Some (Post-)Exascale Challenges

An NumPEx/InpEx international perspective

October 2024, ORAP

Jean-Yves Berthou (INRIA)

NumPEx contributors: M. Krajecki, J-P. Vilotte, A. Buttari (**CNRS**), C. Prudhomme (**U. de Strasbourg**), H. Barucq, G. Antoniu, B. Raffin (**Inria**), R. Namyst (**U. de Bordeaux**), J. Bigot, T. Deutsch, V. Brenner (**CEA**), F. Bodin (**U. de Rennes**), M. Asch (**U. Picardie**)

InPEx-EU contributors: Sergi Girona, Rosa M Badia, Fabrizio Gagliardi (**BSC),** Sanzio Bassini (**CINECA**), Stéphane Requena (**Genci**), Kimmo Koski, Per Oster, Damien Lecarpentier, Janne Ignatius (**CSC**), Bernd Mohr (**JSC**), Claus-Axel Mueller (**GCS**), Martin Schulz (**LRZ/TUM),** Michael Resch **(HLRS)**

The French NumPEx Program

Consolidating and accelerating the construction of a sovereign European **exascale software stack** and **strategic applications (post)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**

Foster national and international collaborations to prepare for the Exascale and post-Exascale era

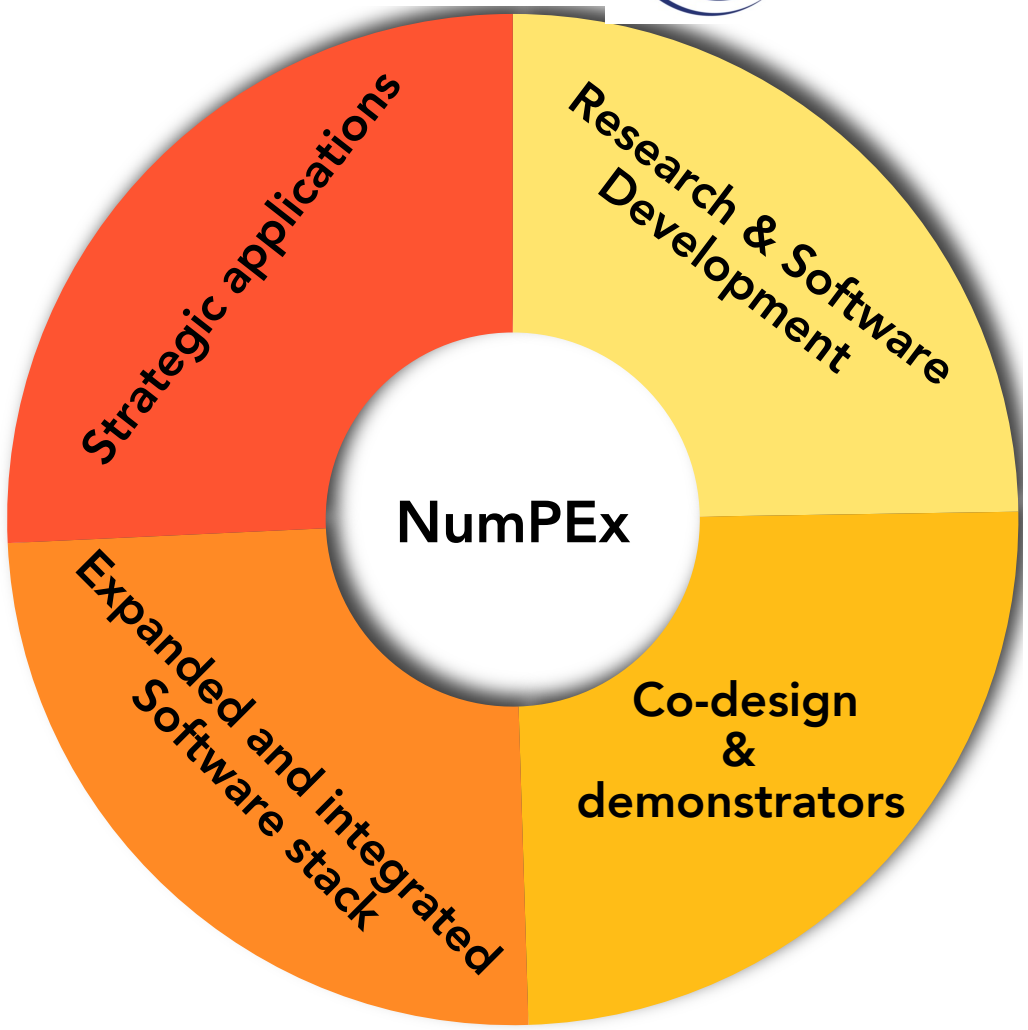

Help aggregate the French Edge/Cloud/HPC/HPDA/IA community

NumPEx by numbers

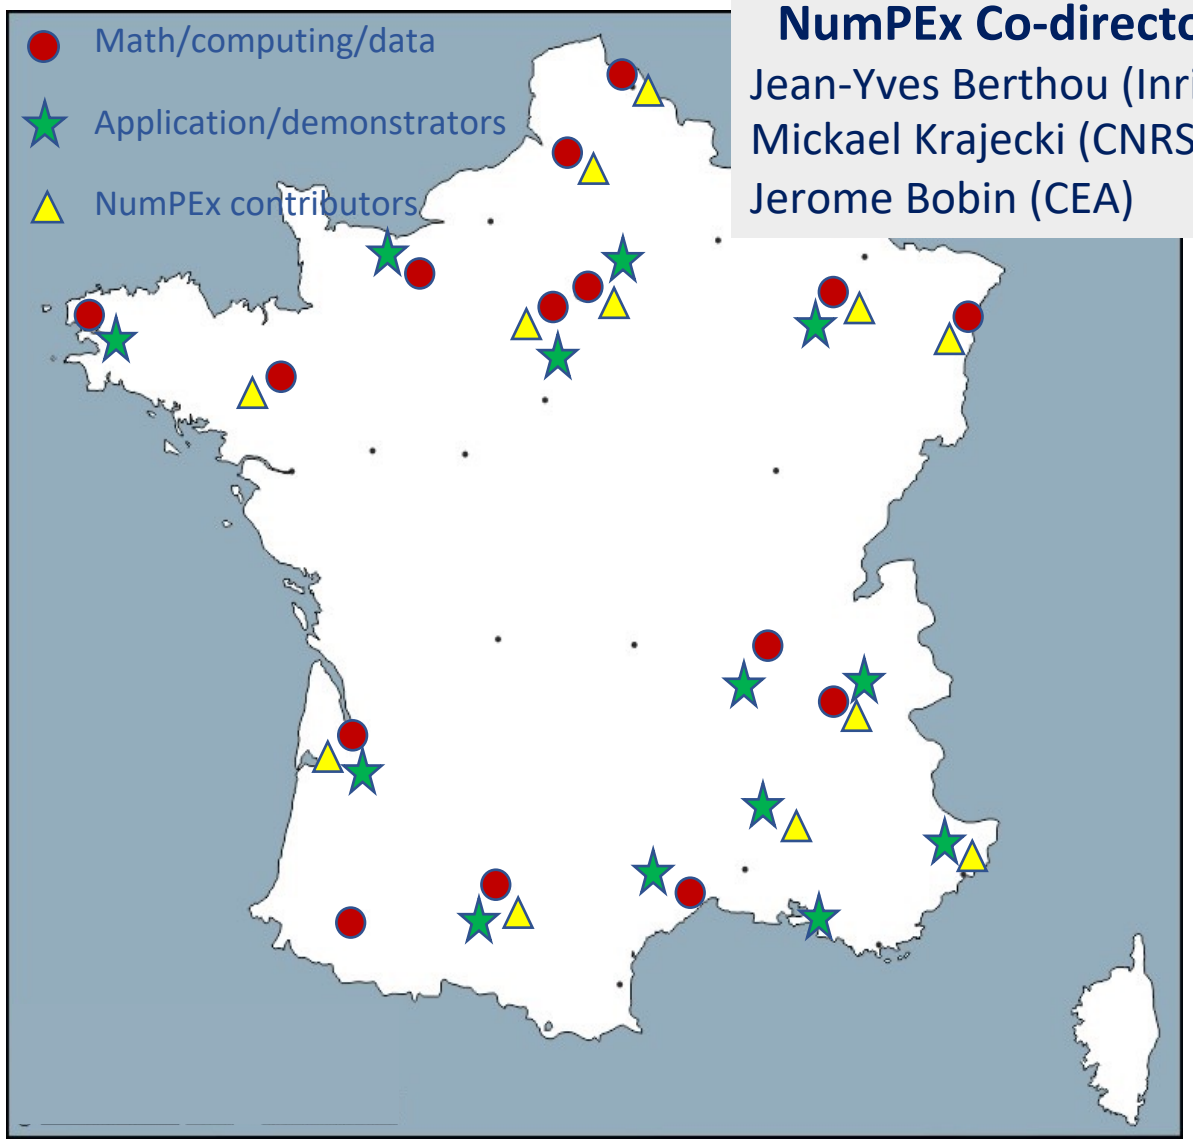

NumPEx Co-directors
Jean-Yves Berthou (Inria)
Mickael Krajecki (CNRS)
Jerome Bobin (CEA)

6 Years
81 M€*

2023-2028
* Funding 41M€=500 person.year non permanent staff
+ 170 person.year permanent staff
Total cost : 81 M€

Core Research Institutions

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

3 Focus Area

Software stack development (PC 1-3)
Wide-area workflows and architecture (PC 4)
Integration and application development (PC 5)

80 R&D teams

The French NumPEx Program - workplan

Applications

ExaDI
Application co-design and software integration
JP.Vilotte/V.Brenner

ExaAtow
Continuum digital

ExaDost
Data
G.Antoniu/J.Bigot

ExaMA
Numerical methods and solvers
C.Prudhomme/H.Barucq

ExaSoft
Computing
R.Namyst/A.Buttari

The French NumPEx Program – transversal WG

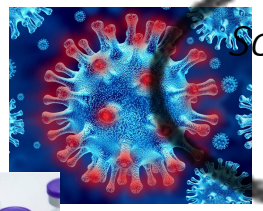

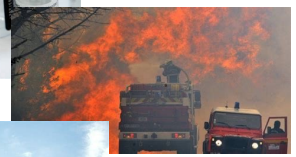




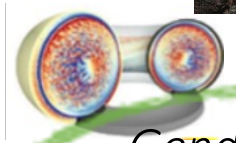

Software production & Intégration
B. Raffin

Resilience
L. Morin

Science4AI – AI4Science
T. Moreau

Accelerated archi. prog.
S. Thibault

Energy mgt&opt
G. Da Costa, A. Guermouche

Gender/Equity/ Diversity
Anne-Laure Pelé & Virginie Grandgirard

Training
M. Krajecki, R. Namyst, C. Prudhomme

GENCI French Tier0 facilities
F.Bodin/N.Lardjanne

Int. Collaborations
JY Berthou, E. Jeannot

The International Post-Exascale (InPEx) Project

InPEx expected outcomes

- Identify future trends/disruptions, missing software components
- Contribute to the share/development of software components: @deployable, @maintenable, @robust, @sustainable => **partnership factory**
- Landmark documents largely exploited, worldwide, for supporting future post-exascale science
- Develop an international network of exascale computing experts and leaders

Actions:

- Dedicated international **working groups**
- International Post-Exascale (InPEx) workshop series

Participants:

Researchers, engineers, industry, funding bodies


Pre-workshop InPEx, October 2023, Reims, F

The International Post-Exascale (InPEx) Project

Inpex.science

Date	(10/2023)	11/2023	06/2024	06/2025
Location	Preparatory phase EU (France)	SC'23 - BOF	Workshop1 EU/BSC	Workshop2 Japan
Date	03/2026	09/2026	06/2027	09/2027
Location	Workshop3 US	Workshop4 EU	Workshop5 Japan	Workshop6 US

The (post)-Exascale race, where are we?

ECP EXASCALE COMPUTING PROJECT

China

Development of applications in preparation for the arrival of the Tianhe3 machine.

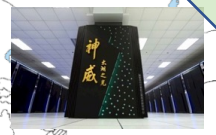

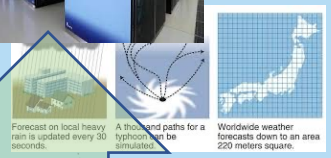

USA

ECP : DOE funded, NSF support (2016-2023)
+ Creation of 6 co-design centers
Still a challenge: Exascale ready app, sustainable software stack
New perspectives: Gen IA for Science, 20B$

Japan

- FugakuNext prod. 2029 (xB$): co-design HW/SW/Apps
- Next Gen AI (x100M$), Quantum-HPC (140M$)
- (post)-Exascale as a Service (AWS/Riken): from Fugaku to Virtual Fugaku

A strong effort in both hardware, software and applications/co-design

The (post)-Exascale race, where are we?

Focus on some post-exascale challenges

AI4HPC – HPC4AI

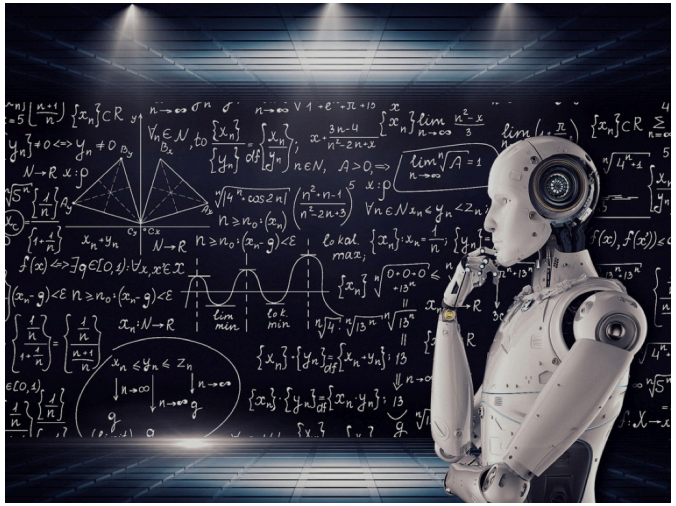

From edge to HPC systems
The digital continuum

Software, the new frontier

Software/application co-design

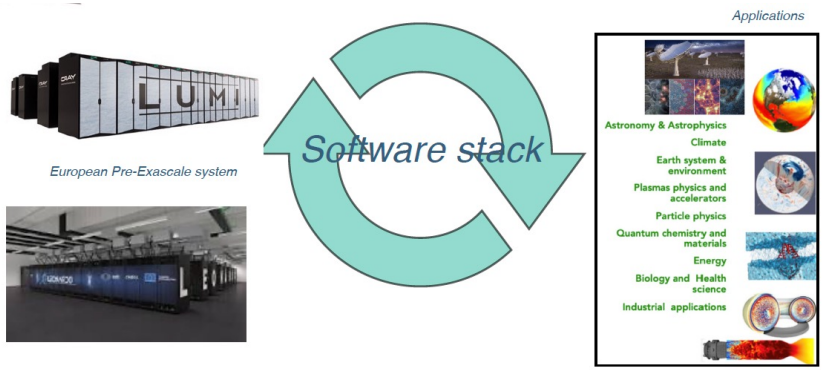

Focus on some post-exascale challenges

AI4HPC – HPC4AI

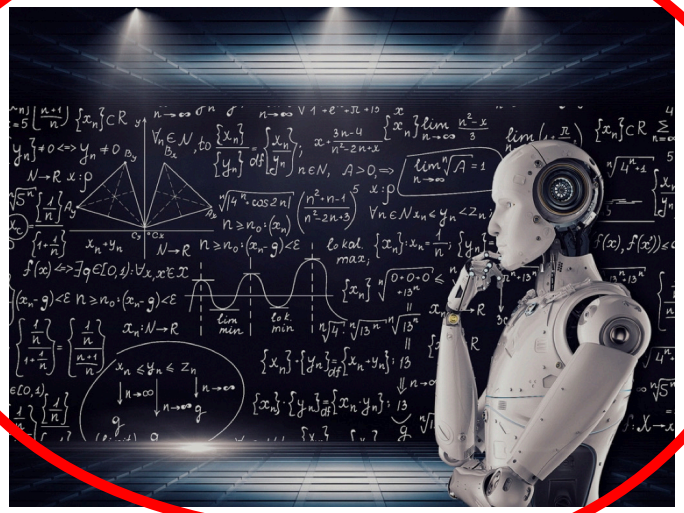

From edge to HPC systems
The digital continuum

Software, the new frontier

Software/application co-design

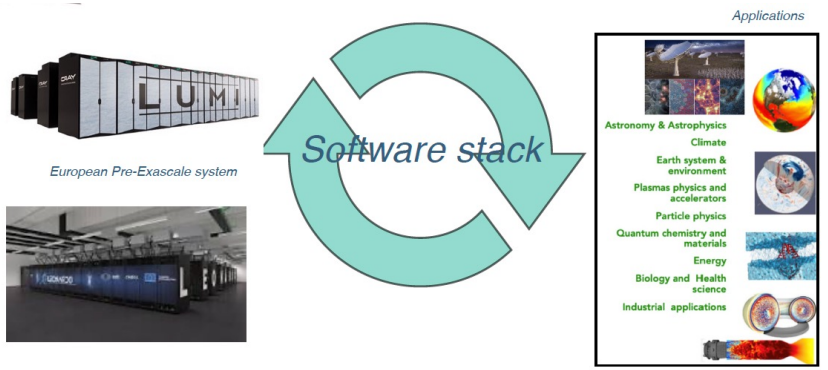

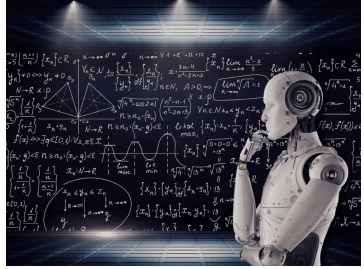

Some post-exascale challenges
AI4HPC

AI everywhere, a game-changer at the post-Exascale era

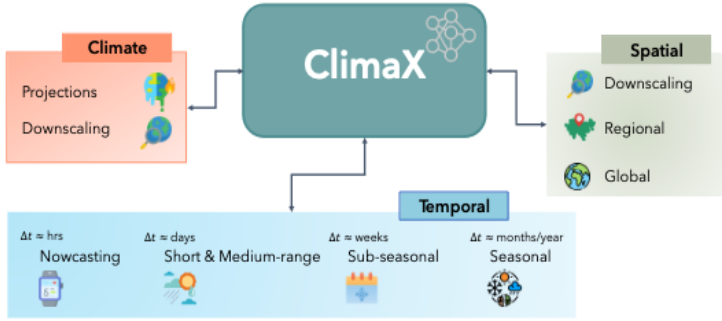

CLIMAX: a foundation model for climate science (Microsoft)

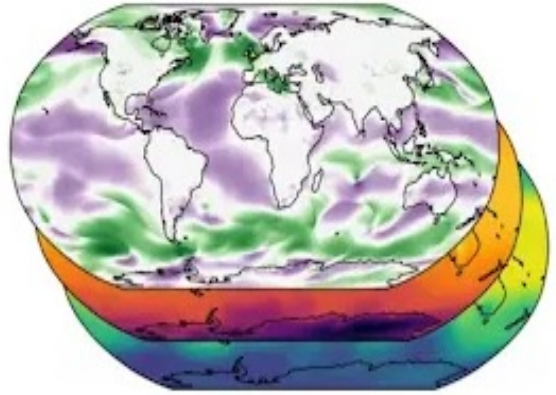

Graphcat for weather forecasting (Deepmind

International collaboration is key and Gen AI/LLMs are of major importance

But …

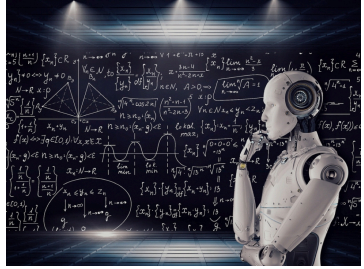

Some post-exascale challenges
AI4HPC

Gen AI/LLMs are of major importance, But …

Additional challenges to be addressed:

- Also domain specific foundation models, xLM …
- What are the right AI methods/algo for a given problem so that resources consumption are minimized?
- Towards an hybrid HPC/AI software stack (AI4HPC and HPC4AI)

  => 10 Exascale performance with x25 EDP+Data science pp and x42 HW improvements, some AI-based solvers can be speed up by 6 orders of magnitude, etc., weather forecast with Graphcast
- How to go towards validated/robust and trustworthy AI models/services?

Other impacts of AI for scientific computing:

- AI for code production, optimisation, portability, efficiency, maintenability
- AI for smart analysis of generated data
- AI for optimizing the management of workflow execution, dynamic ressources allocation, …
- AI for predicting and mitigating potential failures in exascale systems, ensuring robust and reliable operation of scientific applications, …

Some post-exascale challenges
HPC4AI

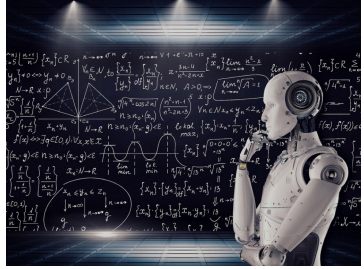

How AI can benefit of decades of HPC research&dev to address new AI challenges?

Efficient and frugal processing of large models:

- Data placement,
- Load-balancing,
- Resilience and fault tolerance,
- Smart execution models,
- Advanced, high level programming models,
- Portable models/architecture abstraction models,

Some post-exascale challenges
AI4HPC / HPC4AI

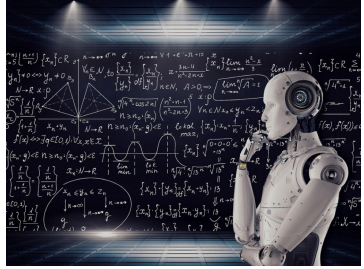

Some regional strategic and sovereignty issues

- How to manage data stewardship/ownership and sharing ?
- Trends towards AI-centric HW (e.g. AI accelerators) puts huge constraints on the HPC SW stack
- Need to build regional eco-system for HW&SW product from **academia** and **industries**

=> need for proper regional agenda and international parternship

Focus on some post-exascale challenges

AI4HPC – HPC4AI

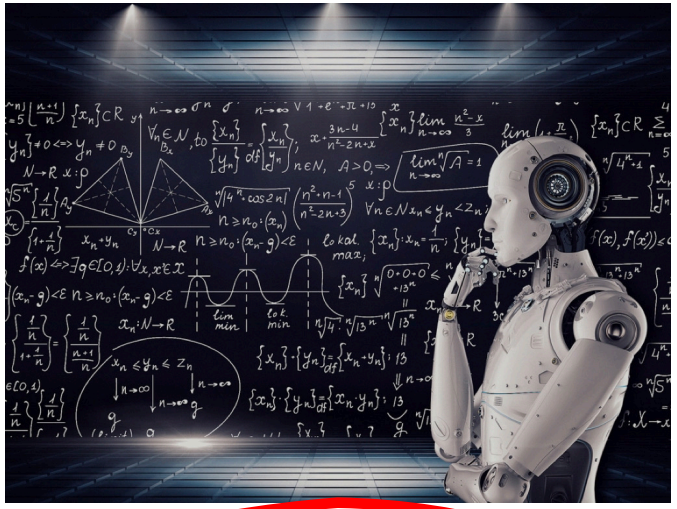

From edge to HPC systems
The digital continuum

Software, the new frontier

Software/application co-design

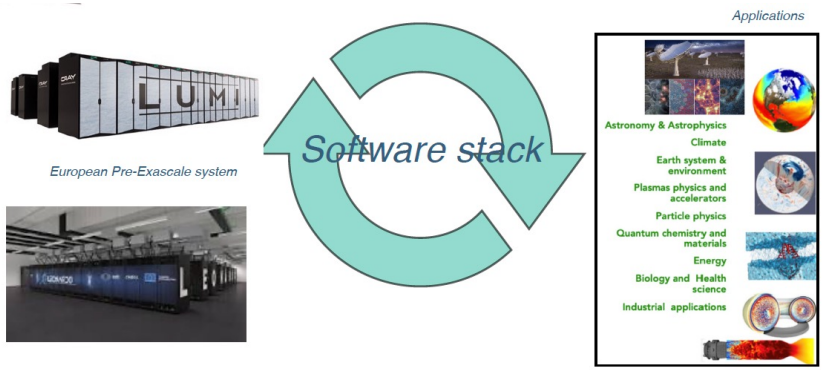

Software, the new frontier

PROJECTED PERFORMANCE DEVELOPMENT

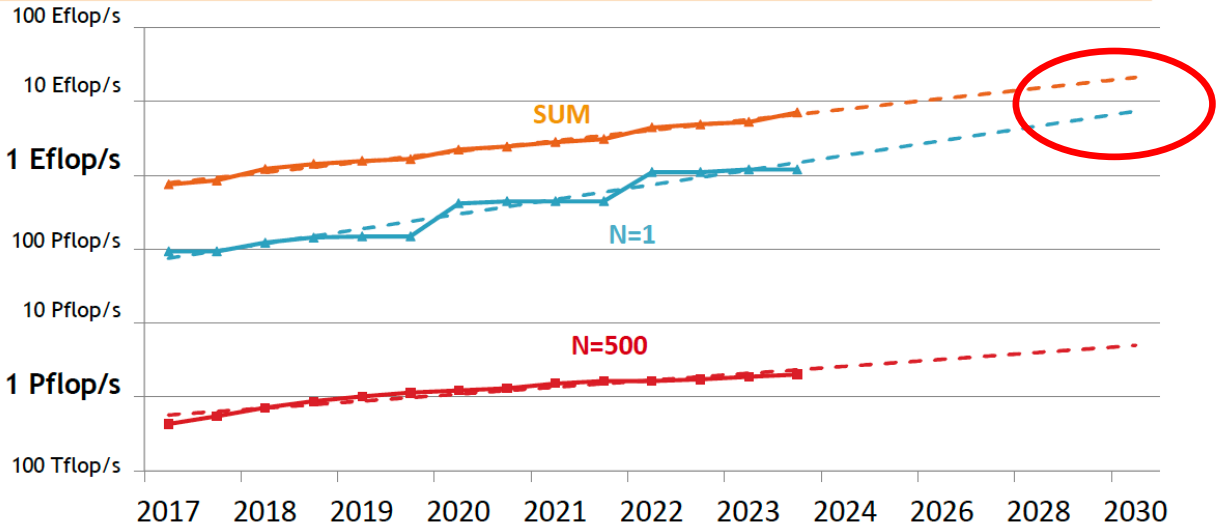

SC23 TOP500 HIGHLIGHTS

- Frontier is #1 and remains only ExaScale system in the TOP500 - for now
  - 4 new systems in TOP10! + one upgrade
  - And 9 additional new systems in TOP50
- First Intel Sapphire CPUs (25) and GPUs (4) on the list
- HPC systems are used longer and replaced less often
  - Due to an increasing number of technological limits
  - Leads to strong concentration at the top (research as well as commercial)
- TOP500 shows new reduced growth-rates since 2017!
  - End of original Moore's Law scaling
  - Unlikely to achieve 10 Exascale by the end of the decade unless we fundamentally change "business"!

AVERAGE SYSTEM AGE

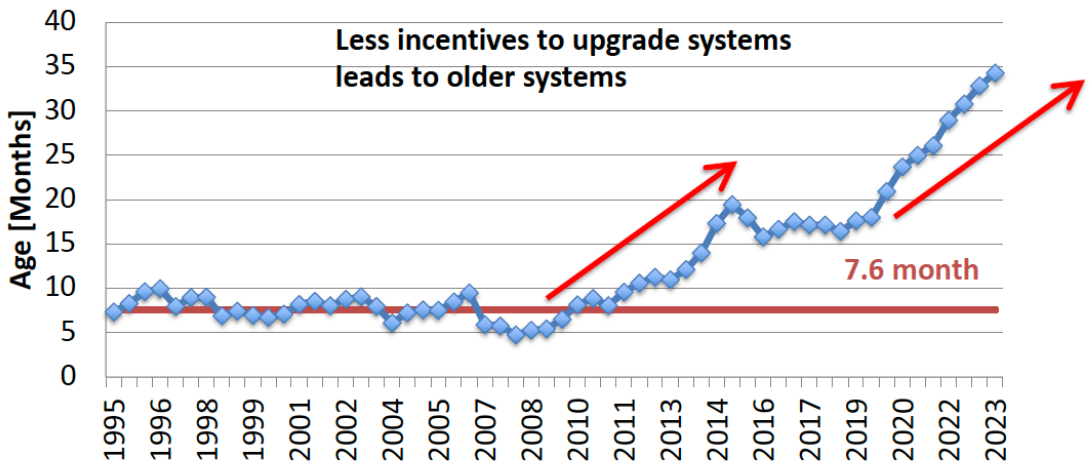

Software, the new frontier

Challenge - the race for power is slowing down: we're aiming for 10 Exaflops in 2030, not 1000

At Stake, gain effective computational power:

- **Maths&algorithms**: new digital schemes, precision mix, innovative discretization methods, energy aware algorithms,…
- **Programming and execution environment @ (post-)Exascale:**
- Converge towards common high-productivity programming interfaces:
  - hide machine complexity to improve productivity, portability and composability
  - transparently asynchronous parallelism with dynamic optimizations to improve scalability
  - AI+Software heritage (https://www.softwareheritage.org/) : AI-assisted code generation for robustness, efficiency, …
- Extend the mini-app ecosystem to guide the development of scalable and efficient software

Software, the new frontier

Challenge - software engineering

HPC applications and machines are **gaining in complexity** (think compute continuum, HPC/HPDA/AI/Cloud/Edge hybridization), leading to high costs to build and deploy applications on supercomputers, impairing portability and reproducibility

=> The standard software installation process on supercomputers is reaching its limits

At Stake: need for **regional exascale software stacks and fostering international partnership**

- Need for a HPC dev-ops methodology for fast deployment while ensuring performance and supporting reproducibility efforts and modern software packaging
- Offer an "industrial" production environment converging towards "As a Service »
- Capitalize on and consolidate software production in Europe and worldwide
- Importance of software sustainability

Focus on some post-exascale challenges

AI4HPC – HPC4AI

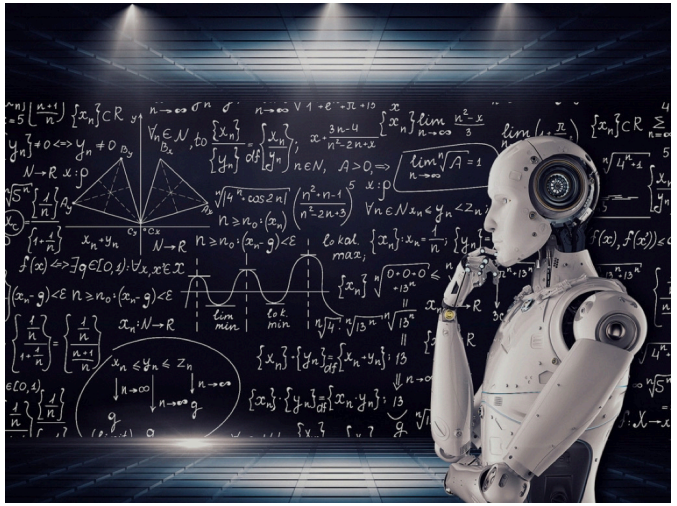

From edge to HPC systems
The digital continuum

Software, the new frontier

Software/application co-design

Software/application co-design

European Pre-Exascale system

Key challenges

- How to get post-Exascale ready applications ?
- How to expand an application-driven SW stack ?
- How to make applications portable and sustainable at the post-Exascale era ?

International context

- Early-binding HW/SW/application co-design approach at Riken (Japan)
- In the USA, DOE co-design centers were a key component of the Exascale Computing project (ECP)
- Centre of Excellence (CoEs) in Europe
- Inspired by ECP, co-design is central in the NumPEx project (FR)

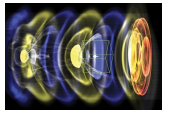

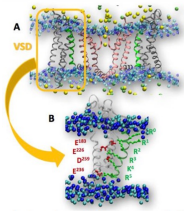

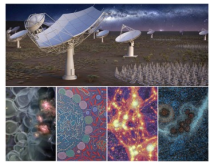

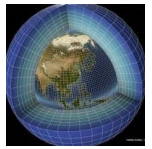

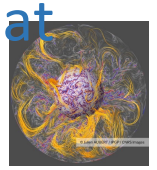

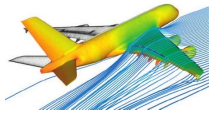

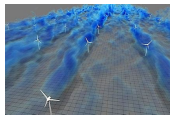

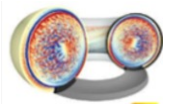

Astronomy & Astrophysics

Earth System Models & Climate

Earth sciences & environment

Computational biology & Life science

Laboratory laser-plasma physics

High-energy particle physics

Quantum chemistry and materials

Digital health

Environmental & societal risks

Urban systems planing

Magnetically confined fusion plasma

Sustainable Transport & mobility

Energy production & transport

Software/application co-design

Help the applications get prepared for post-Exascale challenges

- Towards **shared and common mini-apps/proxy-apps**, centered about technical functionalities/bottlenecks (AMR, I/O, AI, etc.)
- **Shared benchmarks** are central to test/evaluate portability/performance/deployability on different HW.

Towards a post-Exascale application-driven software stack

- Application code development is a long and costly journey, greater **reuse** of **composable/interoperable SW components**
- **Long-term visibility/sustainable development for applications is key.**

Focus on some post-exascale challenges

AI4HPC – HPC4AI

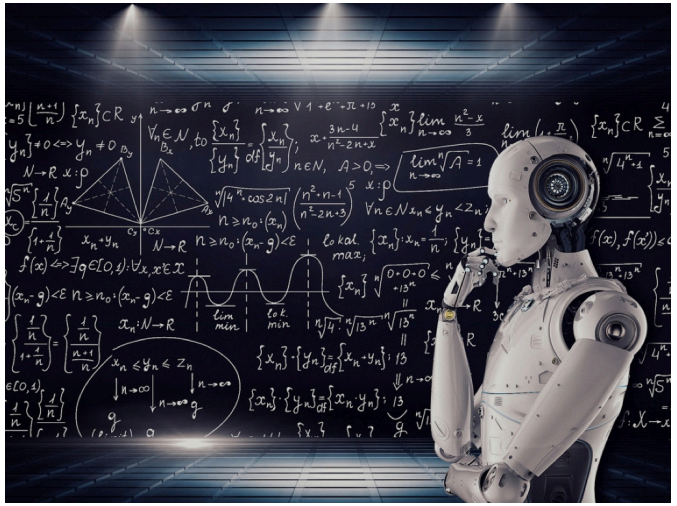

From edge to HPC systems
The digital continuum

Software, the new frontier

Software/application co-design

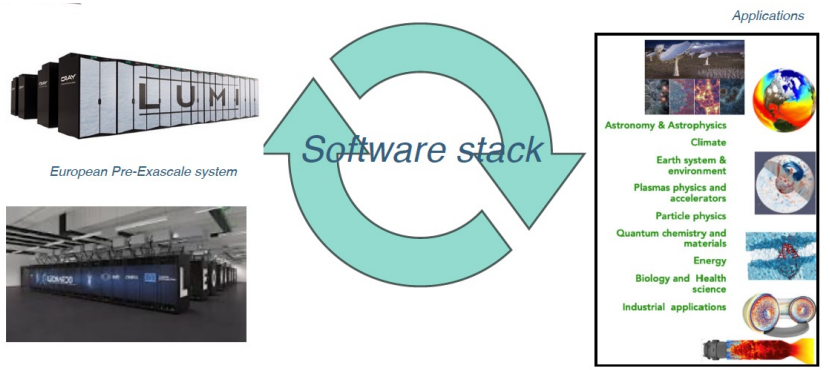

The digital continuum

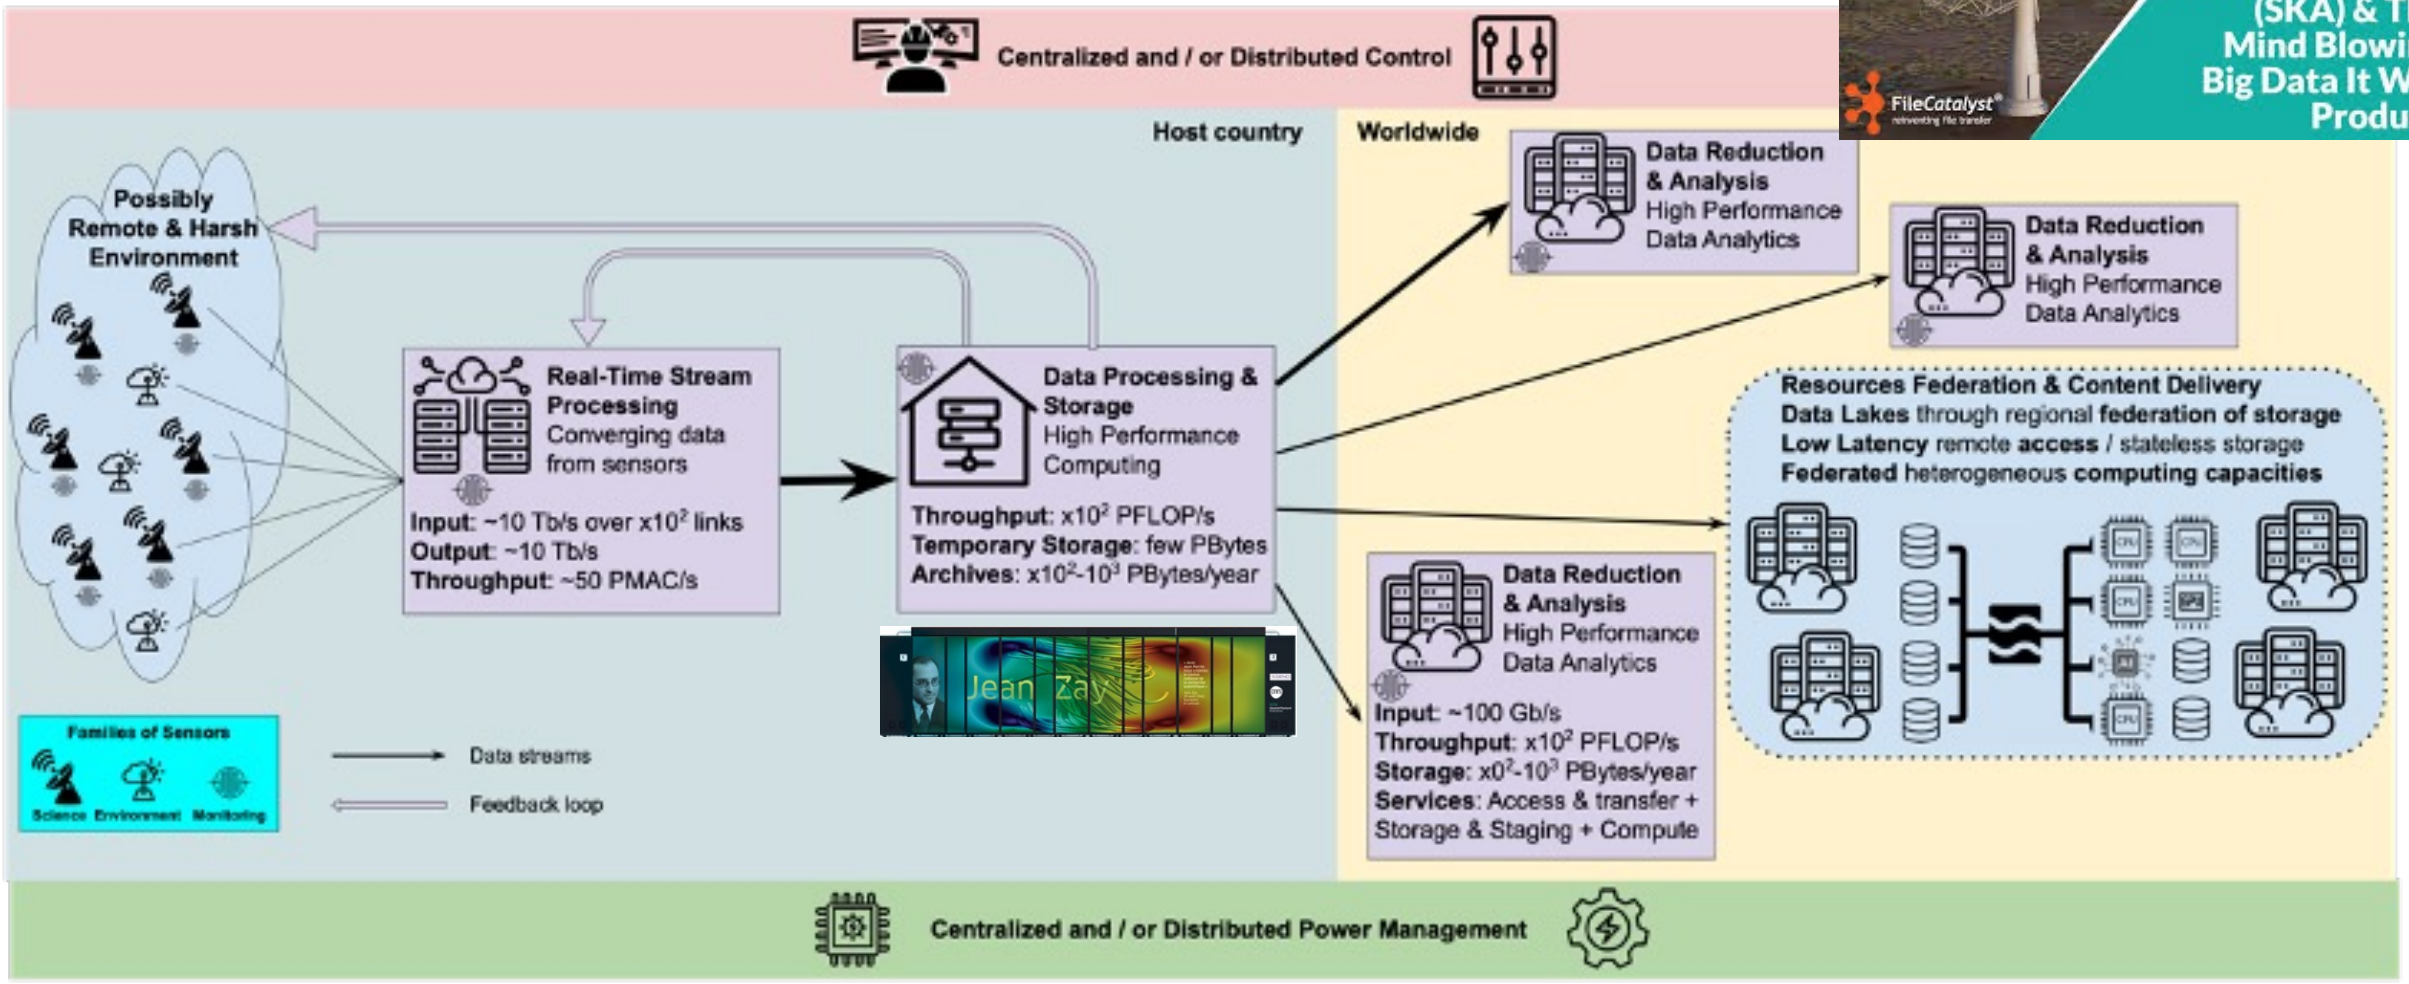

Observation data product delivery
- Data streaming reduction and processing,
- Edge computing and content delivery network

Data Science analysis: R&D component & software development
- Emerging technologies compute, I/O, storage, wide area workflows
- Distributed & heterogeneous architecture, content delivery network
- Power management

The Edge-Cloud-HPC continuum

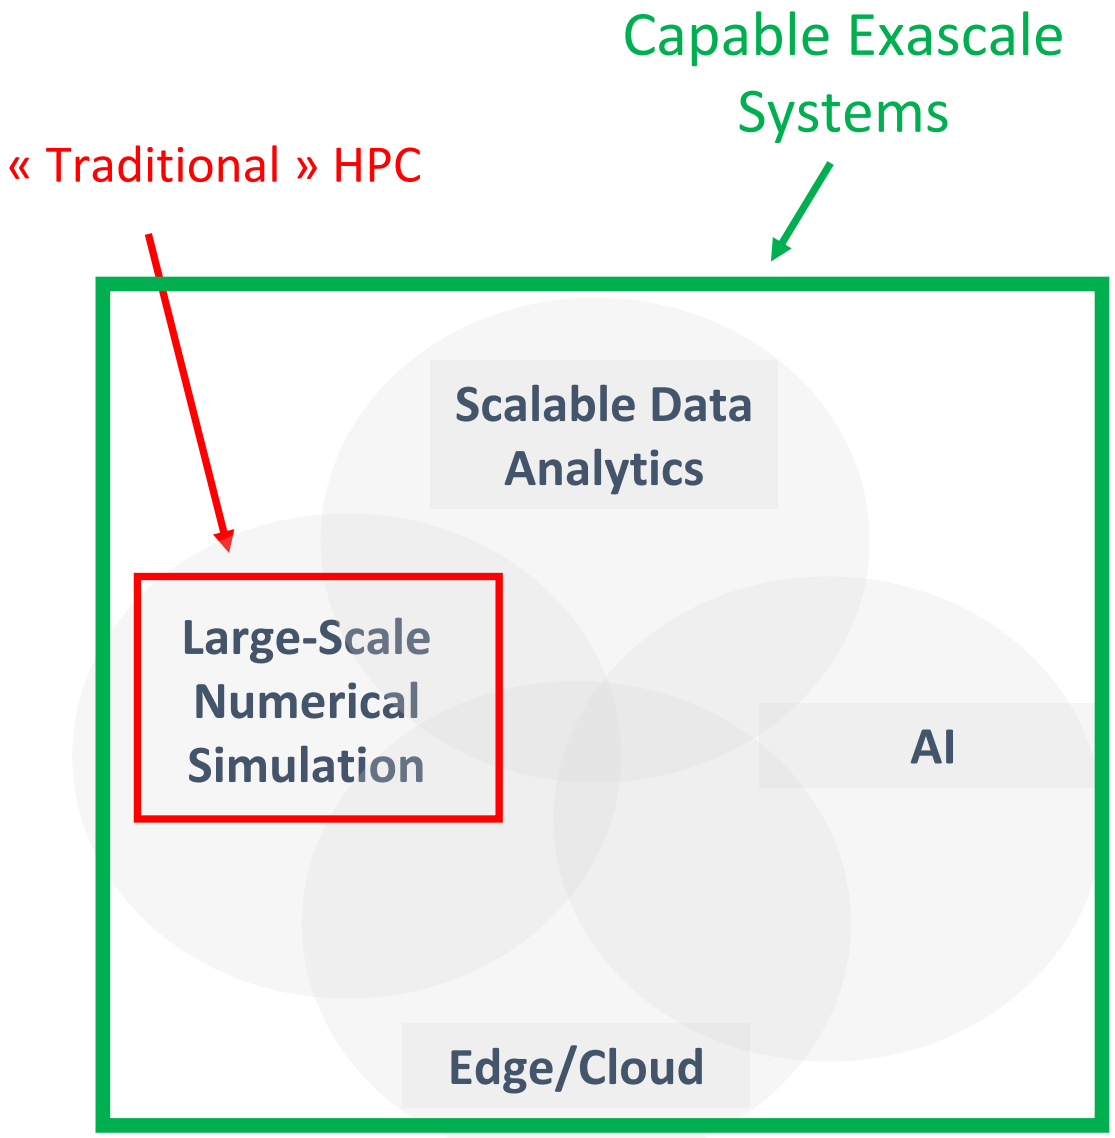

- (post)Exascale machine is part of a “datasphere”
  - Data from large scientific instruments
  - Data from sensors (IoT)
  - Data from simulations

- The primary challenges:
  - developing integrated software ecosystems across this continuum
  - At an extreme scale with stringent **cybersecurity** constraints
  - Develop the concept of **Exascale as a Service**
  - development of a data-everywhere, **FAIR** ecosystem
  - Identifying **new emerging usages** including urgent computing and digital twins

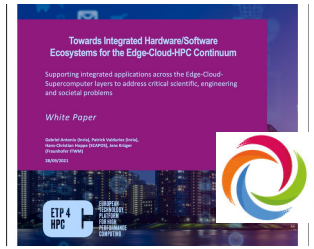

The International Post-Exascale Project

InPEx

- Landmark documents largely exploited, worldwide
- Contribute to the **implementation** of an international, shared, high-quality **computing environment based** on the principles and practices of co-design
- Formation of a solid network of exascale computing leaders, all around the globe

Date	(10/2023)	11/2023	06/2024	06/2025
Location	Preparatory phase EU (France)	SC'23 - BOF	Workshop1 EU/BSC	Workshop2 Japan
Date	03/2026	09/2026	06/2027	09/2027
Location	Workshop3 US	Workshop4 EU	Workshop5 Japan	Workshop6 US

The European Software Post-Exascale initiative

SPE-EU

- Structure, formalize and unify a common European **vision**, including **academia and industry**
- Operate and speak in a coherent and distinct way in the world-wide international InPEx project

=> Contribute to the foundation to support a healthy, sovereign and sustainable European post-Exascale computing ecosystem

SPE-EU

Proposition for a dedicated European Software Post-Exascale agenda

Contributors (March 2024)

Spain-BSC : Sergi Girona, Rosa M Badia, Fabrizio Gagliardi
Italy-CINECA : Sanzio Bassini
France-NumPEx : Jean-Yves Berthou (Inria), Thierry Bidot (Neovia), Jérome Bobin (CEA), Raymond Namyst (Inria), Michael Krajecki (CNRS)
France-GENCI : Stéphane Requena
Finland-CSC : Kimmo Koski, Per Oster, Damien Lecarpentier, Janne Ignatius
Germany-JSC : Bernd Mohr
Germany-GCS : Claus-Axel Mueller
Germany-LRZ : Martin Schulz (TUM)
Germany-HLRS: Michael Resch
Germany-MPG: Erwin Laure
UK-EPCC : Mark Parsons, Michele Weiland
CERN - OpenLab : Maria Girone

CSA SPE-EU, an incubator for a European Software Post-exascale program

The European Post-Exascale initiative
SPE-EU

P1 - AI for High Performance Computing (AI4HPC)

P2 - Digital Continuum and data management for post-exascale HPC

P3 - Cross-cutting post-exascale concepts, methods, algorithms, and software

- Maths and algorithms
- Programming and execution environments
- HPC building blocks for AI large and generative models
- Transversal issues (energy, disruptive tech.)

P4 - Computational Science and Engineering (CSE) Applications post-exascale ready: the post-exascale demonstrator and proxy-apps factory

CSA SPE-EU
An incubator for a European Software Post-Exascale program

1. Building a coherent post-exascale agenda and coordinate its execution.
2. Coordinating software production good practices, integration and management.
3. Contributing to the post-exascale application domain roadmap
4. Synchronizing with other EU initiatives
5. International coordination (including TPC)

Take away messages

Software, the new frontier

Consolidating and accelerating the construction of a exascale software stack (portable, interoperable, reproducible, sustainable)

Support and foster the developpement of disruptive Math & models and prog./excecution environnement

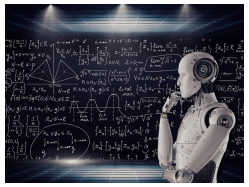

AI4Science – Science4AI

Push an **hybrid AI/HPC software stack**, to accelerate HPC and provide AI at scale

Support AI for Science, foster fully open AI use-cases/benchmarks, **not restricted to GenAI**

From edge to HPC system: the digital continuum

Coordinate efforts to share workflows, solutions and services for the convergence of HPC/Cloud/Edge

EaaS: **Exascale as a Service**, for Tier-0 systems

Develop a data-everywhere, FAIR, ecosystem

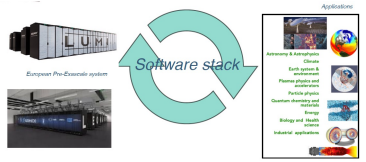

Software/application co-design

HW/SW/application co-design to help the communities get prepared for post-Exascale

Foster the use/reuse of modular/interoperable and portable SW components

Push sustainable SW development model

=> need for proper regional agendas and international parternship