



CEA Supercomputing Complex Site Update

Gilles Wiber




Outline

Site update & software strategy

1. TERA+EXA

2. CCMD

3. TGCC / CCRT



CEA HPC resources for Research, Industry and Defence



1 site (CEA/DIF), 2 facilities and 4 computing centres, <https://www-hpc.cea.fr/>



TGCC
Facility
(open science)



Public research
France (GENCI) and Europe (EuroHPC)



Joliot-Curie
22 Petaflop/s
Alice Recoque
1 Exaflop/s 2026

TOPAZE
10 Petaflop/s



CEA and Industrial partners
(CCRT)



OCRE
R&D Zone



TERA-EXA
Facility
(Defence)



CEA/DAM Defence usages
(EXA)

EXA1-HE
180 Petaflop/s


EXA1-HF
36 Petaflop/s

ECRIN
2 Petaflop/s



CEA and Defence partners
(CCMD)






1 ■ TERA-EXA Site

EXA1 Partitions

- **EXA1-HF: High Frequency** partition (December 2021)
 - Focus on time to solution - mandatory for some codes
 - **Nodes: AMD Milan** (2 proc 64c, 2.45GHz), **256 GB/node**
 - 6,912 nodes
 - **Interconnect: BXI V2**, FatTree 1:2:6
 - **I/O: Lustre 1TB/s**
 - **3 containers of new modular facility**
<https://www-hpc.cea.fr/en/Modular-infrastructure.html>
 - **OCEAN** (CEA administration stack)

 **36 PFlops peak** (average 2 * previous computing power)



EXA1 Partitions



- **EXA1-HE:** High efficiency partition (2024 – 2025)
 - Focus on Energy efficiency: energy to solution (TCO optimization)
 - **672 Grace Hopper nodes in 28 racks XH3000**
 - **Interconnect:** BXI V2
 - **I/O:** Lustre 240GB/s
 - Inside the existing TERA building
 - OCEAN (CEA administration stack)



➔ **180 PFlops peak GPU + 9 Pflops CPU (Linpack: 90.79 Pflops)**

EXA1 software focus



■ RHEL9 migration

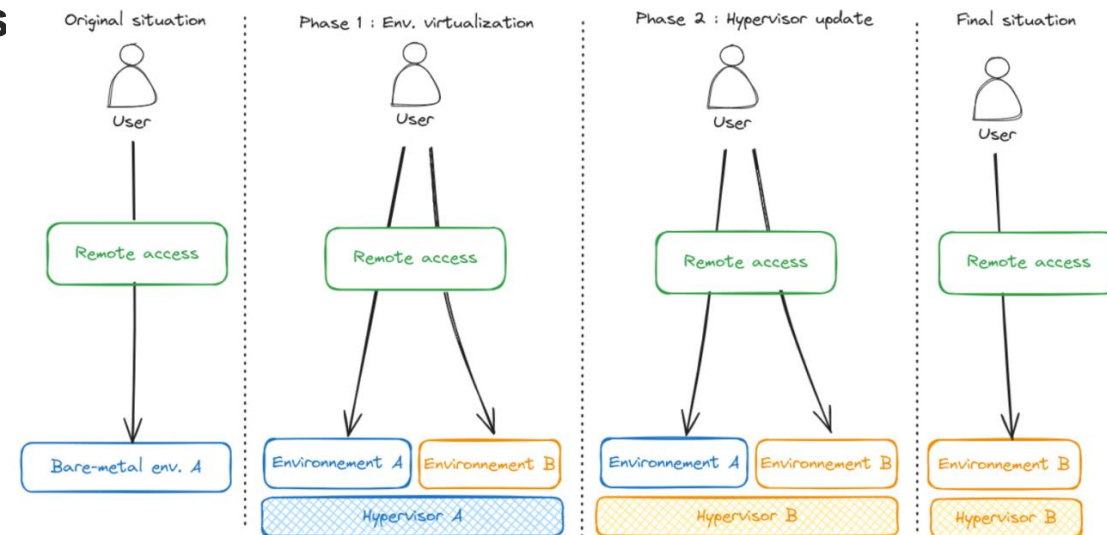
- Migration from RHEL8 to RHEL9 on Linux container (including remote visualization)
 - We are using the opensource PCOCC (our own container/virtual machine tools)
 - Try to limit user environment dependency from the OS

■ All production will be runned under Linux containers

 We get more flexibility

Strategy for all computing centre:

⇒ Using Linux containers for all our production



Network & Storage status

■ Storage:

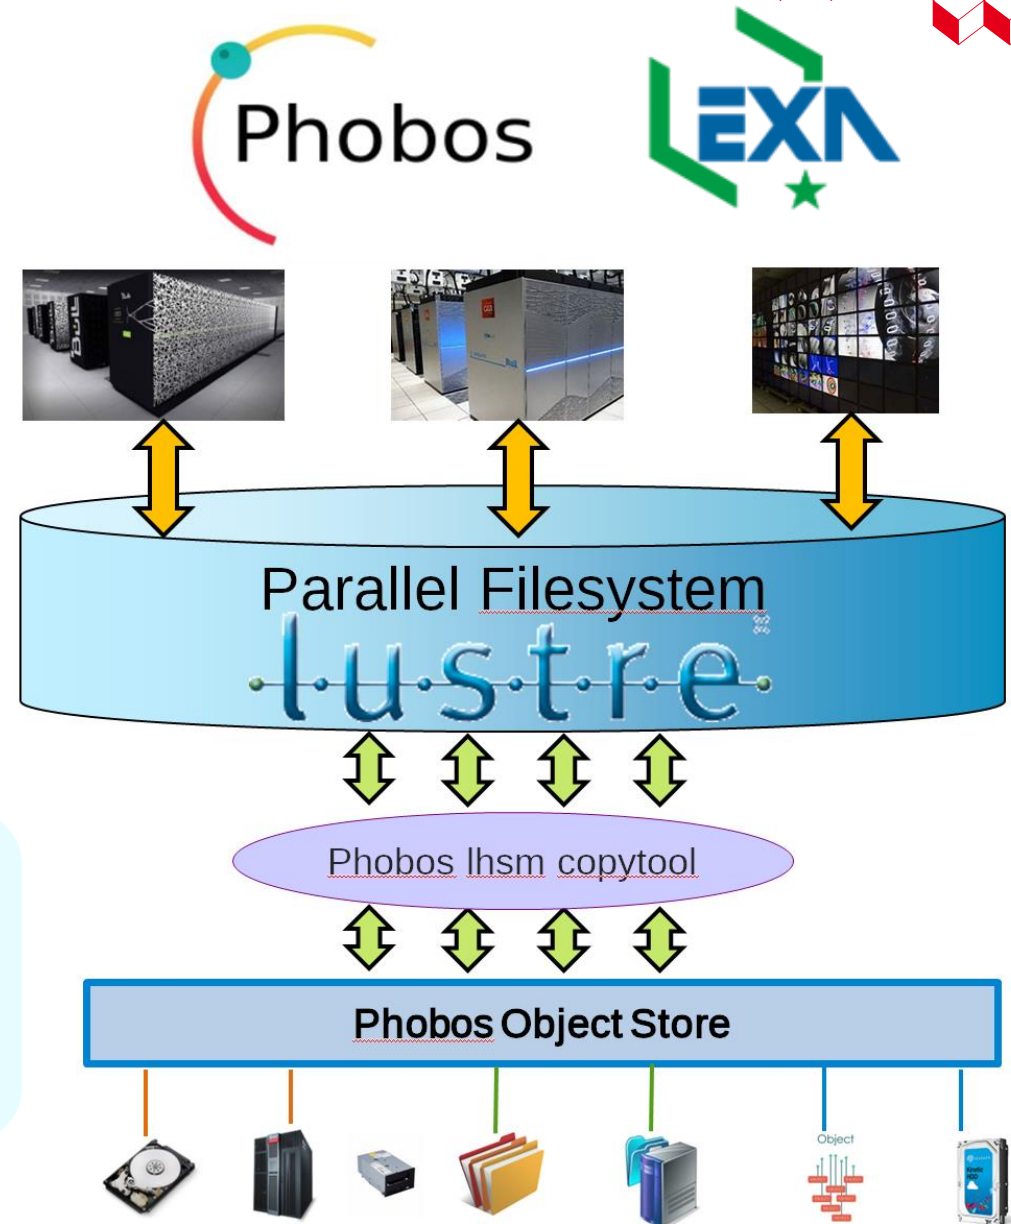
- Data migration from SL8500 to IBM TS4500 is done
- Focus on **Phobos V2** development (HPSS replacement)
 - Final developments
 - Tests have been done and it will be ready to sustain the full production by the end of 2025

■ RTHP network (connecting storage & computer):

- RTHP network is completely migrated from Infiniband to Ethernet (ROCE on Arista hardware)

Strategy for all computing centre:

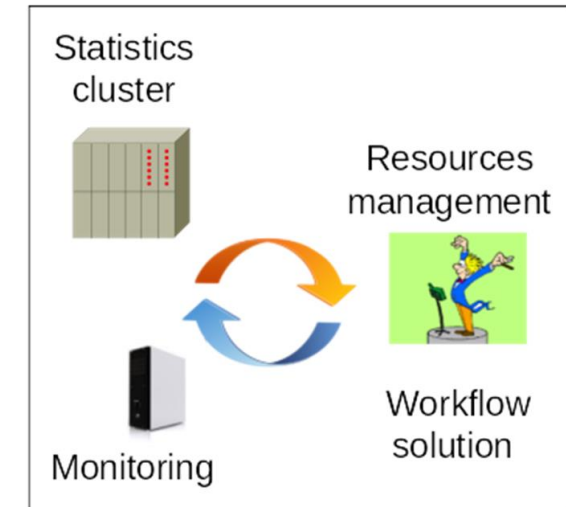
- ⇒ Data centric architecture helps us to add easily new compute partition
- ⇒ Multi-tiers storage approach fully based on opensource



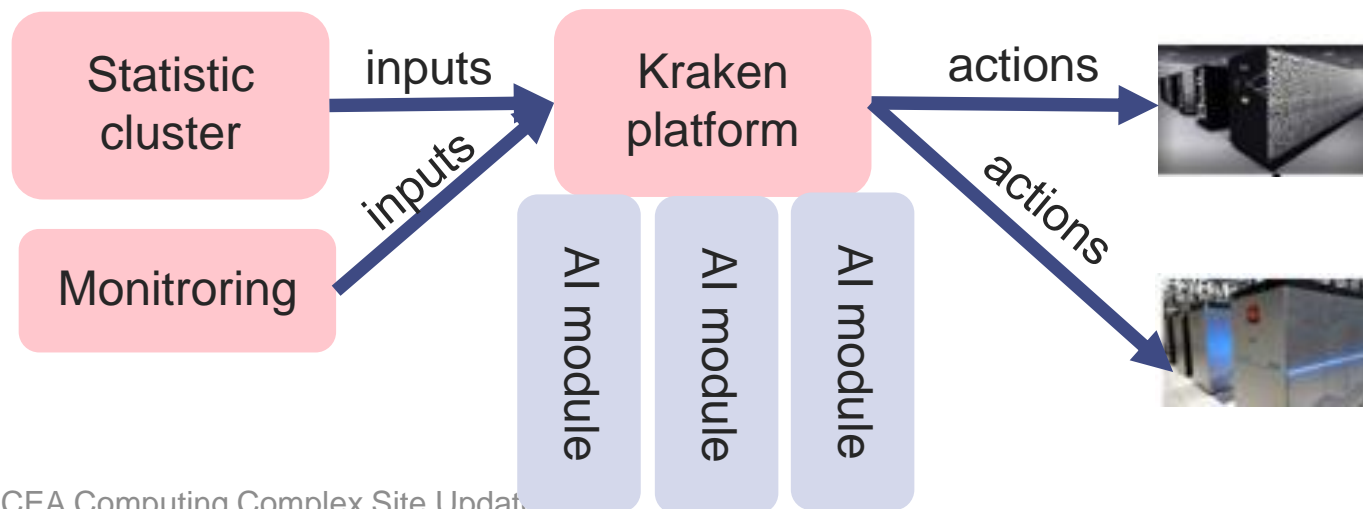
Monitoring, AI and actions (focus)



- **Statistics cluster (new)**
 - Based on OpenSearch
 - Centralization of lot of information from the computing center
- **More information**
 - Pushing tools to analyse job running




➔ **And now automatic actions through Kraken platform**



Strategy for all computing centre:

⇒ **Smart active monitoring**

2. CCMD Site



Computing Center for Confidential Runs

- Needs beyond CCRT scope (Computing Cluster for Research Technologies) partners
 - Partners needing to run more and more confidential runs (which is not possible on CCRT)

➔ Decision to create a brand new computing center (CCMD)

■ CCMD description

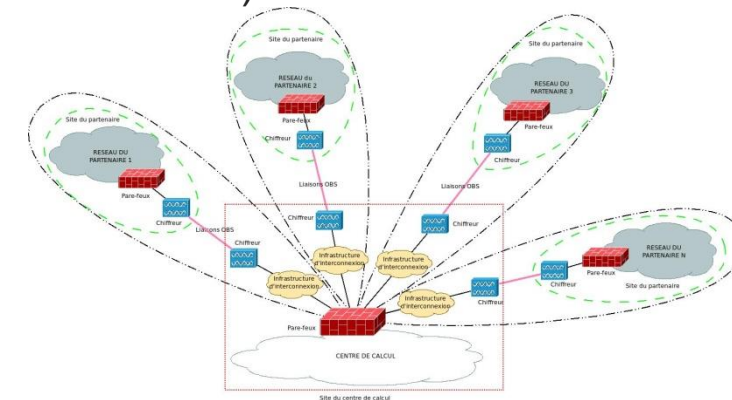
- Special architecture for connecting such network
- Computing cluster:
 - 396 nodes (dual-socket 64-core AMD Milan @ 2.45GHz, 256 GB/node)
 - 48 Nvidia nodes (dual-socket 64-core AMD Milan @ 2.45GHz, 256GB, 4 x Nvidia A100, 4 x BXI NIC)
- Storage: 5 PB (temporary and long term storage)

■ Status

- Focus on security: fully Selinux now, optimized remote visualization (sanzu)
- Fully stabilized and certified
- 4 industrial partners


Strategy for all computing centre:

⇒ Enhanced security based on Selinux and jobs' isolation







3 ■ TGCC Site




TGCC: Awaiting Alice Recoque

■ EuroHPC

- TGCC building will host the second European Exascale machine
- Focus on defining :
 - the RFP → On-going tender by EuroHPC
 - The new computing center architecture
 - All the infrastructure
- Site adaptation work has just been finished
 - Increased floor capacity, 2,8t/m²
 - +24 MW electrical power
 - +20 MW DLC warm water loop



JOLIOT-CURIE
22 Pflops peak



CCRT (Computing Centre Research and Technologies)

- CCRT = HPC club of 23 French industrial partners for more than 20 years
- We are in 5th generation of computing centre
 - ~1000 CPU nodes, 127kcore AMD Milan
 - 300 GPUs Nvidia A100@80GB
- And we still extend it (September 2024)
 - 2 racks of 16 computes blades each (3 Milan AMD bi-socket nodes with 256Go each) with HDR-100
➔ + 480 Tflops
- Workshops on GPU (presentation / training / discussions)
=> Try to push this industrial community to use GPU
- End of 2025, we will begin the 6th CCRT generation RFP
 - Larger HPC needs, both for CPUs and GPUs jobs
 - More and more AI needs for training and inference



TOPAZE (CCRT)
8.8 Pflops peak

TGCC / CCRT software focus

■ Computing centre API

- More and more interactions with other computing centre or external services (authentication from federation, remote jobs submission, ...)

=> Collaboration with RYAX company



■ Some new industrial trends: windows computing for pre/post-processing !!!


- We will push forward our opensource tool PCOCC using the capacity of VM instantiation

➡ But introducing Windows in our computing center will need a lot of work.

Strategy for all computing centre:

- ⇒ Promoting and using this computing centre API for all our tools
- ⇒ Enabling VM instantiation in our computing centre

HPC & QC Integrated Platform



FRANCE
HYBRID
HPC|Q
QUANTUM
INITIATIVE



Joliot-Curie © TGCC



Ruby (Pasqal)



Lucy (Quandela)

PERCEVAL



PULSER



Shared Software Stack

PASQAL
<HPC|Q>

Q2 2025

QUANDELA
EuroQCS-
France

Q4 2025

QPU

EuroQCS-France
GENCI/CEA



October 2025

16



Conclusion

- Four computing centres dedicated each to defense, confidential industry, industry, and research
 - specialization of services on top of a strong kernel of common methods and tools, pooling competences
- **Different directions for one global strategy**
 - Improve flexibility through PCOCC and generalization of containers' usage
 - Increase computing center instrumentation through statistics cluster especially from codes dedicated tools
 - Improve computing centre control with Kraken, a monitoring and action tools (based on AI platform)
 - Increase opensource dissemination through PHOBOS (an HPSS replacement) and Sanzu (an optimized remote visualization tools)
 - Increase security level with Selinux on and real jobs' isolation



QUESTIONS?