Proposition for a dedicated European Software Post-Exascale program

SPE-EU

May 2024

Authors :
France-NumPEx program: Jean-Yves Berthou (Inria), Thierry Bidot (Neovia), Jérome Bobin (CEA), Emmanuel Jeannot (Inria)**,** Michael Krajecki (CNRS), Jean-Pierre Vilotte (CNRS)
France-GENCI: Stéphane Requena
Spain-BSC: Sergi Girona, Rosa M Badia, Fabrizio Gagliardi
Italy-CINECA: Sanzio Bassini
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

High Performance Computing (HPC) have recently breached the Exascale barrier opening a new computational era. This new era, however, does not only manifest itself in unprecedented raw computational power, but also in the increasing convergence of major directions, in particular heterogeneity of computing platforms, traditional modeling and simulation, the use of AI in science and technology and the reliance on big data across disciplines, which drives a change in the usage models. Nevertheless, the growth in computational demand is unbroken and even seems to accelerate, as existing users communities learn to harvest the new opportunities of such converged systems, and – at the same time – new communities unlock previously unseen application areas. This ensures that HPC, as it advances towards post exascale performance, continues to open new opportunities for innovation and progress in a wide range of science disciplines, and becomes an even greater factor in the development of sovereign solutions for industry and society.

However, this growth up to Exascale as well as the evolution to converged systems, has come at the price of unprecedented complexity in system architecture, infrastructure demands, middleware as well as algorithm development. With the end of Moore's law finally in sight, this trend drives us to even more complex solutions that embrace accelerator technologies, heterogeneity, and specialization, not only in computing, but also memory, storage, and network technologies... This is further exacerbated by the need for the development of more energy-efficient and more sustainable systems, the complement of classic 64-bit numerics with AI-driven models and reduced precision computing, all the way to revolutionary approaches such as quantum and neuromorphic architectures. This very fast evolution further is highly challenging for the applications and calls for a radical change in code development methodologies and frameworks.

In such a drastically changing compute ecosystem, existing design methods will seize to be effective – continuing with business as usual is not an option. Instead, we need to build on, if not re-imagine, truly foundational co-design principles bringing together stakeholders across the entire system stack: from application- to hardware-designers, from algorithm theorists to runtime system developers. Only this will enable us to not just attempt to exploit compute resources efficiently, but to design them from the ground-up in a way that they can be adjusted at all levels, from design and construction to configuration and runtime, so that we land on systems that can be efficiently exploited and keep providing the computational capacities that are required and expected.

At the European level, it is of strategic importance to structure, formalize and unify a common European vision to, on one hand, help build our short-, medium- and long-term path, while on the other hand to allow Europe to operate and speak in a coherent and distinct way with our international co-competitors in US and Asia. Only this will allow Europe to to support a healthy, sovereign and sustainable European post-exascale computing ecosystem, which is vital for European industry and society.

We identified 4 main technical pilars P1 to P4 to structure the construction of a European Post-Exascale (SPE-EU) program supported by one coordination action CSA SPE-EU:

- P1 - AI for High Performance Computing (AI4HPC)
- P2 - Digital Continuum for post-exascale HPC
- P3 - Cross-cutting post-exascale concept methods, algorithm, and software
- P4 - Computational Science and Engineering (CSE) Applications post-exascale ready: the post-exascale demonstrator and proxy-apps factory
- CSA SPE-EU - Coordination and international collaboration

CSA SPE-EU, Coordination and international collaboration

A. Building a coherent post-exascale agenda and coordinate its execution.

The first goal of this CSA is to consolidate a coherent dynamic agenda for post-exascale research and applied topics and coordinate its execution (P1 to P4 technical pilars). This will be done in close collaboration with the European stakeholders such as:

- ETP4HPC in charge of elaborating an SRA for EuroHPC,
- Tier-0 sites such as Julich, HLRS, LRZ, BSC, CSC, CINECA, GENCI/CEA
- Research organizations dedicated to digital sciences and technologies such as BSC, Inria, Julich, MPG, CINECA, CSC, CEA, Cerfacs, CNRS, ESRF, etc,
- Industry dedicated to digital sciences and technologies, such as SiPearl, Eviden, DDN, VAST, AMD, Intel, Nvidia, Infineon, LightOn, Mistral AI, Aleph Alpha, Silo AI and EU cloud providers such as Scaleway and OV,
- Industrial end users such as, Airbus, Safran, EDF, Audi, Ericsson, BASF, BMW, TotalEnergies, ENI, Repsol, Noki,
- Research organizations dedicated to specific application domains such as CERN, SKA, EuroFusion, EBI etc.

Some interactions with UK institutions involved in the Exascale race (such as the Excalibur project) already exist. The participation of UK partners should be discussed.

B. Coordinating software production good practices, integration and management.

Exascale and post-exascale applications and machines are gaining in complexity leading to high costs to build and deploy applications on exascale and future post-exascale infrastructures, impairing portability, deployability, maintainability and reproducibility. The standard software installation process on supercomputers is reaching its limits. In this context, the **CSA SPE-EU will have the following objectives:**

- capitalize on and consolidate software production in Europe while ensuring software sustainability using rigorous software development practices
- Supporting and actively drive efforts in standardization
- Fostering the use of HPC dev-ops methodology for fast integration and deployment while ensuring performance and supporting reproducibility efforts through modern software packaging and testing
- Contributing to an "industrial" production environment converging towards "As a Service”

The **CSA SPE-EU** will coordinate efforts to evolve and harden the European Exascale software ecosystem, from low level system software all the way up to full applications and use it as the basis for a post-exascale software stack, using existing software components developed in and/or driven by European partners, and integrate those into future European projects.

C. Contributing to the post-exascale application domain roadmap. The post-exascale demonstrator factory (P4 technical pilar) will demonstrate the capacity of cross-cutting post-exascale concepts, methods, algorithms and software to address the post-exascale challenges the main application domains will face. This demonstrator factory will further support application domains to build their own post-exascale roadmap. The CSA SPE-EU will have the mandate to make this possible, paying specific attention to the CoEs funded or to be funded.

D. Synchronizing with other EU initiatives. EU has many initiatives (past and present) that are relevant to the post-exascale vision: EOSC, EUDAT, AARC Blueprint Architecture, CoEs, etc. The CSA will ensure the coherency of the post-exascale vision with such and future EU efforts.

E. International coordination. The SPE-Eu program need to be coordinated with its international counterparts in the US and Asia. This the fifth objective of the CSA SPE-EU.

Building on earlier efforts of the International Exascale Software Project (IESP), the European Exascale Software Initiative (EESI, EESI2), the European EXtreme Data and Computing Initiative (EXDCI) and the BDEC community, the world-wide InPEx initiative aims to work on the implementation of an international, shared, high quality computing environment that focuses on co-development of reusable and sustainable science-driven software components principles and practices. Within the International Post-Exascle (InPEx) project, Europe (EuroHPC, French NumPEX project, BSC, JSC, CINECA, CSC, GENCI), USA (DOE, NSF, ANL), Japan (FugakuNEXT, Riken-C) have already met (Inpex) and identified a set of areas for international coordination and decided to launch a new series of workshops and working groups dedicated to international collaborations between Europe, USA and Japan on Exascale and post-exascale computing. The CSA SPE-EU will coordinate the European contribution to **InPEx** and will have in this context the following objectives:

- Produce landmark documents largely exploited, worldwide, for supporting future post-exascale science in strong collaboration with ETP4HPC
- Contribute to the implementation of an international, shared, high-quality computing environment based on the principles and practices of co-development of software components
- Form of a solid network of exascale computing leaders, all around Europe, including EuroHPC, ETP4HPC and industrials representatives (Sipearl, Eviden, DDN, Safran, Airbus, EDF, Infineon, Ericsson, Audi, Partec, LightOn, Mistral AI, E4, OpenChip, Silo AI, Nokia, BMW, etc.)
- Study strategic technical issues at European level and contribute to the corresponding dedicated InPEx working groups
- Promote the mobility of researchers, particularly young researchers, between Europe, the United States and Asia.

TPC, the Trillion Parameter Consortium, is an international initiative that focuses on sharing expertise, data and foster collaborations towards building Large Foundational Models. This major effort towards extremely scale AI is ongoing with private as well as several (trans-) national public initiatives thus recognizing the crucial importance of AI for Science and Society. The CSA SPE-EU will coordinate the European contribution to **TPC.** Moreover, this European coordination of TPC operated the CSA SPE-EU can lead two specific objectives that are key to implementing an “international, shared, high-quality computing environment” in the scope of IA4HPC P1 pilar:

1. Contribute to the growth and sustainability of a solid international HPC & AI community by organising international conferences, topical workshops, and hackathons to foster networking between HPC & AI researchers and practitioners and drive the WGs’ collaboration efforts forward.
2. Participate in policy dialogues to develop and foster access in Europe and internationally to HPC infrastructures in the context of the convergence between

HPC systems and resources federation, and the EU AI factories initiative. TPC members actively involved in the WGs' collaboration efforts can advise policymakers at the EU and national levels on how to re-orient HPC systems adapting them to AI development requirements and how this Europe efforts can integrate and benefit from the international dimension of TPC.

Four main technical pilars to structure the construction of a European Post-Exascale (SPE-EU) program

1) P1 - AI for High Performance Computing (AI4HPC)

New uses of HPC in the Exascale era raise methodological and software engineering challenges. This includes AI model learning, physics-based AI and hybrid simulation models, observational data reduction, distributed AI models, intelligent analysis of in-stream or in-situ data, coupling between observational data and digital simulations at work in digital twins, etc.

As this ecosystem gathers massive experimental and observational data, large computing power, and expert practitioners, it offers the opportunity to develop new uses of AI to tackle complex scientific questions. Indeed, recent advances in AI have shown that the availability of these three ingredients is the core component required to develop powerful models. AI is seen as a game-changer, but that **remains often at the proof-of-concept** level and thus needs a clear leap forward a higher maturity with full-fledged scientific validation and implementation at scale.

In this context, the European Software Post-Exascale (SPE-EU) program will serve as a cornerstone for the entire HPC/HPDA/AI ecosystem, with close connections to national supercomputing centers, computer science/engineering and application research communities and vendors.

We list below some scientific and technical challenges to be addressed:

- Define and promote a European strategy for large and complex data processing, Foundation Models (such as Large Language Models), software tools and computing capacities dedicated to AI4HPC.
- Foster the development of robust and trustworthy AI for scientific computing applications (Science Foundation Models) specific to European needs and requirements.
- Develop a European converged HPC and AI software stack with common, interoperable software components and tools.
- Foster AI-centric solutions in scientific and engineering applications
- Help aggregate scientific communities from core-AI and computer science to application domains.
- Investigate AI-centric application co-development of software components with use-case sharing.
- Investigate AI assisted tools for scientific code generation, code robustness and code efficiency.

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

This pilar focuses on the continuum of technologies and practices related to managing data from collection to processing and storage, including edge computing and cyber-physical systems.

Currently, the Digital Continuum is primarily being driven by large cloud providers, while high-performance computing (HPC) centers are competing with these providers to offer alternative solutions for managing and processing large datasets. One of the challenges facing HPC centers is demonstrating their value proposition compared to cloud providers, especially when it comes to real-time data processing requirements.

Another challenge for the Digital Continuum is that it is a multi-tenant environment, where collected data is used for multiple purposes and computing infrastructure is shared. As such, building trust in this continuum of entities is crucial and the related issue of cybersecurity becomes central.

Digital Continuum and Data Management encompasses:

- Post-exascale as a cloud/federation of exascale systems
- Data-centric workflow composition across hybrid HPC-Cloud-Edge infrastructures
- Sustainable Interoperability Across the HPC-Cloud-Edge continuum (requiring middleware able to interoperate and execute in hybrid scenarios)
- Develop frugal architectures for computation and storage for post-exascale workflows executed across the continuum
- Develop the concept of EaaS, Exascale as a Service, for Tier-0 European systems
- Develop a data-everywhere, FAIR, ecosystem (in Europe)

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

The evolution of Top500 supercomputers reflects the current trend towards massively accelerated architectures composed of heterogenous nodes featuring numerous super-powerful GPUs. The average number of nodes is actually decreasing, as the computing power of each node is increasing. Typical modern architectures include 512+ Gigabytes of host memory, hundreds of cores, 4 GPUs, and a 200/400Gbs interconnect.

This profound evolution means that the HPC community must master a new form of scaling – the explosion of GPU cores – and has to cope with all the complexity of heterogeneous programming: SIMT execution model, code generation, task scheduling, data movements between accelerators, etc. The challenges go from mathematical models and algorithms to compilers, runtime systems and programming models.

Maths and algorithms:

- new solvers,
- mixed precision,
- innovative discretization methods,
- energy aware algorithms,

Programming and execution environments:

- Converge towards common high-productivity programming interfaces:
    - Hide machine complexity to improve productivity, portability, in particular performance portability, and composability.
    - Support transparently asynchronous parallelism with dynamic optimizations to improve scalability.
    - Utilize AI+ i.e., AI-assisted code generation for robustness, efficiency, etc.
- Extend the mini-app ecosystem to guide the development of scalable and efficient software
- Offer an "industrial" production environment converging towards "As a Service",

HPC building blocks for AI large and generative models:

- HPC contribution to the design, optimization, energy aware, efficiency of SLM to LLM, domain specific foundation models, Multimodal models on post-exascale systems
- Programming model, execution models for AI at post-exascale
- parallel differentiable programming for scientific computing and AI at scale
- identify the right AI methods/algo for a given problem so that resources consumption is minimized

- high performance stochastic methods, algorithms, software and frameworks for next gen AI, scientific computing and data analytics.

Transversal issues:

- Investigate and improve energy and environmental impact and sustainability
- Identify future and disruptive SW&HW technologies and usages

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

Capable computing ecosystems need to integrate CSE exascale and post-exascale applications, software technologies and hardware innovations with an expanded software stack that can be easily deployed onto exascale and post-exascale facilities. For these to be successful, we require innovations in productivity, performance portability and sustainability. This in turn has the build on improvements in scientific software development methodologies, the adoption of new mathematical approaches, algorithmic or model improvements, and software development methodologies leveraging optimised libraries and frameworks, extreme-scale programming and execution environments and data logistics. Only this will help improve exascale capabilities and thus ensure new science and engineering insights.

Preparing strategic application codes for the Exascale and post-exascale requires substantial engineering efforts that the scientific communities alone can generally not afford. Tangible development progress requires close coordination among application, algorithm and software development to address key application development challenges. This represents a major and essential effort, and doing this independently for each application, and/or repeating it for the same application as technology changes is not an option. Preparing applications for the current and future extreme-scale computing systems builds upon two pillars: i) the co-development of software components, which can be re-used across applications, and ii) a strong commitment to the application communities to support the evaluation, testing and integrating of new software components through shared proxy-apps. This will lower the complexity of CSE application software by co-developing cross-cutting efficient software building blocks (libraries, frameworks, tools) , as well as to improve software development methodologies in the application communities.

- Organize and mutualize application code development effort within co-development initiatives, centered about mutualized, cross-cutting software and algorithmic building blocks.
- Co-identify and co-development of common software and algorithmic building blocks with the mathematics and computer science communities and experts common.
- Foster the co-development of modular, interoperable and re-usable software components that can be used adaptively to answer the domain-specific requirements and needs.
- Develop shared proxy-apps and benchmarks together with science-driven performance evaluation methodologies and results to support future system software and architecture developments, as well as to drive upcoming European procurements.
- Support the integration of logical collection of co-developed, post-exascale ready, software components in the CSE application codes to improve their performance and portability in collaboration with application owners and application efforts like the CoEs.
- Train the application communities to improved development methodologies to leverage high quality software components, portable programming and execution models.