1 2020 : HTCondor for HTC grid jobs
- Q1 2022 : Slurm for HTC, HPC and GPU non grid jobs


At the end : very positive

- Deployment schedule fully respected
- Some few adjustments required to completely meet expectations
- Since the progressive withdrawal of UGE, at least 260 k€ have been saved
- The team gradually gained expertise throughout the projects to reach a high level of operation

Job slots



HS06



■ HTCONDOR ■ SLURM HTC

Computing perspectives



Another way to access HTC Slurm resources

- The Jupyter Notebooks Platform is using Dask to submit Slurm jobs
- In production since April 2023
- Heavily used by LSST

Some questions about the HPC and GPU future in 2024 ?

- Hardware need to be renewed (7 years old for HPC, 5 years for GPU)
- HPC activity is very (very) low
- Except for L2IT activities, GPU activity infrequent

Evolution and perspectives

- How to reduce the power consumption : Direct Liquid Cooling study and ARM based servers evaluation
- Work in progress on jobs and resources allocation optimization

Computing resources



Platform Slurm - HTC partition for "local" activities

- **307 076 HS06, 19 904 jobs slots**
- 192 servers Dell C6525, 6 144 cores
- 68 servers HPE XL225n, 3 808 cores

Platform Slurm – GPU partition

- 18 x Dell PowerEdge C4140 / 4 x NVIDIA V100 32GB (01/2019)
- Very low activity, 2 servers have been shutdown in 2023 to save electricity

Platform HTCondor – HTC for GRID environments (EGI/WLCG)

- **500 015 HS06, 34 752 job slots**
- 216 servers Dell C6420, 4320 cores
- 240 servers Dell C6525, 7 680 cores
- 98 servers HPE XL225n, 5 376 cores

GPU in the Jupyter Notebooks Platform

- 10 x Dell PowerEdge C4130 / 4 x NVIDIA K80 24GB (06/2016)
- Transferred from UGE to JNP in 2022
- Need to be “upgraded” in 2024

GPU in the Slurm batch cluster

- 18 x Dell PowerEdge C4140 / 4 x NVIDIA V100 32GB (01/2019)
- Very low activity, 2 servers have been shutdown in 2023 to save electricity

Purchased for and by the L2IT (direct use, not in batch farm)

- 1 x Dell PowerEdge R725 – 2 x NVIDIA RTX8000 48GB (10/2020)
- 1 x Dell PowerEdge R725 – 2 x NVIDIA A100 80GB (12/2021)
- 1 x HPE ProLiant DL385 Gen10+ – 2 x NVIDIA A100 80GB (12/2022)
- High activity of research and development



Mass storage evolution



2019 : 4 Oracle SUN StorageTek SL8500 tape libraries

- Used by the backup service IBM Spectrum Scale (3 PB)
- Used by the mass storage system HPSS (80 PB)
- Faced to Oracle position on discontinuing SL8500 tape library, decision to replace all Oracle hardware by 2024

2020 : installation of the first Spectra Logics TFinity tape library

- Installation during confinement (March 2020)
- Start to move data from Oracle to Spectra Logics

2021 : purposes, status and plan presented at the Evaluation and Survey Committee

- Extend the first Spectra tape library
- Due to Oracle maintenance contract issue, speed up the data migration and prepare the purchase of a second tape library

Mass storage in 2023



2023 : 2 Spectra Logics TFinity tape libraries

- Oracle tape libraries retired on January 2023
- The primary copy (3PB) of the backup system has been transferred on disks.

Capacity and usage

- Capacity 270 PB (2 x 135) in Spectra libraries
- 150 PB currently used on tapes and continuously growing
- 2 accessors, 48 x IBM TS1160 drives, 6965 cartridges JE 20TB per library



A the end :

- 60 PB moved in the first 18 months, the remaining 20 PB were moved in 2022
- The production quite stable
- Some problems encountered with robotics on very high load
- Activity balance across both libraries required in the future

Storage perspectives



Evolutions and perspectives on storage :

- Preparing a procurement and a purchase of a third tape library during 2024
- Monitoring, adapting to guaranty performances of storage platforms both on disks and tapes
- Enhancement of the data lifetime cycle for all collaborations with a mandatory Data Management Plan
- Growing usage of Ceph for internal usage (Openstack, Kubernetes), but as a replacement of file spaces (replacement of GPFS and Isilon)

Infrastructure deployment methods



Bare metal

- Computing nodes, storage servers, support services

Cloud IAS Openstack

- R&D, support services for CC-IN2P3
- Service hosting for experiments



Cloud PAAS OKD (Kubernetes + Docker)

- R&D, support services for CC-IN2P3
- Growing demand by experiments



Operating system

- Still deploying CentOS 7 for the computing platform
- Rocky Linux / RHEL 8 & 9 for core services

Evolution of the external throughput planned on 2023, but deferred due to issues encountered by RENATER

Additional 100Gbps link for LHCOPN

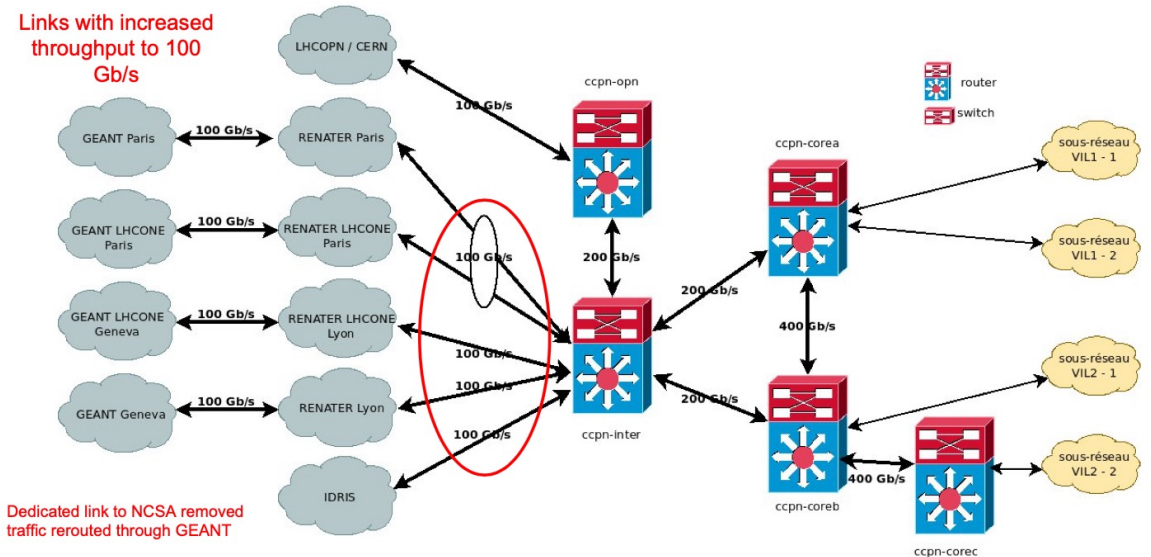
- 15/01/2024 => 200 Gbps

Additional 100Gbps link for LHCONE

- 29/01/2024 => 200 Gbps

Ready to cope with the LHC Data Challenge 2024 starting in February

- It was very tight !



More than 400 TB data are transferred each day with our partners

Summary



Many concerns about the evolution of the price of the electricity

- 2024 should be better than 2023

FITS project is ongoing on a good way for the new building and the distributed storage

What will be the future for HPC and GPU usages at CC-IN2P3 ?

- HPC requests for 2024 are low
- GPU requests seem to be more driven by the JNP platform and R&D than the batch activity

Waiting for LHC collaborations to provide OS specification for replacing the current CentOS 7 into computing farms

- EL9 should be the future, but when ?

We are ready for the coming soon data challenge 2024, this is the most important for now.



The logo for CNRS, consisting of the letters "cnrs" in a white, lowercase, sans-serif font, enclosed within a dark blue circular background.

Thank you