Introducing Jules Verne, toward a common French and European Exascale ambition

GENCI – 'GRAND EQUIPEMENT NATIONAL DE CALCUL INTENSIF'

A Very Large Infrastructure for Research (IR*)

In charge of implementing the national strategy for **high performance computing, artificial intelligence (AI), data processing** and **hybrid quantum computing (QC)**

- **18 members of staff**
- **Civil society established in 2007,** associates MESR, CEA, CNRS, FU and Inria
- **Public Operator ESR**
- **Annual budget : 32.5M€**

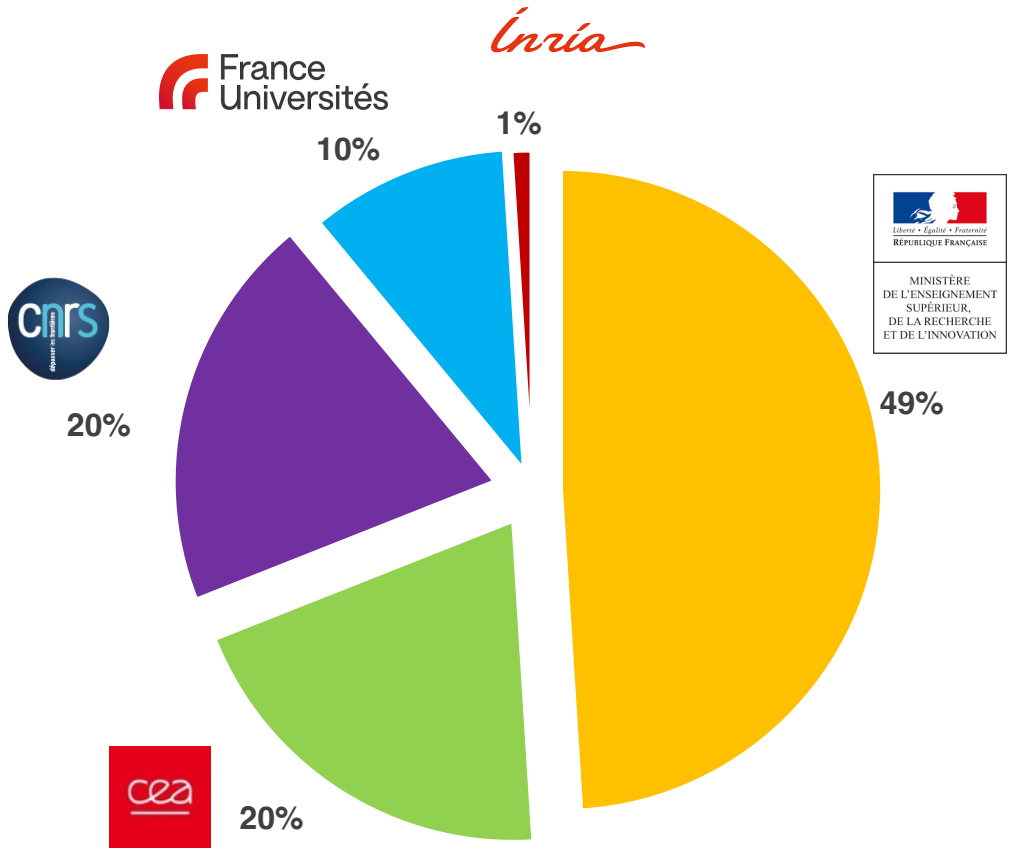

GENCI, A FRENCH HPC RESEARCH INFRASTRUCTURE

Serving yearly 1300 research projects in HPC and AI (academia, industry)

TGCC/CEA - Ile de France
- **European French node (PRACE)**
- **Candidate Hosting Site** for a future Exascale system (**EuroHPC**)
- Hosting Site for the **1st hybrid HPC + Quantum computing infrastructure** (HQI, HPCQS, EuroQCS-France)

IDRIS/CNRS - Ile de France
- **1sr converged HPC/AI system in France** *(#AIForHumanity)*
- Bring **sovereign** computing facilities / services to AI research community
- **600 yearly projects in AI**
- **> 3100 GPUs in 2023**

CINES/FU - Montpellier
- **> 70 PF** with AMD next gen GPUs and CPUs
- Available Q1 2023
- **Next step before French Exascale system**

#3 THE GREEN 500

JEAN ZAY À L'IDRIS

Une machine éco responsable : récupération de la chaleur fatale

- Signature de la convention l'EPAPS <-> CNRS le 16/02/2021

Philippe van de Maele, directeur général de l'EPA Paris-Saclay, et Alain Schuhl, directeur général délégué à la science du CNRS signent la convention CNRS/EPA Paris-Saclay visant à récupérer la chaleur fatale du supercalculateur Jean Zay. Photo CNRS IdF Gif-sur-Yvette

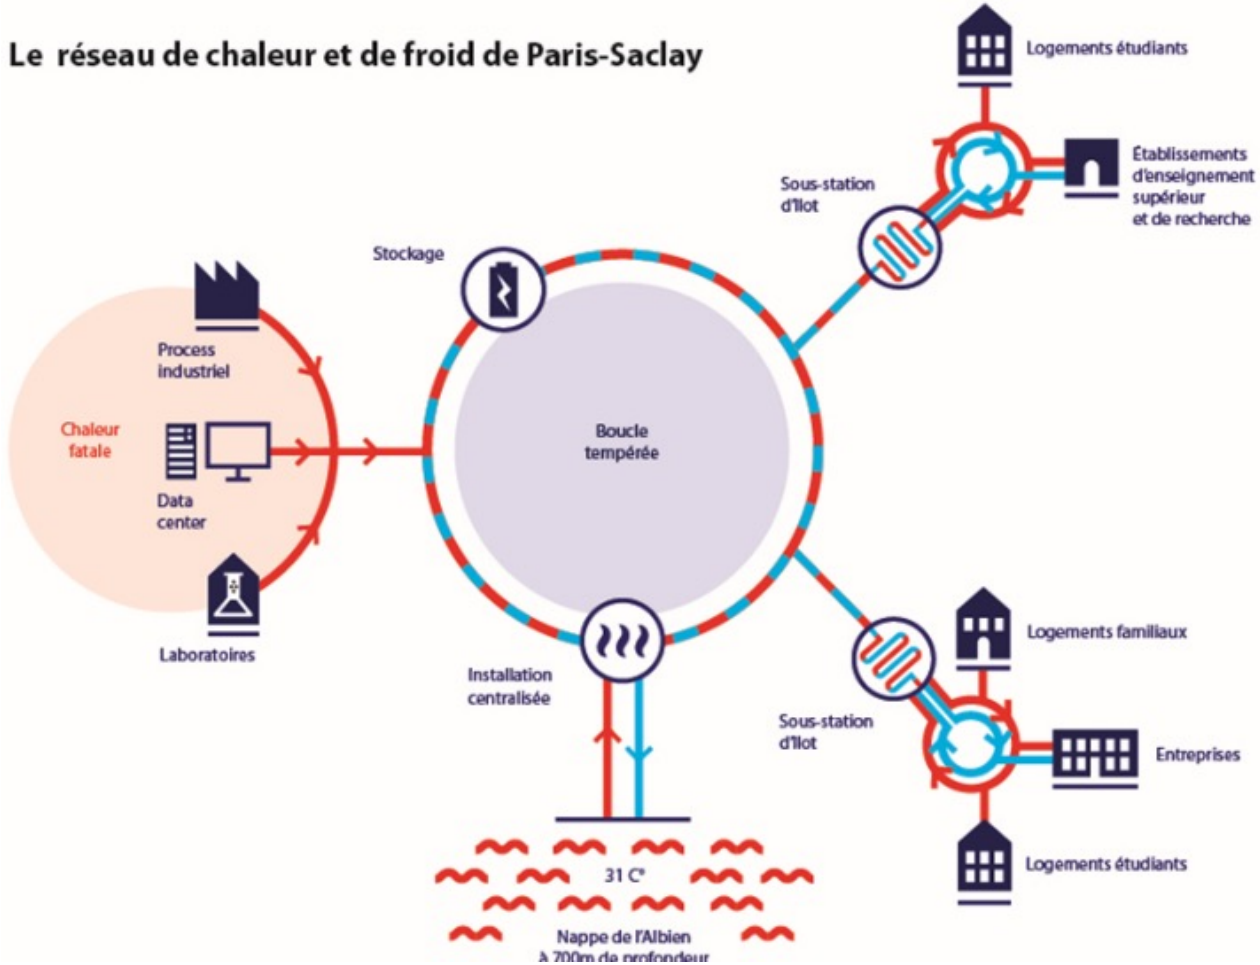

- Objectif : récupérer la chaleur de Jean Zay pour alimenter les bâtiments du Campus urbain
- 4000 MWh/an = **1 000 logements neufs**

AND SOON IN THE CLOUD WITH CLUSSTER

Cloud **U**nifié **S**ouverain de **S**ervices, de **TE**chnologies et d'infrastructu**R**es

❑ Unified sovereign and secure portal for AI (and HPC and Quantum)

- Federating existing infrastructures from private and public operators
- Continuum Open Research, Confidential Research & For Profit Activities
- Academic Sector, Industry and public authorities
- Integration European ecosystem: GAIA-X, European Open Science Cloud and Fenix

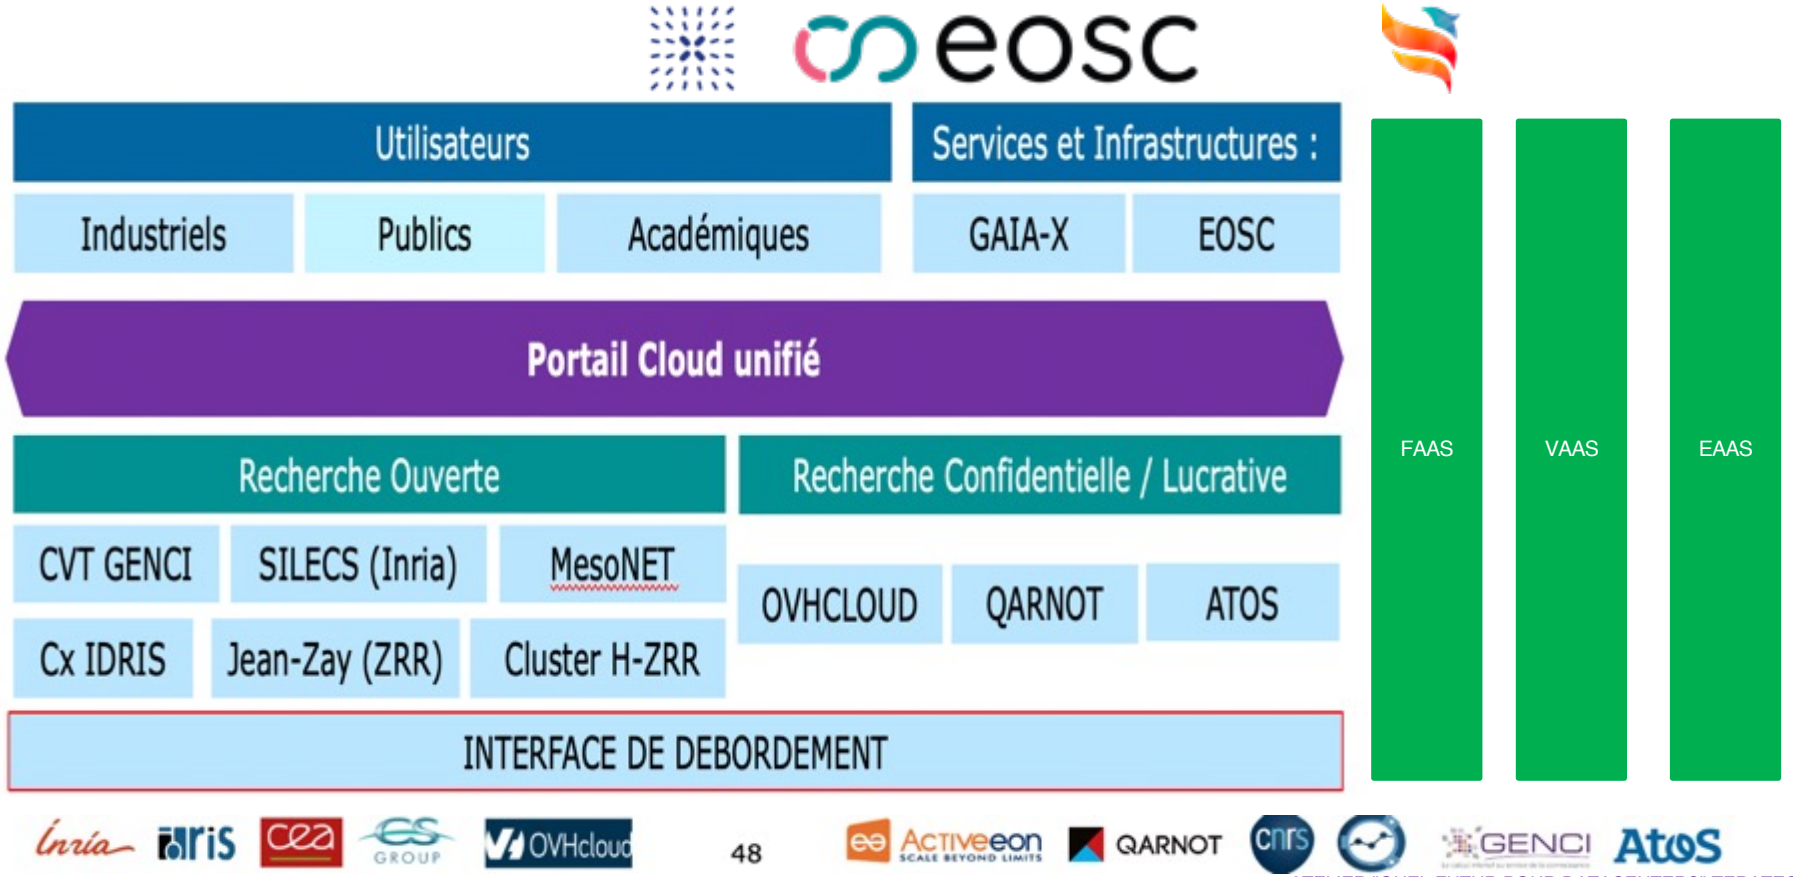

AND SOON IN THE CLOUD WITH CLUSSTER

Examples of some application-based Cloud verticals

NLP As a Service

Sovereign NLP service based on **Bloom (and others)**

- Trained on Jean Zay
- LLM (176B), open and multilingual (46 languages+13 prog. languages)
- Led by HuggingFace and >1000 researchers from academia and industry, FR and worldwide

💡 Propose a **full integrated service :**

- Training and inference
- Research and business uses
- Via GENCI, OVHCloud / ATOS

AI-based Radiology As a Service

Following COVID19 projects, wish to AP-HP to access to sovereign HPC/AI services for AI-guided radiology as a service.

💡 Propose a **full service** :

- Relying on the Bernoulli joint laboratory between Inria & AP-HP
- Publish into CLUSSTER AP-HP processing workflows
- Address booth standard and sensible data (HDS)
- Via GENCI, OVHCloud / ATOS

GaiaData As a Service

💡 Extend spatial processing vertical (from CS Group) for integrating data services from the **GaiaData** project:

- Started in March 2022
- Joining CLIMERI-France, PNDB, DATA TERRA + 21 partners
- Data portals, multi-sources processing of data from Universe and Earth Sciences
- Publish services into CLUSSTER and give access to GENCI ressources for (re) processing of data using AI

RISE OF QUANTUM COMPUTING WITH THE HQI INITIATIVE

FIRST SERVICES AVAILABLE SINCE THE END OF 2022 AT TGCC

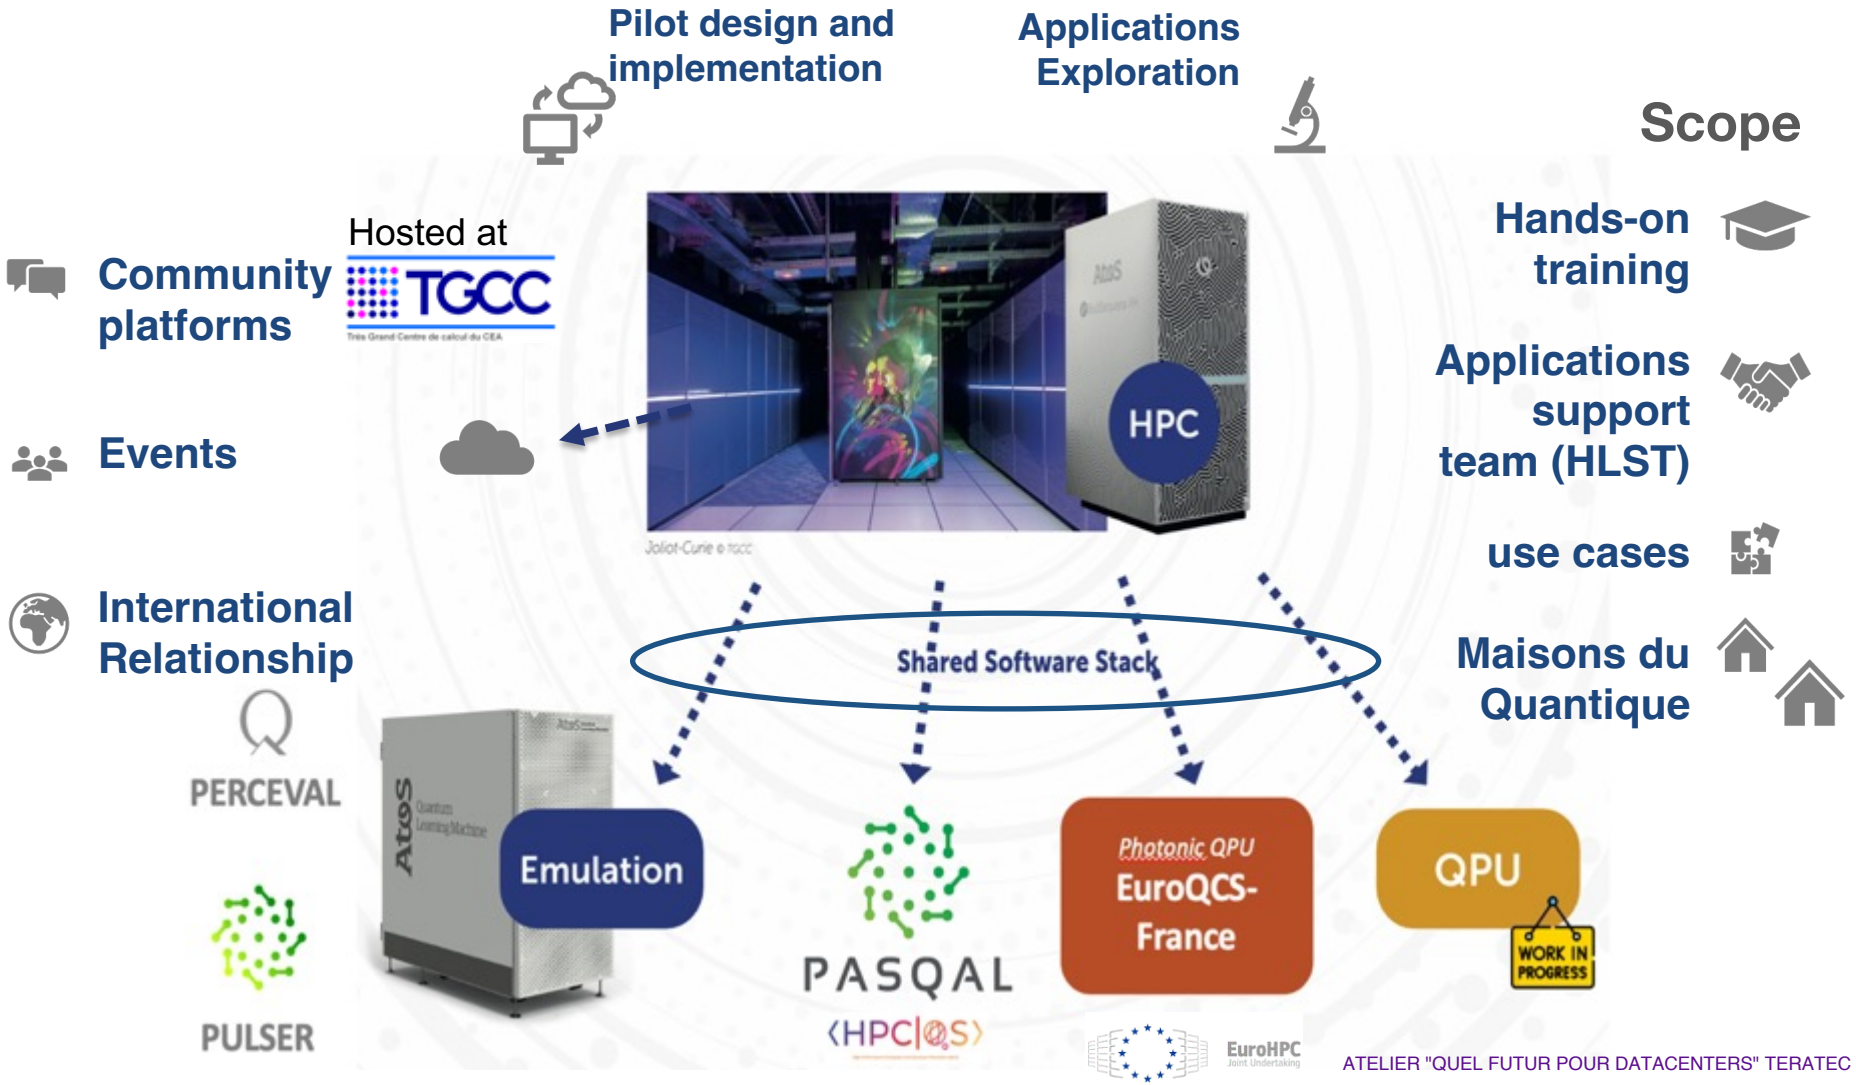

FEW RECENT RESULTS

In academia and industry

DGEN-380 engine Large Eddy Simulation at take-off conditions

by C. Pérez Arroyo, J. Dombard, F. Duchaine, L. Gicquel, N. Odier and G. Staffelbach (2022)

AVBP

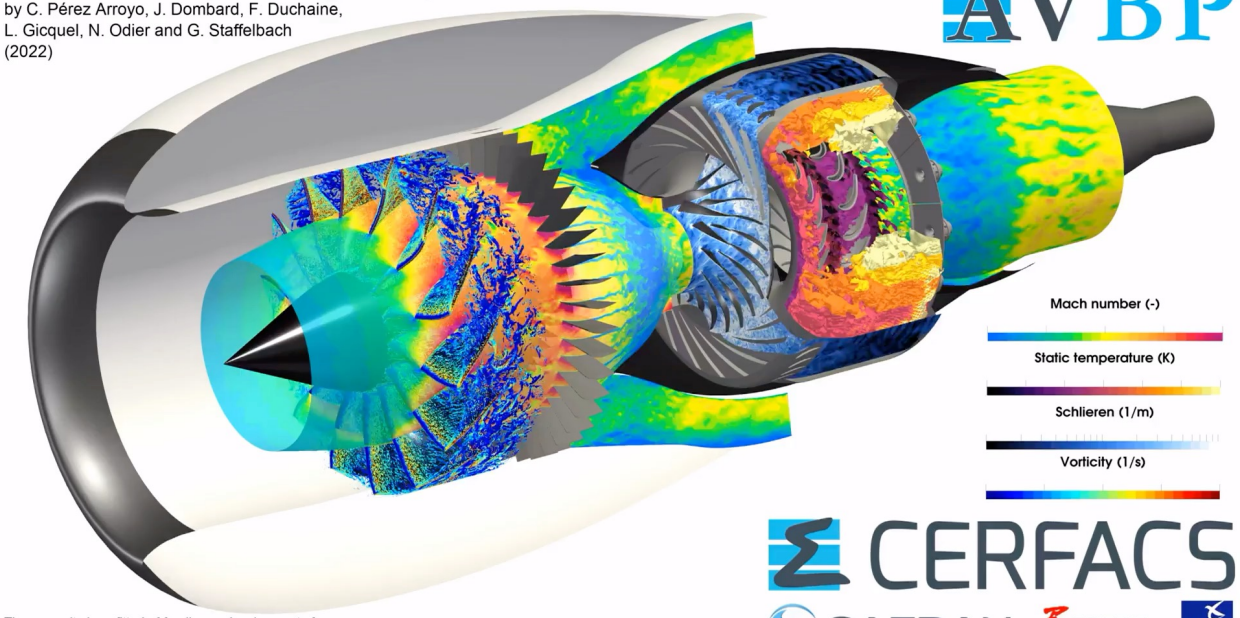

CERFACS

SAFRAN AKIRA dgac

These results benefitted of funding or developments from:
project ATOM (DGAC/SafranTech No 2018-39), PRACE (20th Call Project Access FULLEST), EXCELLERAT (H2020 823691), EPEEC (H2020 801051) and GENCI (A0122A06074).

World first-ever :
Full engine combustion model using 13k cores on Joliot Curie

One of the longest simulation of M^{pro} protease of SARS-Cov2 on Joliot Curie

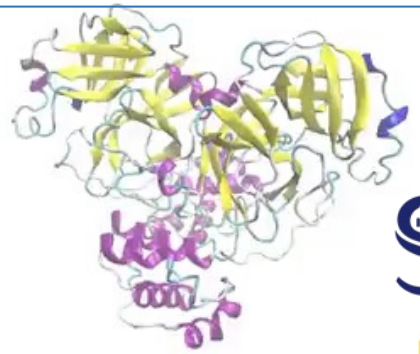

SORBONNE UNIVERSITÉ

Qubit pharmaceuticals

Update: Introducing The World's Largest Open Multilingual Language Model - BLOOM 🌸

- Training on Jean Zay of the biggest open NLP model
- Global collaboration (>1000 researchers), 46 natural and 13 programming languages
- 176B parameters, more than 400 GPUs used

FEW RECENT RESULTS

Assessing the evolution of the climate : French simulations for CMIP6

- **GENCI supported IPSL teams for their (post)CMIP6 commitment**
  - >500 million core hours allocated since 5 years on Curie and Joliot Curie **@TGCC**
  - A dedicated storage infrastructure of 14 PB and user support
  - Specific optimisation of the workflows as well as each individual model
- **As a result**
  - More than 80 000 years simulated
  - First results disclosed in September 2019, France (IPSL and MeteoFrance) **as first contributors** of the IPCC community for CMIP6
    - Confirmation of the acceleration of global warming
    - Better view at regional scales
    - Up to +7° in worst case scenarios

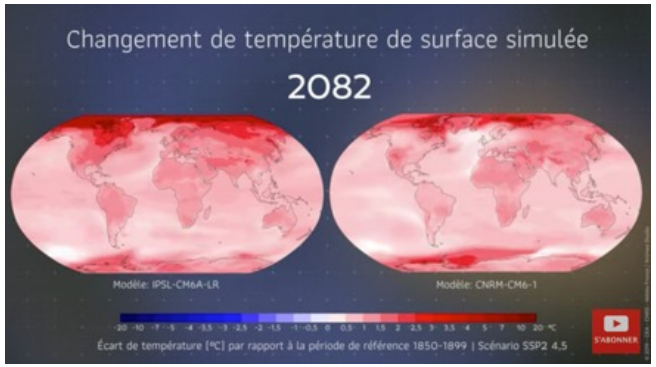

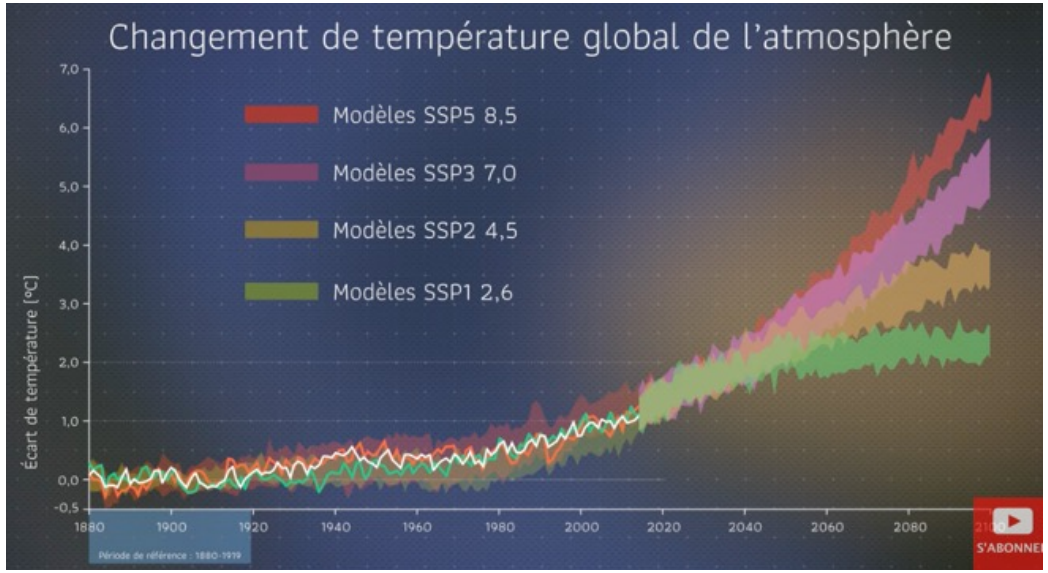

- IPCC full report published in 2021

FEW RECENT RESULTS

ANR Project « Du Carbone à l'Or Olympique »

- Call for proposals « Sport de très haute performance »
- 6 scientific partners
  - ESPCI Paris-PSL (led), École Navale, École Nationale de Voile et des Sports Nautiques, Ifremer, University of Nantes, and the LadHyX laboratory (CNRS, Ecole Polytechnique)
- And associated ones
  - RiseSailing, AIM45, GENCI and Institut Aérotechnique de Saint-Cyr
- Objective
  - Interdisciplinary project bridging R&D teams in CFD, structure mechanics and humanities (cognitive ergonomic in sports)
- 500 000 core hours allocated on GENCI' HPC facilities

CFD simulations in aero/hydrodynamics coupled with studies in structure mechanics

FEW RECENT RESULTS

Europe and France are in the game 🚀

🥇

ACM

Gordon Bell Prize

2022

PRESENTED TO

Luca Fedeli, Kevin Gott, France Boillod-Cerneux
Thomas Clark, Axel Huebl, Conrad Hillairet,
Stephan Jaure, Adrien Leblanc, Rémi Lehe, Neil Zaïm
Andrew Myers, Christelle Piechurski, Henri Vincenti
Mitsuhisa Sato, Weiqun Zhang, Jean-Luc Vay

FOR Pushing the frontier in the design of laser-based electron accelerators with groundbreaking mesh-refined Particle-In-Cell simulations on Exascale-class supercomputers

ACM PRESIDENT | AWARDS COMMITTEE CO-CHAIR | AWARDS COMMITTEE CO-CHAIR

NEXT STEP IS EXASCALE

A major stake for European research and innovation

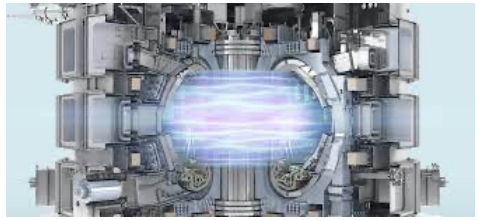

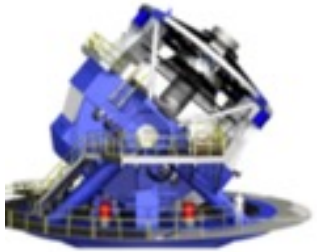

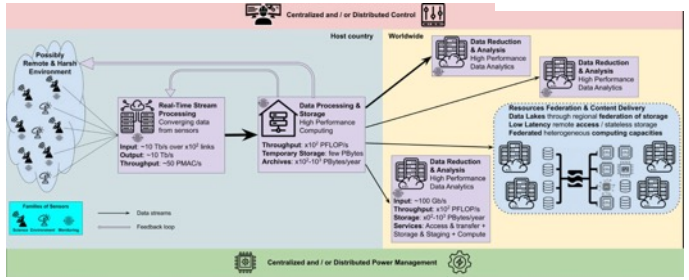

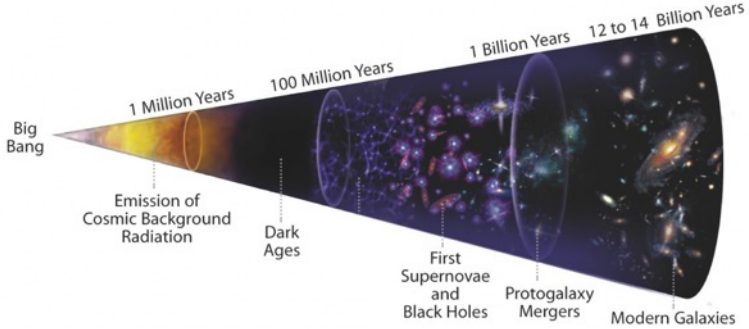

SAFRAN

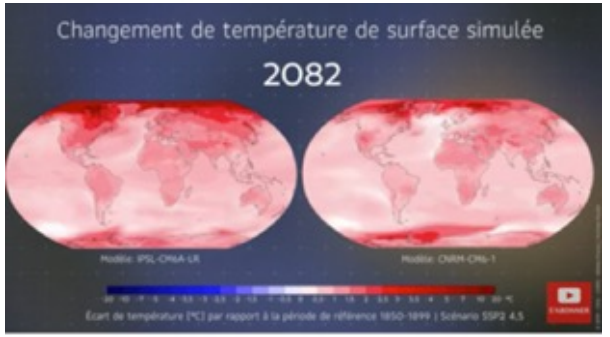

FOR EXASCALE = HYBRID ARCHITECTURES IS A GENERAL TREND

In Europe, US or Japan (but not in Japon – 40MW, 0.5EF)

428 PF peak, 309 PF (LUMI-G),
4096 compute nodes: 62% hybrid,
38% scalar **(#3 top500**)

Leonardo

255 PF peak, 174 PF HPL
4992 compute nodes:
69% hybrid, 31% scalar **(#4 top500**)

Marenostrum 5 (314 PF peak) to come with similar ratios

1.68 EF peak, 1.2 EF HPL **(#1 top500**)
9472 nodes, 100% hybrid, 21MW

>2 EF peak, > 9000 nodes, 100% hybrid (end 2023 ?)

>2 EF peak
Q1 2024

Sunway OceanLight
1.3 EF peak, 1.05 EF HPL (?)
>107k nodes, 42M scalar cores but 35MW (!)

NUDT Tianhe-3
1.7 EF peak, 1.3 EF HPL (?)
> Phytium 2000+ ARM + Matrix 2000+ MTP accelerators

EUROHPC

EC + 33 member States + 3 private partners

#EuroHPC
Joint Undertaking

The European High Performance Computing Joint Undertaking (EuroHPC JU) will pool European resources to develop top-of-the range exascale supercomputers for processing big data, based on competitive European technology.

Member countries are Austria, Belgium, Bulgaria, Croatia, Cyprus, Czech Republic, Denmark, Estonia, Finland, France, Germany, Greece, Hungary, Iceland, Ireland, Italy, Latvia, Lithuania, Luxembourg, Malta, the Netherlands, North Macedonia, Norway, Poland, Portugal, Romania, Serbia, Slovakia, Slovenia, Spain, Sweden and Turkey.

EuroHPC Joint Undertaking

ETP 4 HPC — EUROPEAN TECHNOLOGY PLATFORM FOR HIGH PERFORMANCE COMPUTING

BDV BIG DATA VALUE ASSOCIATION

QIC European Quantum Industry Consortium

EUROHPC, MAIN DIRECTIONS FOR #2 REGULATION (2021-2027)

2 main technological directions, 7B€ for the second phase

EuroHPC Joint Undertaking

➢ **Infrastructure deployment**

#EuroHPC (high performance computing) Joint Undertaking

	2019 & 2020	2021	2022	2023	2024	2025	2026	2027
HPC Infrastructure	pre-exascale + petascale HPC systems	Several pre-exascale systems and exascale HPC systems				exascale and post-exascale HPC systems		
Quantum Infrastructure	Quantum simulators interfacing with HPC systems	1st generation of quantum computers + quantum simulators interfacing with HPC systems				2nd generation of quantum computers + quantum simulators		

1st Exascale CFEI published by EuroHPC on dec 2021
→ Selection of the German proposal (at FZJ) called JUPITER
→ 2nd Exascale CFEI published on **14 December 2022**

1sr CFEI for quantum computers from EuroHPC
→ EuroQCS-France application (with FR, GE, IE et RO) using photonics based solution integrated into HPCQS

➢ **Successive R&D calls for proposals including :**

- **an Exascale pilot** – Consortium EUPEX (lead Atos)
- A quantum pilot – Consortium HPCQS

EU PEX

HPCQS

JULES VERNE : THE FRENCH LED EXASCALE PROJECT

Organization of the french application

- GENCI *Hosting Entity*
- CEA *Hosting Site*
- SURF (NL) as member of consortium

Name of the consortium/supercomputer : Jules Verne

Full TCO over 5 years : 500 to 600M€ (50% EuroHPC, 50% consortium)

- French public contribution
- NL contribution

-Seeking more partners on the consortium to reach 250M€

- **International partners**
- French research institutions
- French industrial partners (end users)

OUR VISION FOR AN EUROPEAN EXASCALE MACHINE

Addressing societal and scientific challenges (such as universe sciences, climate change, health, new energy, innovative materials, transport or smart cities/systems) via large scale numerical simulations and massive data analysis using artificial intelligence

- An accelerator of European Science and Innovation
  open to all scientific and industrial collaborations, supporting new services including Cloud based interactive supercomputing / visualisation, containerisation and urgent computing for fast decision making
- A converged HPC/HPDA/AI system with a modular, balanced and energy efficient architecture
  based on accelerated, scalar and HPDA partitions within a tiered data centric infrastructure
  integrating state-of-the-art post-exascale quantum accelerators and related services for specific workloads
- A system fully embedded inside the digital continuum
  ready for secured end-to-end workflows from instruments / edge devices to long term sovereign storage
- A system with European Technology and Skills
  integrating European hardware / software technologies in terms of computing, storage, network, infrastructure, middleware, applications with global support of AST to engage/support communities.

A system ready to harness European technologies and the best breed of opensource software in a highly secure environment

EuroHPC Joint Undertaking

#NextGenerationEU

EXASCALE SYSTEM ARCHITECTURE OVERVIEW

Possible reference designs

Large scientific instruments

Academia, industrial and public services users

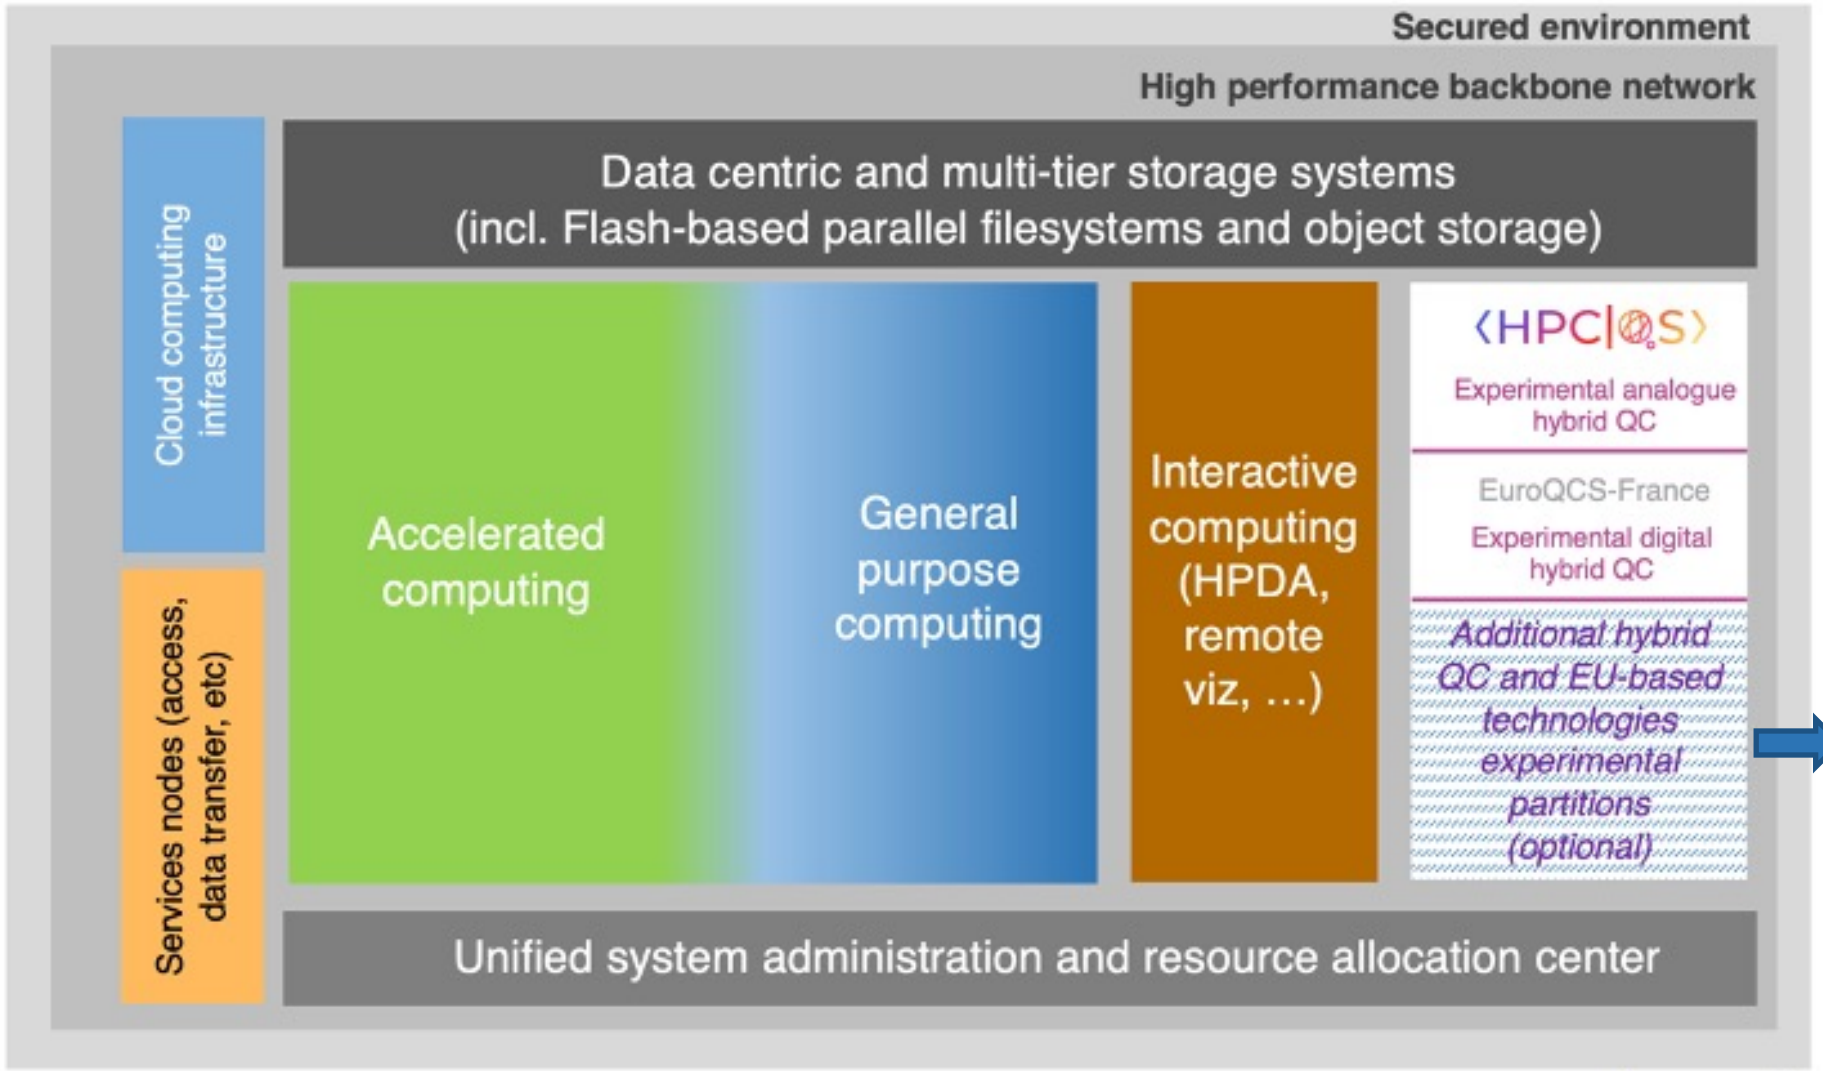

RISC-V ?
Cryo/ neuromorphic computing ?
Immersive cooling ?
Photonics ?
NDA storage ?
>1000 qubits or pre LSQ ?

15 to 18 MW (HPL) expected

EXASCALE AS A KEY APPLICATION ENABLER FOR EUROPEAN SCIENTIFIC AND SOCIETAL CHALLENGES

Fundamental Research

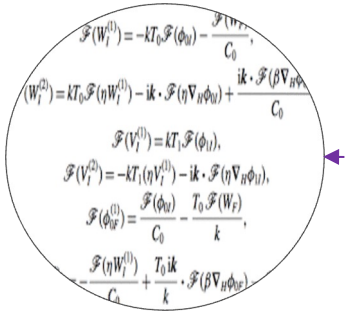

Universe Sciences

Life Sciences

Energy

Industry of the Future & Digital Transformation

Digital Sciences and usages

First French Exascale Scientific Case (>70 apps)
https://hal.archives-ouvertes.fr/hal-03736805/document

Toward Exascale applications and usages : The NumPEx project

Aggregate the French HPC/HPDA/IA community

Contribute and accelerate the emergence of **a European sovereign exascale software stack** and **strategic applications exascale capability** in a **coherent and multi-annuel framework**

Integrate and validate **co-designed** innovative methods, libraries and software stack with demonstrators of strategic applications.

Accelerate science-driven and engineering-driven developers **training and software productivity**

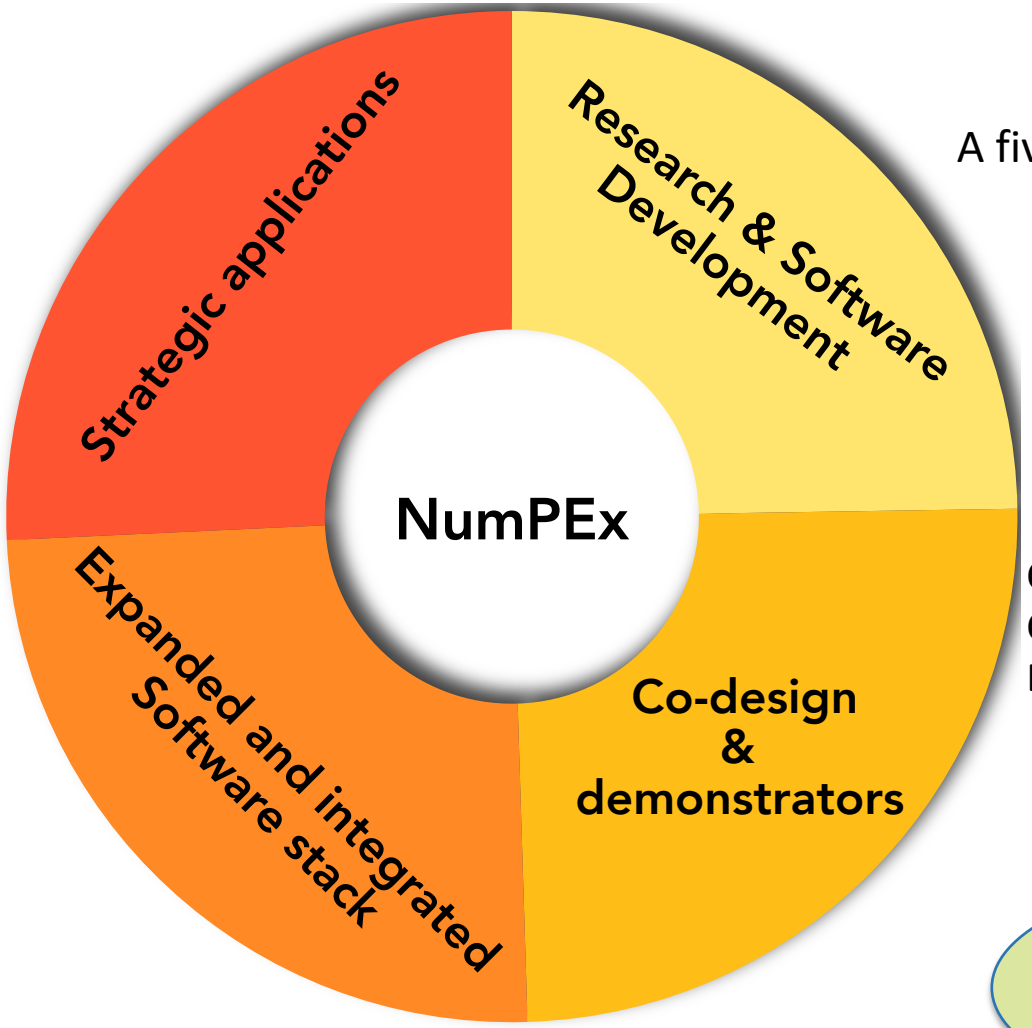

5 Years
40,8 M€

A five-year, € 40,8 M budget

Core Research Institutions

Core Research Institutions: CNRS, CEA, INRIA, Universities, Engineer schools, Industry

80 R&D teams
500 Researchers

JULES VERNE PROPOSAL – NEXT STEPS

From Call for Expression of Interest (CEI) to commissioning the Exascale supercomputer

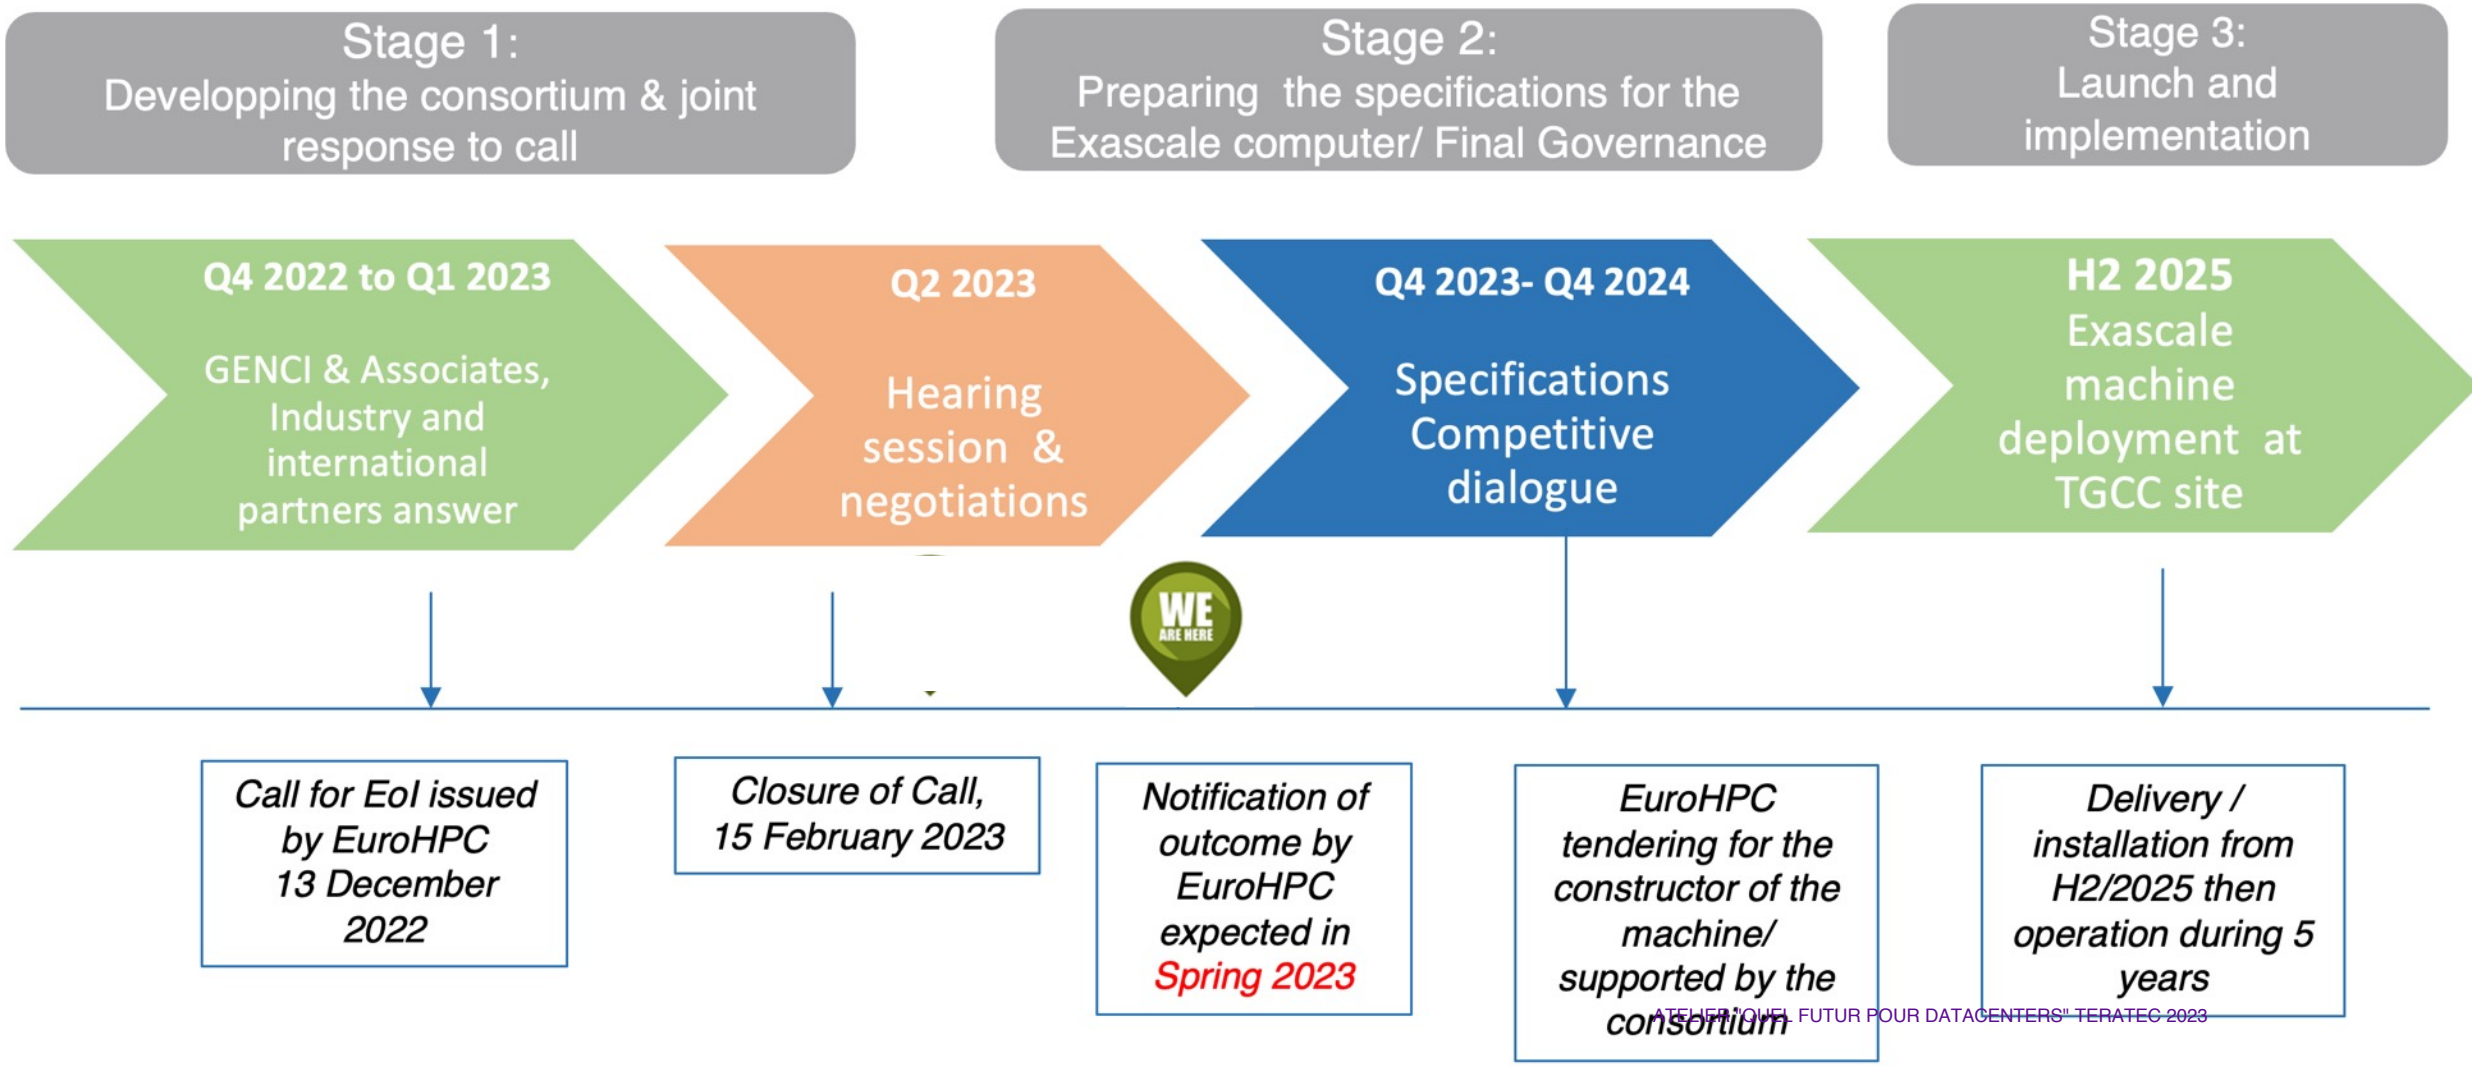

Let’s foster science and innovation together 🚀