

# Tape libraries at CINES

05/04/2017

Jérôme Chapelle <chapelle@cines.fr>

Hervé Toureille <toureille@cines.fr>

# Agenda

- The CINES
- Applications
  - Our datacentric architecture
- The tape libraries
- Some facts
- Conclusion

# CINES

## The national computer center of french higher education

- French public organisation under the supervision of the French Ministry in charge of higher education and research.
- Provides the french public research community with computing resources and services.
- Located in Montpellier,  
55 persons : technicians, engineers  
and administratives.



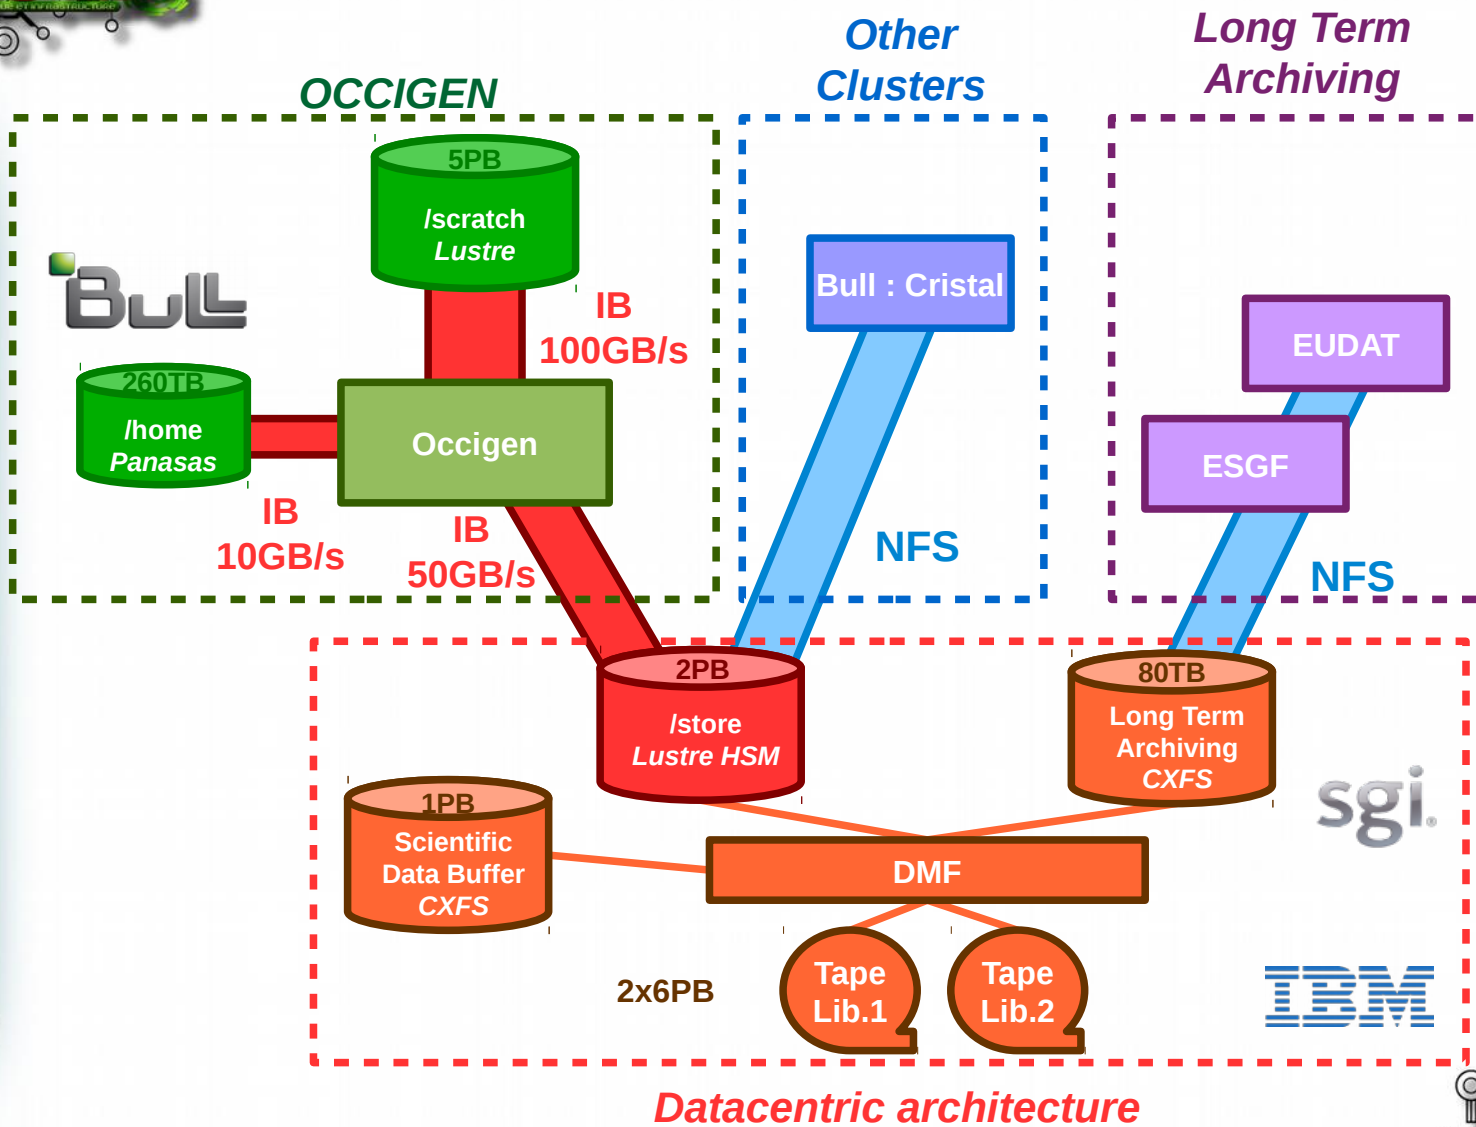
# CINES

*The national computer center of french higher education*

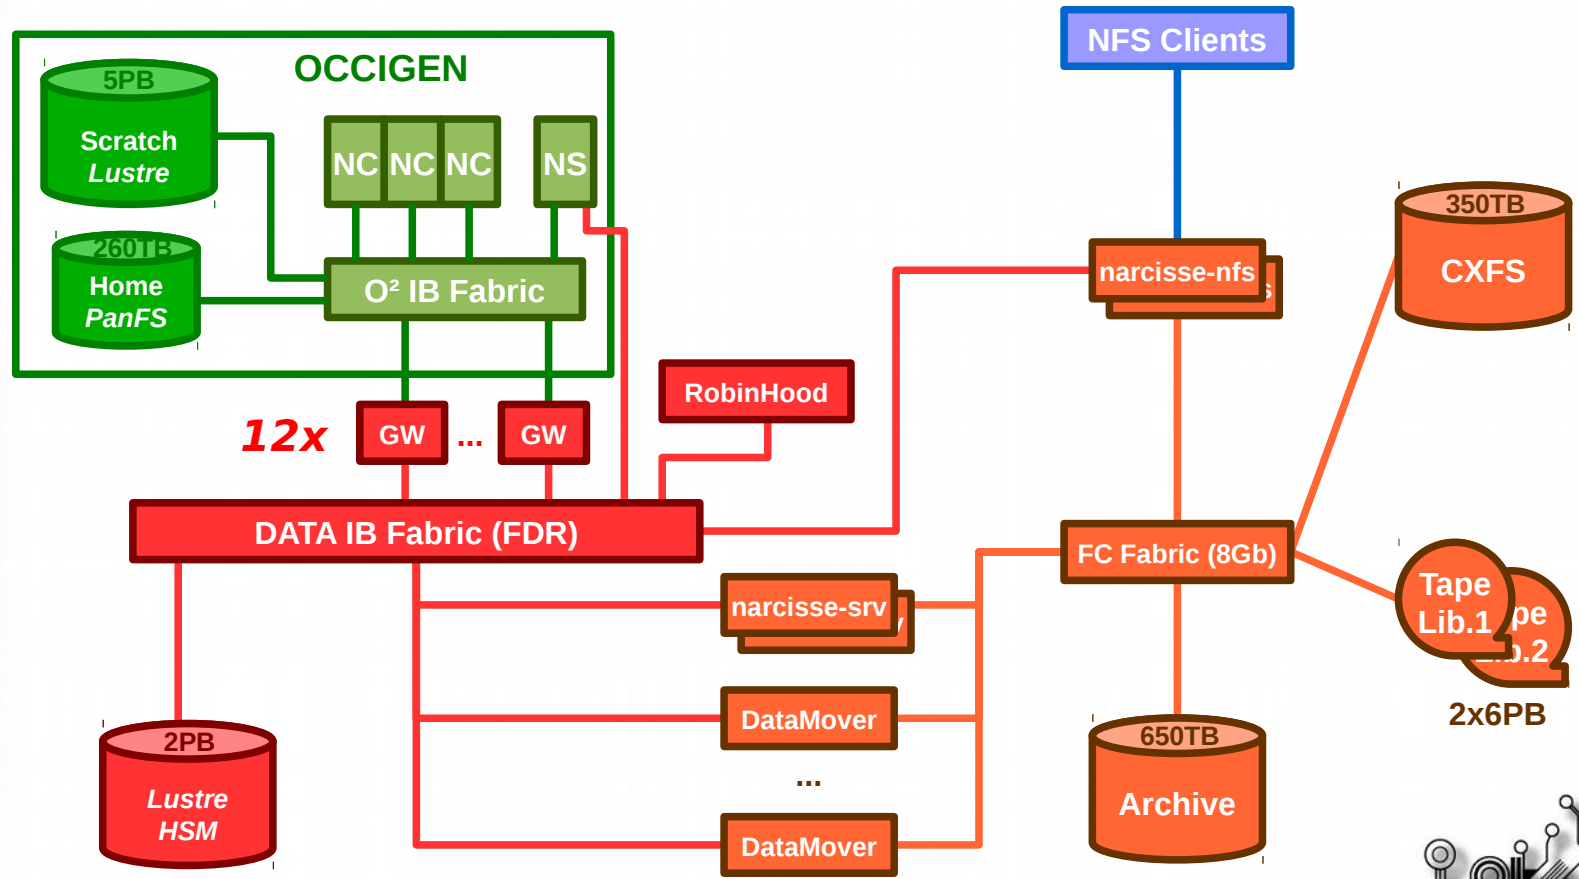
## Missions :

- High Performance computing
- Long term preservation of data and digital documents for universities and public research institutions
- Data center hosting for french national level academic institutions : 10 partnerships, 30 IT cabinets






# Current Architecture - 2018



## Applications

- Scientific Data (Lustre+DMF): 6 PB (2x11 drives)
- Backup (Bacula): 200 TB (3+2 drives)
- Longterm Archiving (Arcsys): 80 TB (1+1 drives)
- Prototypes: 80 TB (2+2 drives)




# Calcul

| Espace       | Volume Total | Volume Utilisé | Nombre Inodes |
|--------------|--------------|----------------|---------------|
| SCRATCH      | 5 Po         | 3 Po           | 250 Millions  |
| STORE        | 2 Po         | 4 Po           | 12 Millions   |
| HOME (panfs) | 260 To       | 25 To          | 13 Millions   |




## Volume



■ CT1 ■ CT2A ■ CT2B ■ CT3 ■ CT4 ■ CT5  
■ CT6 ■ CT7 ■ CT8 ■ CT9 ■ CT10 ■ CTX

## Inodes



■ CT1 ■ CT2A ■ CT2B ■ CT3 ■ CT4 ■ CT5  
■ CT6 ■ CT7 ■ CT8 ■ CT9 ■ CT10 ■ CTX

| CT                            | Inodes            | Volume (Go)         |
|-------------------------------|-------------------|---------------------|
| CT1 environnement             | 4 203 680         | 2 070 139           |
| CT2A                          | 1 171 284         | 289 577             |
| CT2B                          | 988 940           | 204 604             |
| CT3                           | 58 740            | 8 591               |
| CT4 Astronomie et géophysique | 1 599 749         | 988 310             |
| CT5                           | 543 609           | 35 906              |
| CT6                           | 26 468            | 914                 |
| CT7                           | 1 117 321         | 104 485             |
| CT8                           | 789 942           | 36 484              |
| CT9                           | 875 941           | 15 059              |
| CT10                          | 119 594           | 12 680              |
| CTX                           | 361 507           | 340 869             |
| <b>Total</b>                  | <b>11 856 775</b> | <b>4 107 622,64</b> |

## Tape Libraries

- Two IBM TS3500, with ~6PB of tapes in each one
  - Main one : designed for speed
  - Second one : designed for redundancy
- Good drive performances : from 100 MB/s to 200 MB/s
- Low TB cost + extensibility (extended twice since 2010)
- Shared between many applications (virtual libraries)



## Main Tape Library

- The main library: *"fast and serious"*
  - 10 frames
  - Drives : 9 IBM 3592E06 + 8 IBM 3592E07
  - Two accessors (robotic arms)
  - Two tiers
  - ~ 3000 JAG3 Tapes: 1 TB with E06, 1.6 TB with E07
  - ~ 600 JAG4 Tapes: 4 TB with E07 only



*Enterprise tapes allow fast access to data thanks to 32 indexes per tape.*

## Second Tape Library

- The second library: *"small but capacitive"*
  - 6 frames
  - Drives : 10 LTO4 + 6 LTO6
  - Single arm
  - Up to five tiers
  - ~ 3000 LTO4 Tapes: 800 GB
  - ~ 1500 LTO6 Tapes: 2.5 TB




## Some facts

- Average data growth = 1.6 TB/day
- Data recall up to 1.7 GB/s for big data sets (using 10 drives)
- In the last 7 years:
  - 1 failed Jaguar tape out of 3000 in the last 7 years
  - 12 failed LTO4 tapes out of 4500 in the same time
  - Only 1 unrecoverable tape (data was recovered thanks to dual copy)



## Good things

- IBM hardware is robust: it works !
- Library can be safely shared between many applications thanks to “virtual libraries”
- Hardware maintenance is fast and reliable: the support staff in Montpellier is really effective:
  - Parts are replaced quickly
  - Technicians know well the machine
- Hardware can be extended:
  - We extended twice the number of frames
  - We add drives twice



## Things that could be improved

- For IBM support libraries are “just hardware”:
  - Very hard to solve a software error (such as the RTC drift bug we had)
- Some functionalities would be useful (available in next generation)
  - Software: NTP, LDAP, SNMPv3
  - Hardware: more import/export slots
- Hard to manage a problem requiring US labs
- Frame extension process may be “tricky”



Thank you

*Questions ?*

