




# **France within the international Exascale ecosystem**

France.boillod-cerneux@cea.fr



# From petascale to the exascale



EuroHPC  
Joint Undertaking



France  
Universités



## TGCC/CEA - Ile de France

- Future hosting site for one of the EuroHPC European Exascale machines
- Site hosting the 1st brick of a federated and integrated quantum infrastructure with an HPC infrastructure (HPCQS)



## IDRIS/CNRS - Ile de France

- 1st AI machine in France in response to the #AIForHumanity plan
- Bringing sovereign power to French research in AI
- 350 annual AI projects
- > 3100 GPUs at the start of 2022



## CINES/CPU - Montpellier

- > 70 PF with next-gen AMD CPUs and GPUs
- Available early 2023
- Last big step before the exascale for France

# EuroHPC exascale machines: next step for France



- Organization of the french application
  - GENCI is the Hosting Entity
  - CEA is the Hosting Site
  - SURF (NL) is member of the consortium
- Full TCO over 5 years : 500 to 600M€ (50% EuroHPC, 50% consortium)
  - French public contribution
  - NL contribution

Global performance targets for the supercomputer: Sustained HPL performance > 1 EFlops

Target composition :

- 60% accelerated nodes and 40% scalar nodes but accelerated nodes will bring > 90% peak performance
- >100 PB Flash/HDD and > 200 PB archive
- Power consumption < 20 MW
- European technology!

## #EuroHPC (high performance computing) 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.

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



# 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 and balanced 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
  - integrating European hardware and software technologies in terms of computing, storage, network, cooling and infrastructure monitoring

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

# TIMELINE FOR THE Exascale PROJECT



Stage 1:  
Developping the consortium & joint  
response to call

Stage 2:  
Preparing the specifications for the  
Exascale computer/ Final Governance

Stage 3:  
Launch and  
implementation

Joint response GENCI &  
Associates, Industry and  
international partners  
Q4 2022 to Q1 2023

Hearing  
session &  
negotiations  
Q2 2023

Machine  
specifications/  
Competitive  
dialogue  
Q4 2023 – Q4 2024

Exascale machine  
launched at  
TGCC site  
H2 2025

Call for EoI issued  
by EuroHPC  
**13 December 2022**

Closure of Call,  
**15 February 2023**

Notification of  
outcome by  
EuroHPC  
expected in  
Spring 2023

EuroHPC tendering  
for the constructor  
of the machine/  
supported by the  
consortium

Delivery /  
installation from  
H2/2025 then  
operation during 5  
years

**Finalisation of the French-led consortium : May 2023**

# EuroHPC exascale machines: next step for France



- Organization of the French application
  - 3 “sub-projects”
    - Sub project 1: in charge of legal aspects
    - Sub project 2: in charge of design and integration of the exascale supercomputer
    - **Sub project 3: in charge of the applications**

Focus on applications

# Team dedicated to the applications



France  
Boillod-Cerneux



Michel  
Daydé



Anne  
Laurent

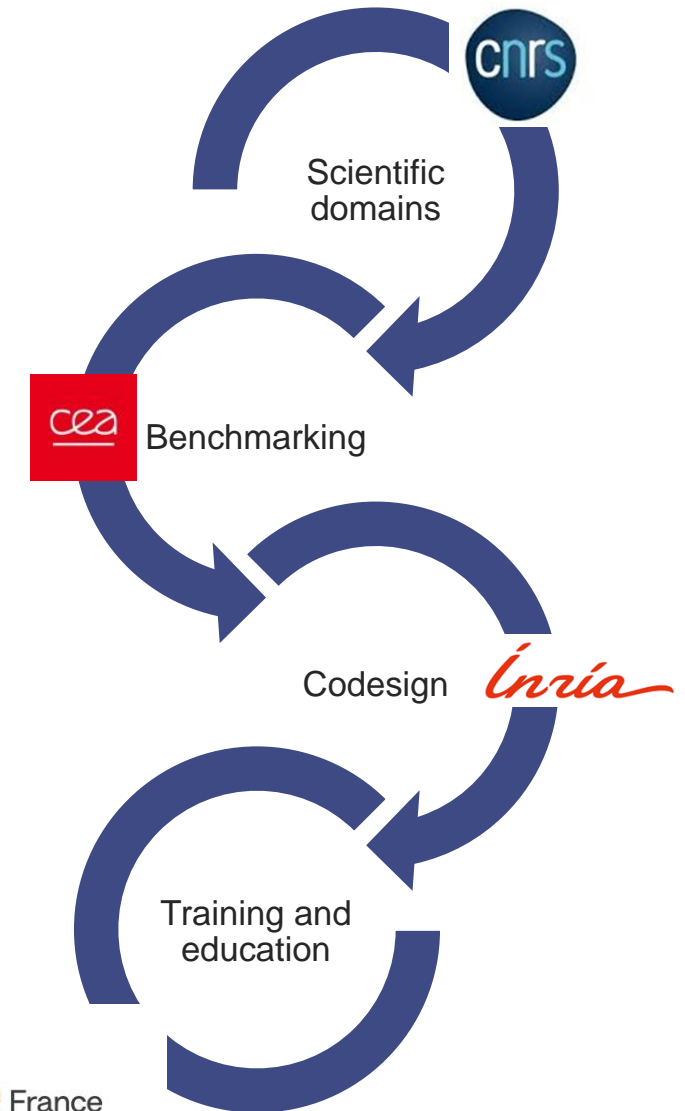


Bruno  
Raffin

A team, representing the application community within the exascale project and which brings together GENCI partners

# Our missions

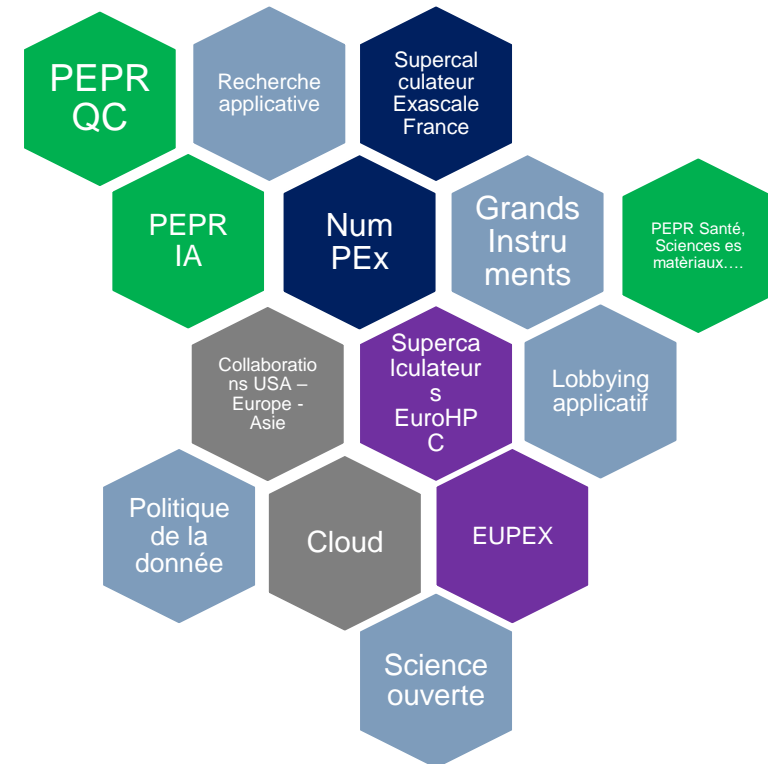
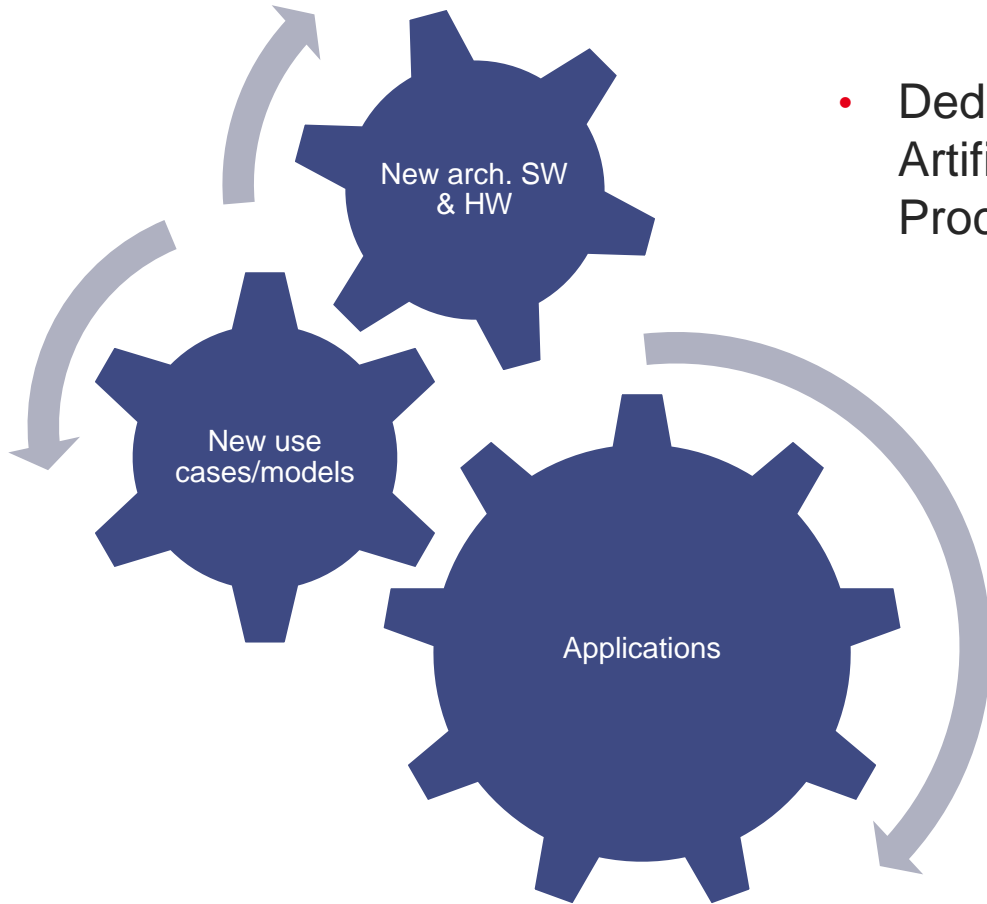
- Study and represent the needs of the users concerned by the exascale
- Propose and implement actions to help the application community to
  - Taking ownership of the exascale
  - Proposing areas for improvement
  - Formulate the needs and blocking points for the applications
- Support the research community with
  - Training support
  - Support in terms of resources (human, material, etc.)
- Identify the applications that will be present in the call for expressions of interest





# Lifetime of supercomputers vs the Science

- At the intersection of various projects and collaborations around the Exascale
- Dedicated to scientific applications (High Performance Computing, Artificial Intelligence, Data Analysis, High Performance Computing Processing, etc.)



# HPC, AI and HPDA community in France



- The French research community is very rich in scientific themes with an important place in the digital sector
- Applications are a key point to maintain France at a world rank in various strategic fields of scientific research: climate, energy, health, new materials...
- This involves considering the entire research chain in France, including:
  - Upstream research: applications in the "prototype" state, allowing applications to be maintained on innovative and futuristic architectures
  - Scientific production: maintaining the production of scientific results, essential for research and innovation
  - The notion of sovereignty: it is about being able to act and master the application chain, ensuring that our researchers can and are able to exploit scientific applications (and their ecosystem as a whole)
- Thus, maintaining scientific production within the framework of the exascale requires a "general" mobilization of the community

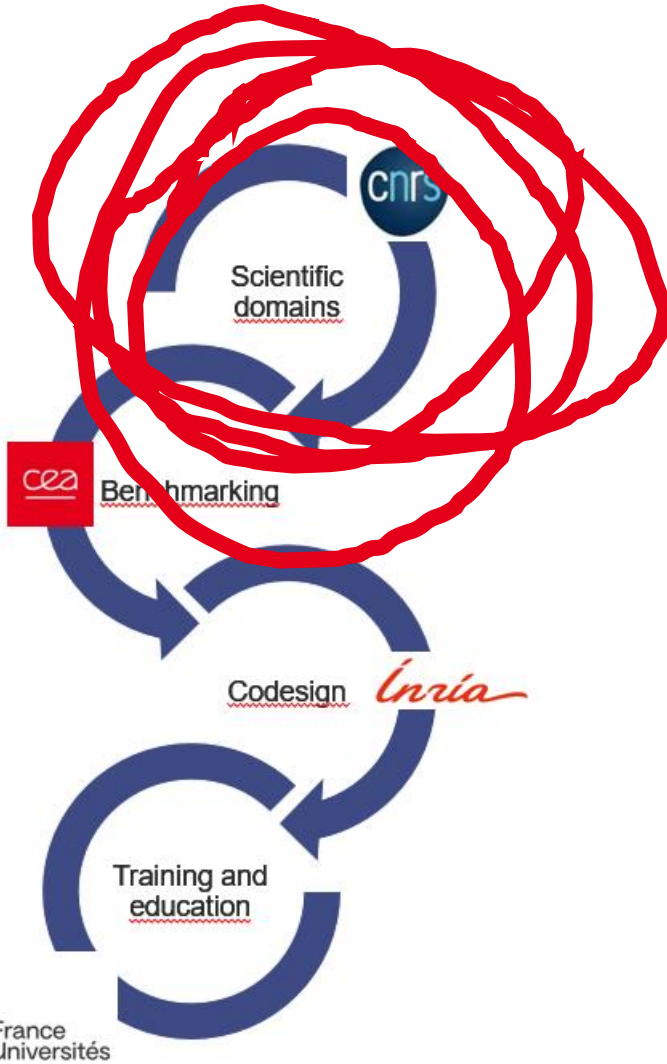
# Our work

- The SP3 organized the writing of a deliverable to identify applications from the French research community:

<https://hal.archives-ouvertes.fr/hal-03736805/document>

- The document is intended to provide a clear vision of the applications of the French research community

- Current state/status of the applications
  - Who's using/developing the application?
  - Is it open source?
  - Is it portable? Scalable?
- Prospects and projections for the exascale: The projection and ambitions of the teams with regard to exascale (and beyond)
- Cost induced by the evolution of codes (material, human, ...)
  - Requirements for switching to exascale, meaning:
    - accelerated architecture and/or improved scalability via algorithmic developments and/or integration of new digital models, human resources, etc.



# Exascale as a key application enabler for European scientific and societal challenges



**ÉNERGIE**  
Groupe de travail animé par  
Ansar CALLOO (CEA) et Denis VEYNANTE (CNRS)

**INDUSTRIE DU FUTUR & TRANSITION NUMÉRIQUE**  
Groupe de travail animé par  
François BODIN (Université de Rennes 1) et Bart LAMIROY (Université de Reims Champagne-Ardenne)

**RECHERCHE FONDAMENTALE**  
Groupe de travail animé par  
Thierry DEUTSCH (CEA) et David LOFFREDA (CNRS)


**SCIENCES DE L'UNIVERS**  
Groupe de travail animé par  
Marie-Alice FOUJOLS (CNRS) et Sylvie JOUSSAUME (CNRS)

**SCIENCES DU VIVANT**  
Groupe de travail animé par  
Marc BAADEN (CNRS) et Alexandra CARBONE (UPMC)

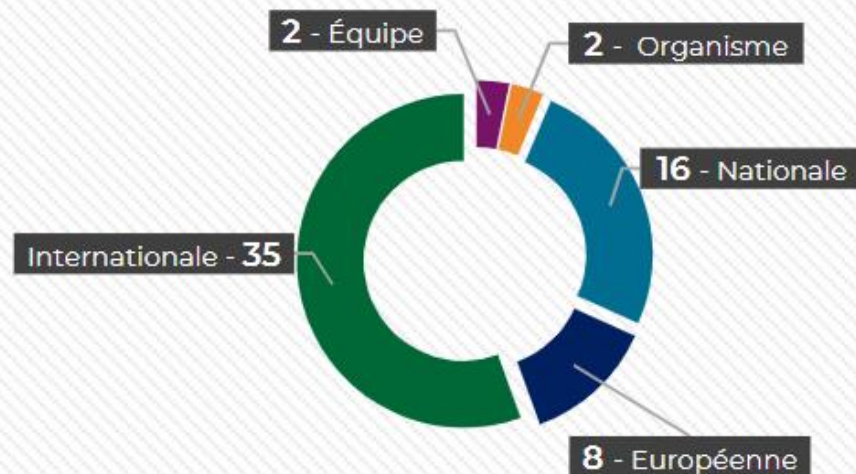
**SCIENCES ET USAGES DU NUMÉRIQUE**  
Groupe de travail animé par  
Raymond NAMYST (Université de Bordeaux) et Bruno RAFFIN (INRIA)

# Overview of codes


## NOMBRE DE CODES RECENSÉS POUR CHAQUE GROUPE DE TRAVAIL



## COMMUNAUTÉ DES CODES RECENSÉS DANS LES GROUPES DE TRAVAIL




## STATUT DES CODES RECENSÉS DANS LES GROUPES DE TRAVAIL

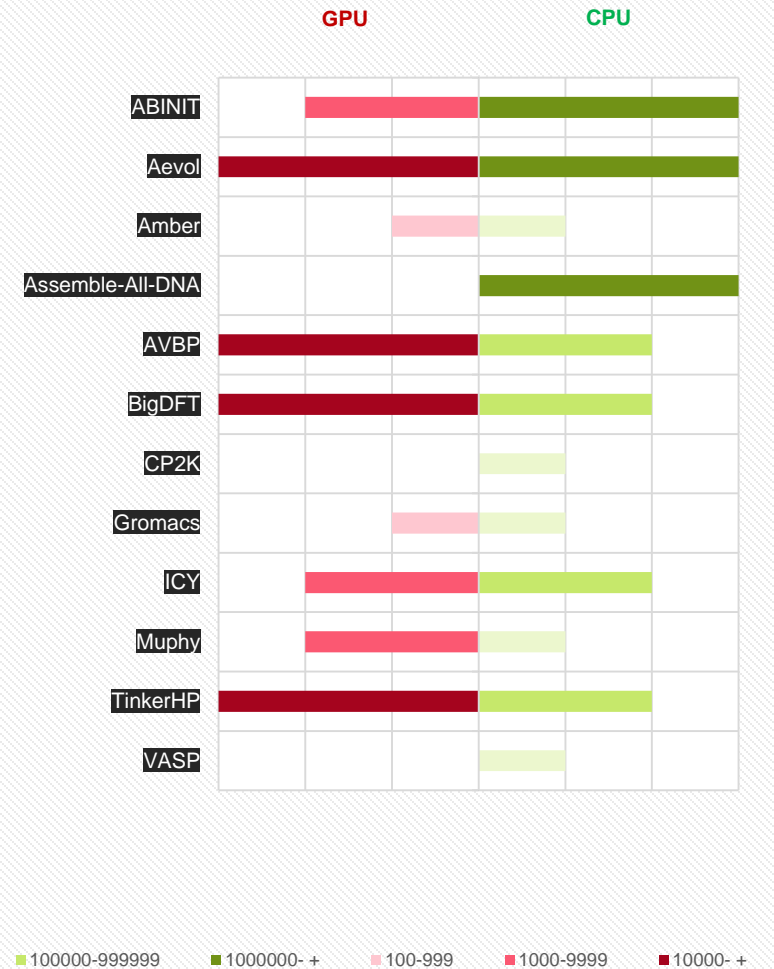


# Scalability

SCALABILITÉ À L'HORIZON 2025 DES CODES RECENSÉS PAR LE GROUPE DE TRAVAIL ENERGIE



SCALABILITÉ À HORIZON 2025 DES CODES RECENSÉS PAR LE GROUPE DE TRAVAIL SCIENCES DU VIVANT



10000-99999 100000-999999 1000000-+ 100-999 1000-9999 10000-+

# A united community



About 230 researchers involved

Le SP3 remercie très chaleureusement tous les participants aux groupes de travail qui ont mobilisé plus de 215 personnes et nous remercions chacune d'entre elle pour leur contribution, leurs efforts et leur soutien. Ce travail n'aurait pu aboutir sans la mobilisation de la communauté « Computational Science » et « Data Science ».

Le SP3 tient à remercier tous les animateurs de chaque groupe de travail : leur dévouement et leur motivation ont été un point clé du succès des groupes de travail.

Enfin, le SP3 remercie GENCI, qui nous a apporté un soutien logistique qui a permis le succès de ces travaux. L'aide et le soutien de GENCI tout au long de cette réalisation ont été déterminants.

Jean-Paul AMPUERO  
Gabriel ANTONIU  
Dominique AUBERT  
Ludovic AUBRY  
Édouard AUDIT  
Nicolas AUNAI

Marc BAADEN  
Mohd-Afeef BADRI  
Rémi BARON  
Corinne BEAL  
Olivier BEAUMONT  
Arnaud BECK  
Marina BECOULET  
Rachid BENSHILA  
Marjorie BERTOLUS  
Édouard BERTRAND  
Guillaume BESLON  
Gérard BIAU  
Julien BIGOT  
Quentin BLETERY  
Benoît BLOSSIER  
François BODIN  
France BOILLOD-CERNEUX  
Olivier BOUCHER  
Frédéric BOURNAUD  
Éric BOYER  
Olivier BRESSAND  
Patrick BROCKMANN  
Romain BROSSIER  
Allen-Sacha BRUN  
Hugo BUFFERAND  
Gaël BURGOS  
ALFREDO BUTTARI

Ansar CALLOO  
Yann CAPDEVILLE

Alessandra CARBONE  
Fabien CASSE  
Frédéric CAZALS  
Jean-Pierre CHABOUREAU  
Emmanuel CHALJUB  
Jérôme CHAPELLE  
Damien CHAPON  
Frédéric CHEVALLIER  
Rayan CHIKHI  
Philippe CIUCIU  
Jean-François CLOUET  
Guillaume COLIN DE VERDIERE  
Lucille COLOMBEL  
Laurent COLOMBET  
Benoît COMMERÇON  
Juan CORTES  
François COSTE  
Sabine CREPE-RENAUDIN  
Paul CRISTINI  
Laurent CROUZET

Michel DAYDE  
Alexandre DE BREVERN  
Florent DE MARTIN  
Laurent DEBREU  
Jean-François DELEUZE  
Maxime DELORME  
Thierry DELZESCAUX  
Julien DEROUILLAT  
Thierry DEUTSCH  
Yohan DUBOIS  
Thomas DUBOS  
Laurent DUFLOT  
Arnaud DUROCHER

Bruno ESPINOSA  
Juan ESCOBAR MUNOZ

Caroline ESSERT  
Daniel ESTEVE  
Catherine ETCHESBET  
Emmanuel FAURE  
Luca FEDELI  
Evelyne FOERSTER  
Cauthier FOLZAN  
Vincent FAUCHER  
Miguel FERNANDEZ  
Marie-Alice FOUJOLS  
Alexandre FOURNIER  
Pascal FREY  
Patrick FUCHS

Filippo GATTI  
Fabrice GAUDIER  
Thomas GASTINE  
Hugo GLOTIN  
Thierry GAUTIER  
Luigi GENOVESE  
Eric GOGLIN  
Virginie GRANDCOIRARD  
Mickael CRECH  
Sergei CRUDININ  
Abdou GUERMOUCHE  
Fabien CILLET-CHAULET  
Laurent GUILLOT

Hubert HANSEN  
Laurent HELIOT  
Yann HÉRAULT  
Guido HUISMANS  
Guillaume HUPIN

Olivier JAMOND  
Sébastien JAN

Emmanuel JEANNOT  
Pierre JOLIVET  
Laetitia JOURDAN  
Sylvie JOUSSAUME  
Laurène JOUVÉ

Masa KAGEYAMA  
Charles KERVRANN  
Bruno KLAMOLZ  
Samuel KOKH  
Michele KRAJECKI  
Gerhard KRINNER

Damien LAACE  
Elodie LAINE  
Bert LAMIROY

Stéphane LANTERI  
François LANUSSE  
Guillaume LATU  
Pierre-François LAVALLEE  
Philippe LAVOCAT  
Pierre LEDAC  
Laurent LEPEVRE  
Arnaud LEGRAND  
Geoffrey LESUR  
François LEVRIER  
Mathieu LOBET  
David LOFFREDA  
Celine LOSCOS

Pierre-Etienne MACCHI  
Fabienne MAIGNAN  
Jean-François MANGIN  
Marilice MIHAL-COSMIN  
Olivier MARTI  
Michel MASELLA  
Sébastien MASSON  
Benoît MATHIEU  
Yves MECHULAM  
Claudine MEDIGUE  
Luca MESSINA  
Yann MEURDESOLF  
Raphaël MITTEAU  
Vedim MONTEILLER  
Luca MONTICELLI  
Richard MORENO  
Raymond NAMYST  
Cyril NGUYEN  
Florence NIEDERGANG  
Michele NILGES  
Gregory NUEL  
Lucas NUSSBAUM

Jean-Christophe OLIVO  
Anne-Cécile ORGERIE

Jean-Noël PATILLON  
Perrine PAUL-GILLOTEAUX  
Christian PEREZ  
Jean-Philippe PIQUEMAL  
Cyril POUPON  
Simon PRUNET

Bruno RAFFIN  
Pierre RAMET  
Giuseppe RASTIELLO  
Elisabeth REMY  
Arnaud RENARD  
Zaccharie RAMZI

Yann RASERA  
Stéphane REQUENA  
Jean ROMAN  
Jonathan ROUZAUD-CORNABAS

Roland SABOT  
Marco SAITTA  
Jean SALAMERO  
Mathieu SALANNE  
Semuel SALVADOR  
Emilia SANCHEZ-GOMEZ  
Yanic SARAZIN  
Claude SCARPELLI  
Nathanaël SCHAEFFER  
Roos SMETS  
Vittorio SOMA  
Aymeric SPIGA  
Jean-Luc STARCK  
Fabio STERPONE  
Radek STOMPOR  
Antoine STRUGAREK

Patrick TAMAIN  
Éric TANNIER  
François TESSIER  
Maxence THEVENET  
Simon THORPE  
Marc TORRENT  
El-Chazali TALBI  
Helene TOUZET  
Pascal TREMBLIN

Sophie VALCKE  
Patrice VERDIER  
Thibaut VERY  
Denis VEYNANTE  
Jean-Pierre VILOTTE  
Henri VINCENTI  
Dorothea VOM BRUCH  
Mathieu VRAC  
Rodolphe VUILLEUMIER

Philippe WAUTELET  
Sevvas ZAFIROPOULOS  
Neil ZAIM  
Christophe ZIMMER

Ainsi que nos partenaires Industriels qui ont contribué à ce document, et bien d'autres encore que nous remercions chaleureusement.

# What I was supposed to present



- All the collaborations (within EU and worldwide) around exascale. We can discuss it at the coffee break
  - Strong work within all EU: EuroHPC and bilateral collaboration with FZJ (*but if you want to join we are happy to work with you*)
  - Strong collaboration with Japan (HPC, AI and HPC-QC coupling): CEA/RIKEN collaboration (*but if you want to join we are happy to work with you*)
  - Strong collaboration with USAs (HPC and AI)

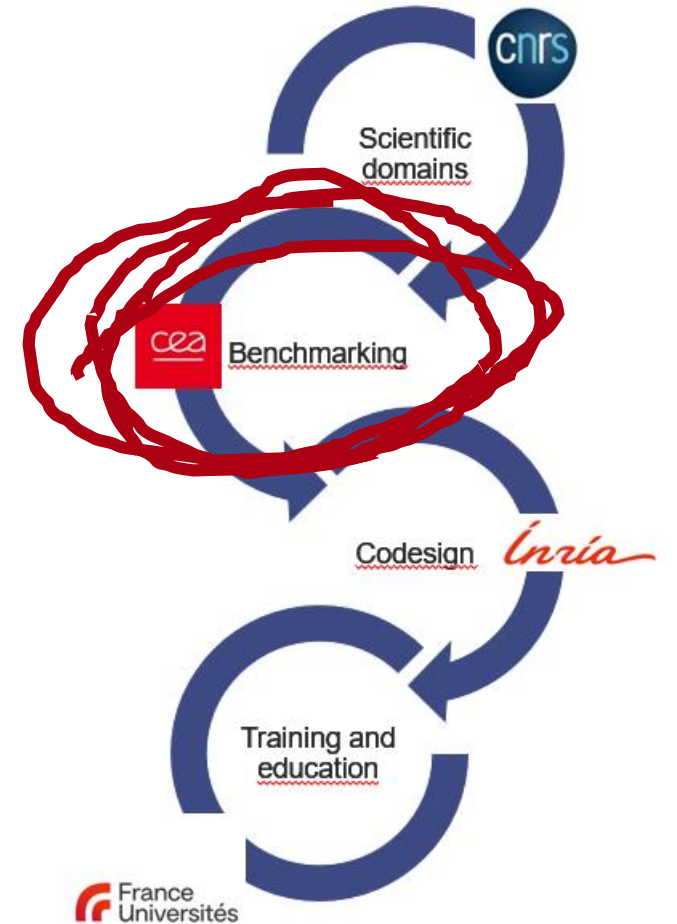
## What I will talk about

- How the hell do we do benchmark and codesign for an exascale machine? What does even this mean?



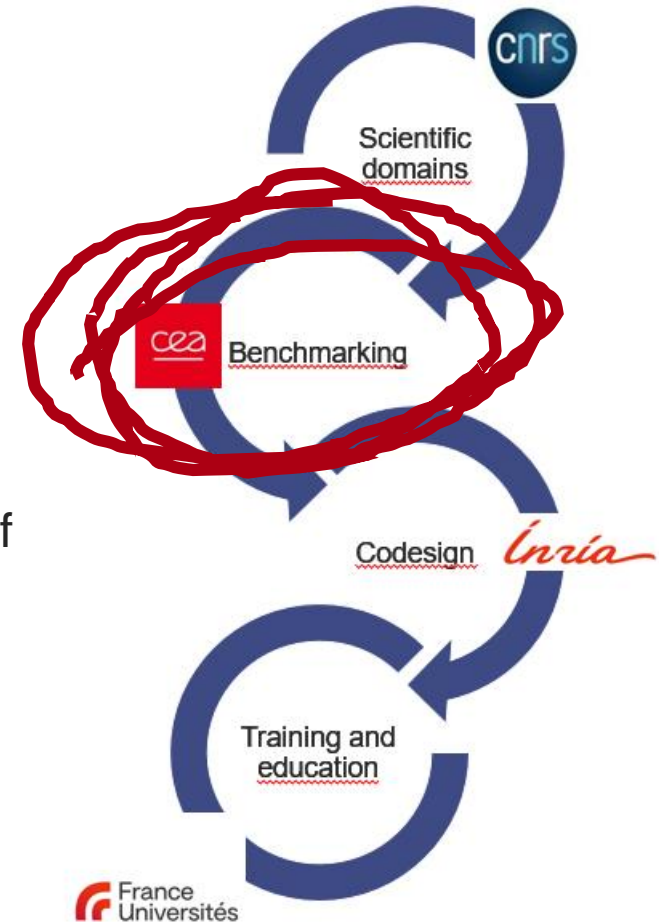
# What's next?

- Reminder: the exascale project is funding the supercomputer (TCO), NOT the users.
- In this context, a benchmark is « just » a use case within an RFP
  - Which will be confidential. Therefore, the outcome of a benchmark in this context is 0 for the scientific team.
- Plus, the supercomputers must achieve an exaflop: meaning, the HPL must go beyond 1 Eflops.
  - In consequence, the codesign is GPU constrained
- Then, why shall we consider the benchmark and codesign in this project?
  - Opportunity to execute mixed/complex workflows is open
  - Data issue is considered



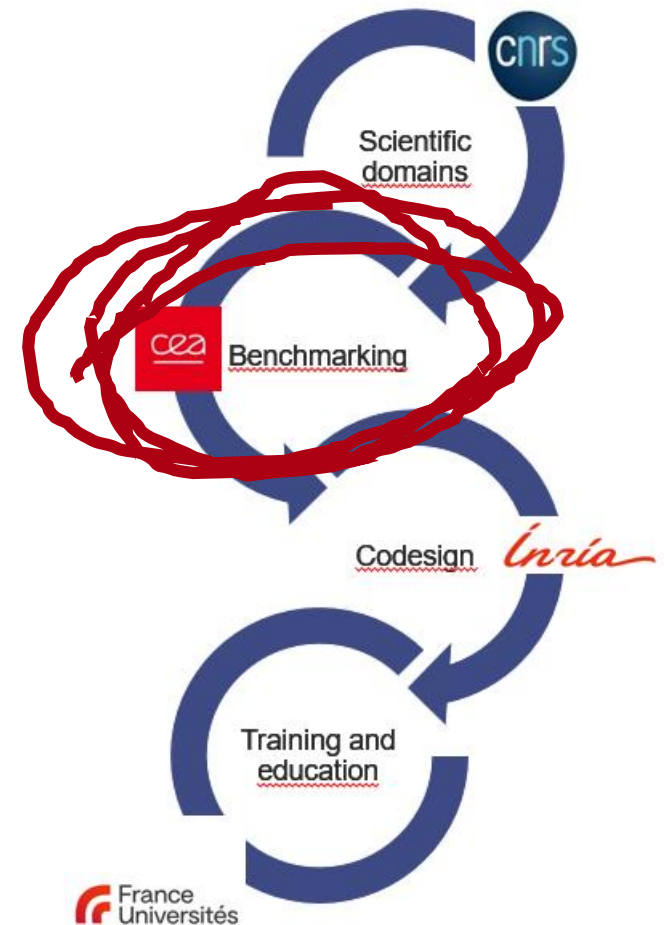

# What's a benchmark?

- User: a benchmark is a use case I execute as a daily routine on the supercomputer.
- Vendor: a benchmark is an execution from which I will reduce as much as I can the TTS and make sure this is compliant with my products (among which, the compiler I sell)
- Sys-admin: a benchmark is a mini-app extracted from a representative application that is stressing a specific characteristic of the supercomputer.



# What's a benchmark?

- 3 different priorities therefore 3 different definitions.
- 1 common point though: the science



# One use case to rule them all

- How do you go from this



work on app1

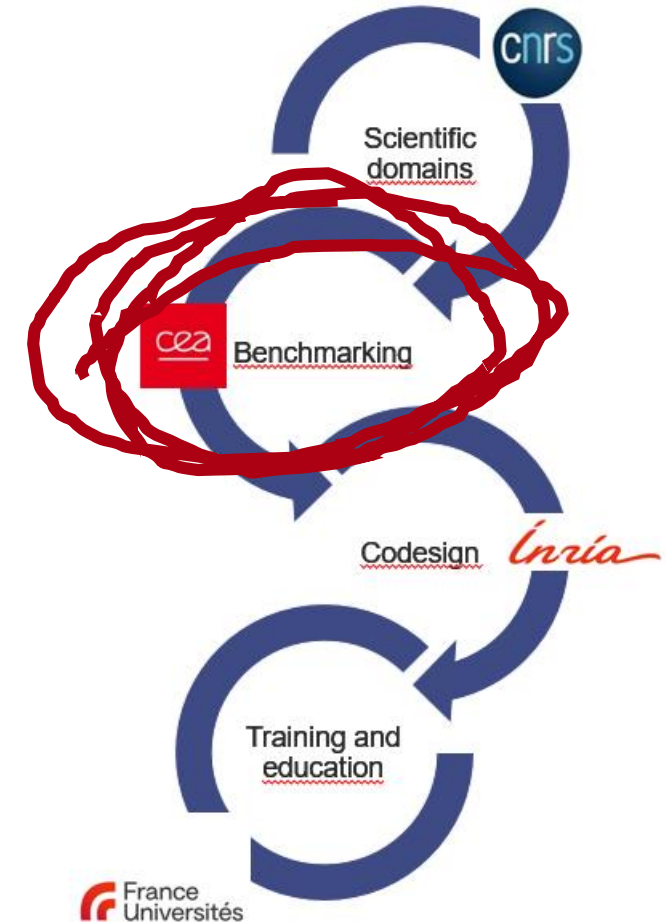


work on app1'




work on app1''

- To this



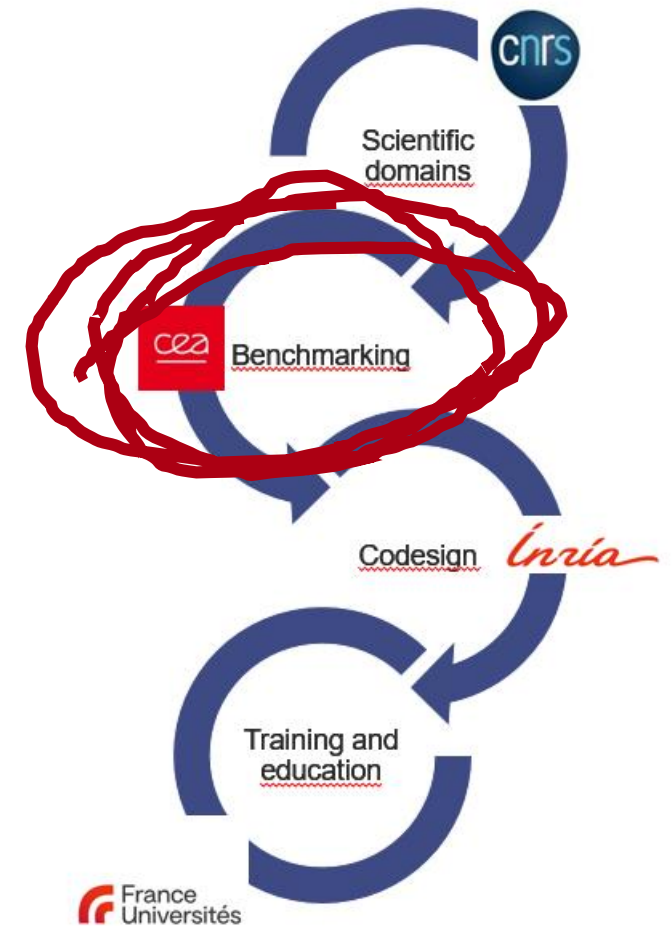


# Thinking...

- Start from the start
  - What's the community? The three groups we mentioned. And they MUST work together.
  - What's a benchmark? Gather the community to build the definition
  - When you get the definition, you can collect the needs.
  - When you get the needs, you can start thinking of a solution.
- Make sure the community is maintaining the 4 steps above during all the process





# Still thinking...

- What's a benchmark? Gather the community to build the definition
  - The community shall provide a list of metrics to characterize a benchmark
    - Ex: portability, possible characteristics of the bench (io bound?), other...
    - This can be viewed as « metadata » of a benchmark
  - Depending on « who's using » the benchmark (and therefore what the computing cat has to do with is)
    - The priority on the listed metrics are different
      - Sys admin cat highest priority is « possible characteristics of the bench »
      - Vendor cat highest priority is « TTS » and « runability » of the code
      - Scientific cat highest priority is « scientific truthworthy »




# Where does it go wrong?


- A benchmark is **WRONG** when you miss one of the metric...
- ... The priority on each of the metric can be different ...
- ... **But even for the lowest priority metric, you must make SURE this one is still true**



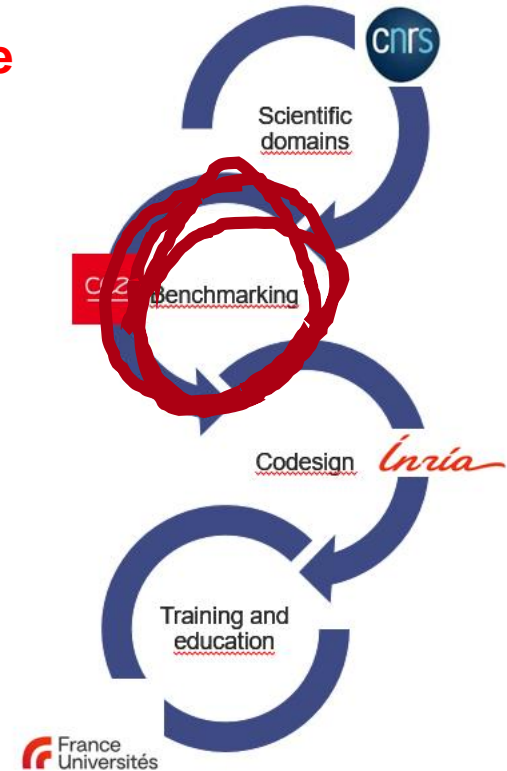
The sys-admin cat and vendor cat methodologies are wrong when they forgot to consider the « scientific truthworthy » of the benchmark.




The scientific cat methodology is wrong when it forgot to consider the portability metric



The sys-admin and scientific cats methodology is wrong when it forgot to consider the metrics relevant to the vendor cat



# Where does it go wrong?




The sys-admin cat and vendor cat methodologies are wrong when they forgot to consider the « scientific truthworthy » of the benchmark.

Expertise developped on irrelevant use case



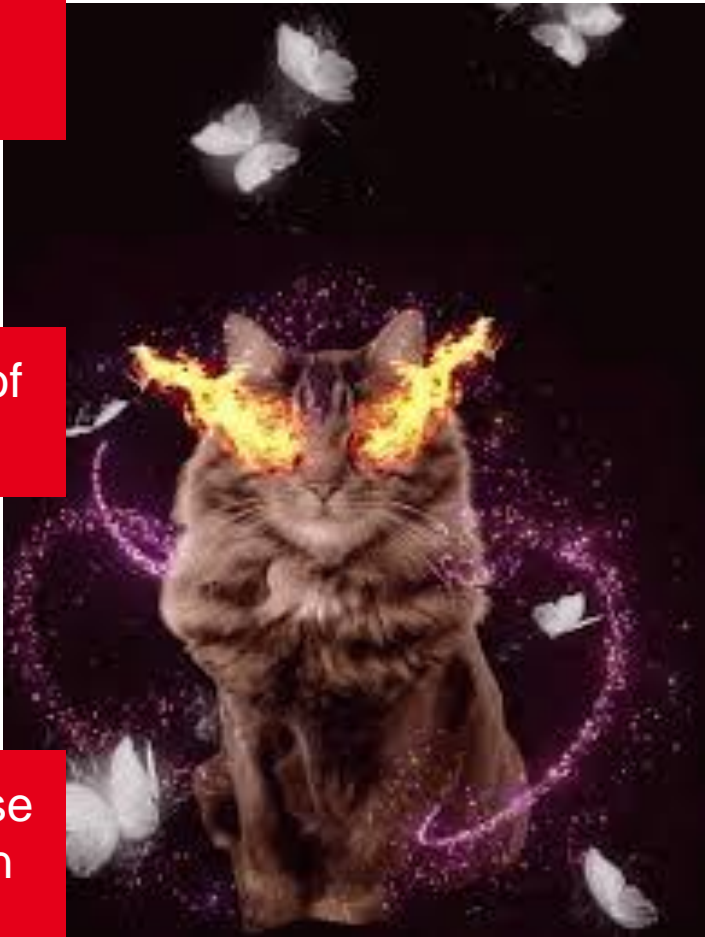
The scientific cat methodology is wrong when it forgot to consider the portability metric

Possible decrease (or stop) of the scientific production



The sys-admin and scientific cats methodology is wrong when it forgot to consider the metrics relevants to the vendor cat

Expertise developped on a use case that will be irrelevant on future HW & SW



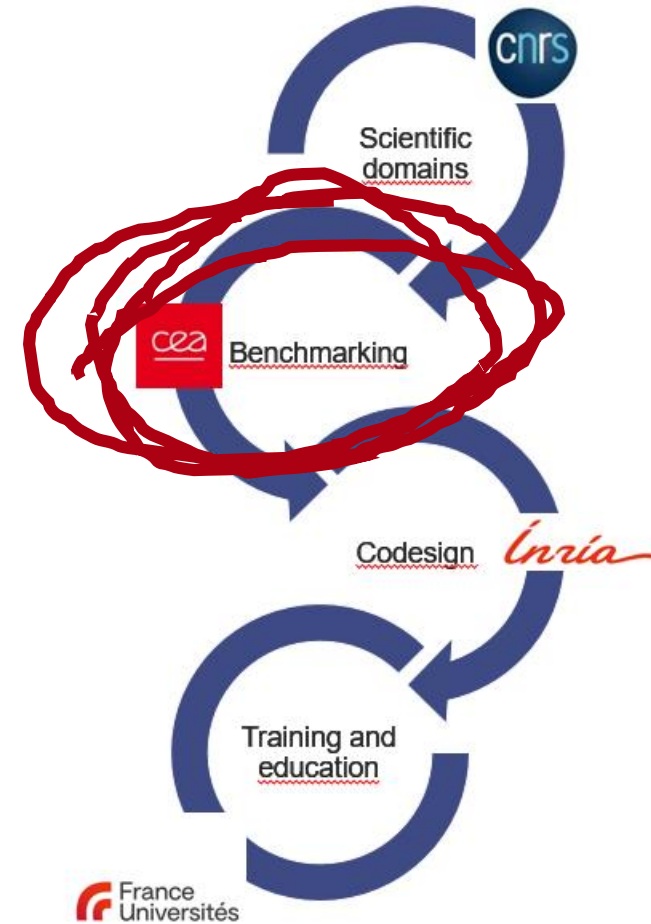


# Can we fix it?



- Bring the community around the table
- Let's discuss the metrics
  - Based on the ongoing work developed within DOE-MEXT collaboration
  - Gather the international community around this very important work

• **Please join in Paris (or online) in September 2023: 25-27 very likely.**



# One community

The scientific cat needs to maintain its scientific production and anticipate its future scientific production



Regression tests

The vendor cat wants to make sure the Supercomputers that will be on the market in 5 to 10y from now will be able to execute scientific applications



Design new models

Optimize the application

*more innovation*

Optimize the SC usage

Characterize the execution

The sys-admin cat wants to make sure the Supercomputer is at its best level to execute scientific production



*Green Computing*



# Thank you for your attention French community is waiting for you !

*i am here*

