

NumPEx “white paper”

from Exascale to post-Exascale computing

I – The NumPEx roadmap

1. Context
2. NumPEx vision and strategy

II – Detailed scientific programme

1. Exa-MA : Methods and Methods for Exascale
 - i. Context and challenges
 - ii. Project workplan
 - iii. Software ecosystem and contributions
 - iv. Internal organization, cross-project collaborations and deliverables
 - v. Post-PEPR scientific prospective
 - vi. Involved research teams
 - vii. Budget
2. Exa-Soft - HPC softwares and tools
 - i. Context and challenges
 - ii. Project workplan
 - iii. Software ecosystem and contributions
 - iv. Internal organization, cross-project collaborations and deliverables
 - v. Post-PEPR scientific prospective
 - vi. Involved research teams
 - vii. Budget
3. Exa-DOST - Data-oriented Software and Tools for the Exascale
 - i. Context and challenges
 - ii. Project workplan
 - iii. Software ecosystem and contributions
 - iv. Internal organization, cross-project collaborations and deliverables
 - v. Post-PEPR scientific prospective
 - vi. Involved research teams
 - vii. Budget
4. Exa-ATOW - Architectures and Tools for Large-Scale Workflows
 - i. Context and challenges
 - ii. Project workplan
 - iii. Internal organization, cross-project collaborations and deliverables
 - iv. Post-PEPR scientific prospective
 - v. Involved research teams
 - vi. Budget
5. Exa-DI - Development and integration project
 - i. Context and challenges
 - ii. Project workplan
 - iii. Internal organization, cross-project collaborations and deliverables
 - iv - Involved research teams

v - Budget

III – Transverse actions across the programme

IV – Post-NumPEX prospective

I – The NumPEX roadmap

1. Context

The challenges related to high performance computing (HPC), high performance data analytics (HPDA) and Artificial Intelligence (AI) are colossal from scientific, societal (e.g. energy, health, environment), economic and financial (e.g. industrial competitiveness) and ethical (e.g. biology) points of view.

Modelling, simulation, large-scale data analysis and Artificial Intelligence are increasingly emerging as decision support tools for a number of societal issues and critical situations (natural hazards, government and industrial decision-making, etc.). They are also a major challenge for maintaining the competitiveness of R&D in France in both academia and industry.

All the areas of science, technology and industry are therefore now, to varying degrees, concerned with the use of digital resources (computing, storage, data processing, etc.) linked to intensive processing (requiring high performance hardware resources) as in:

High performance computing making it possible to go further and further in modelling the phenomena to be analysed with multi-scale and multi-model approaches, as well as optimization/uncertainty quantification methods, or with the implementation of simulation at scales that have so far been inaccessible.

High-performance data analytics (HPDA) with heterogeneous, complex and/or massive data (either observational, experimental or sensor data or resulting from large scale simulations) and Artificial Intelligence. This is a consequence of the explosion in the production of scientific data from sensors, large instruments or experimental platforms.

And finally, the emergence of increasingly sophisticated methods and tools for analyzing and processing data, including those derived from Artificial Intelligence (machine learning or other) but also with hybrid simulations incorporating coupled numerical simulation and AI.

The answer to these needs, which are exploding in the academia and industry, is to use upcoming “exascale” computers. Petaflop computers have been installed since 2010 in the world and in Europe. Fantastic scientific improvements were made thanks to them. Now a new generation of Exascale (10^3 Petascale) massively parallel computing systems composed of several tens of millions of heterogeneous cores are expected for this decade. All international experts consider that an efficient use of these incredible capabilities could not be achieved only by incremental research but by outstanding technological breakthrough possibilities, opening unknown areas in designing new products or optimizing existing ones in almost all societal domains.

In this context, the French Exascale program NumPEX aims at designing and developing the software components that will equip future exascale machines. NumPEX will deliver Exascale-grade numerical methods, softwares, and training, allowing France to remain one of the leaders in the field. It will contribute to take bridging the gap between cutting-edge software development and application domains to prepare the major scientific and industrial application codes to fully exploit the capabilities of these machines. Application domains of the NumPEX program include, but are not limited to, weather forecasting and climate, aeronautics, automotive, astrophysics, high energy physics, material science, energy production and management, biology and health.

In June 2023, EuroHPC announced that a second exascale system will be hosted in France. More precisely, it will be managed by GENCI (as hosting entity), the French national agency for High Performance Computing, and operated at the TGCC computing centre by the CEA (as hosting site), the French Alternative Energies and Atomic Energy Commission, in Bruyères-le-Châtel. NumPEX contributed, as its software development component, to the Jules Verne consortium response to the next EuroHPC call for expressions of interest, with a view to hosting and operating at CEA/TGCC one of the two European exascale machines planned in Europe by 2025.

2. NumPEX vision and strategy

The NumPEX programme is a 6-years/40,8M€ project that started in 2023. It has a total budget of 81 M€, including permanent staff, and gathers more than 500 researchers and engineers nation-wide. It originates from an objective analysis of the current status of the HPC/HPDA community at international and national level. One of the key drivers of the NumPEX programme is the current shift of paradigm of the HPC system architectures with fastly emerging new technologies and new usages (e.g. digital continuum and AI), which mandate developing and adapting the HPC software stack for the Exascale supercomputer to come. This also includes preparing for the post-Exascale systems and usages.

Furthermore, the exploitation by academic and industrial researchers of such a leading-edge, sovereign supercomputer, with its potential for innovation and competitiveness, requires a lot of work to adapt and port the simulation codes because of the mostly accelerated nature of the architecture, and also with the integration of new computation models such as those used in high performance data analysis and AI (data analysis, implementation of Machine Learning and Deep Learning models, etc.). Often the numerical methods and algorithms are not suitable to such architectures and require research efforts to investigate new approaches.

These development efforts and this research are interdisciplinary by nature and can only be carried out by IT specialists, computer scientists, applied mathematicians and scientists in the application fields and may require a few months up to a few years of research and integration and validation work.

It should be kept in mind that the European and international competitors have, to a very large extent, already provided strong support for applications in their own actions. It is also important to note that investments in applications make it possible to capitalize knowledge and expertise over the long term, beyond the lifetime of a supercomputer. Sovereignty and scientific and industrial competitiveness involve technological sovereignty and mastery of applications and various software stacks.

The adaptation of applications to these new architectures is therefore a major scientific and technological issue that requires both strong research efforts and support. This involves rethinking the algorithmics, the associated numerical methods, the way to express and execute parallelism, or even the modelling itself of the studied physical phenomena, while taking into account the evolution of data flow-oriented uses and the implementation of AI approaches.

Moreover, the supercomputer is becoming now a component embedded into a digital continuum formed by a chain of processes guided by the flow of data. It becomes essential to evolve the software layers and tools to design and manage the secured execution of complex application end-to-end workflows that can be deployed at the edge on a set of equipment integrating sensors, to various data storage sites and computing facilities, including heterogeneous supercomputers. If the proposed technological solutions are to achieve the desired power levels in sustainable energy envelopes, energy efficiency aspects are also fundamental in a global approach.

Finally, it is of strategic importance for France and Europe to train the future engineers and researchers who will design and produce future exascale and post-exascale systems, as well as the software components and applications that will benefit from their performance. This is why NumPEX is also leading a global reflection aimed at identifying and designing the future corpus of exascale and post-exascale training courses, in close collaboration with our European and international colleagues.

Therefore, the NumPEX programme has six major objectives:

- **Providing a major contribution to the European Exascale software stack**, with a special focus on the Jules Verne project that will yield the second European Exascale system in 2025. To that respect, NumPEX is a production with a clear short-term deadline to deliver an exascale software stack.
- **Preparing the building blocks for post-Exascale software solutions**: in the long-term, NumPEX aims at exploring innovative software solutions to cope with the fast evolution of complex HPC systems, the growing data flow-oriented uses and the implementation of AI approaches.
- **Preparing the academic and research applications to enter the Exascale era**: whether they originate from academic or industrial applications, only few codes can be deemed Exascale-ready. Adapting codes

for the Exascale era requires major developments that need years for whole teams for developers to be accomplished, which are quite impossible in most cases. In this context, and in the spirit of the ECP co-design approach, NumPEX aims at preparing applications by co-designing algorithmic and software motifs that can be mutualized across applications.

- **Contributing to the creation of a French Exascale community:** This exploratory program is positioned upstream of the applications with the ambition of covering all the software layers from as close as possible to the machine up to the applications in order to allow a major leap in the efficient use of exascale and post exascale machines. This initiative is interdisciplinary by nature; working in silos is not an option. It therefore requires gathering all the actors for HPC/HPDA/AI from fundamental research, software engineers to the hardware vendors and the application teams. It therefore also essential to contribute to the creation of a French Exascale community including public and private actors. For that purpose, a key action of the NumPEX programme will be on training from HPC experts to early-training students.
- **Designing the future corpus of exascale and post-exascale training courses:** NumPEX is leading a global reflection aimed at identifying and designing the future corpus of exascale and post-exascale training courses, in close collaboration with our European and international colleagues
- **Fostering national, European and international collaborations:** NumPEX will further aim at creating and fertilizing collaboration at national and European level with existing initiative, such as French PEPRs, EuPEX, EPI or the CoE to only name a few. Cross-fertilisation with the international community is essential to prepare for Exascale and post-Exascale era. To that end, NumPEX will actively participate to the organization of international events with collaborators Europe, Japan and the USA.

To that purpose, the NumPEX programme builds upon four scientific projects Exa-MA, Exa-Soft, Exa-Dost and Exa-Atow. A fifth project, coined Exa-DI, will focus on co-design to bridge the gap between the scientific and software development teams and the applications.

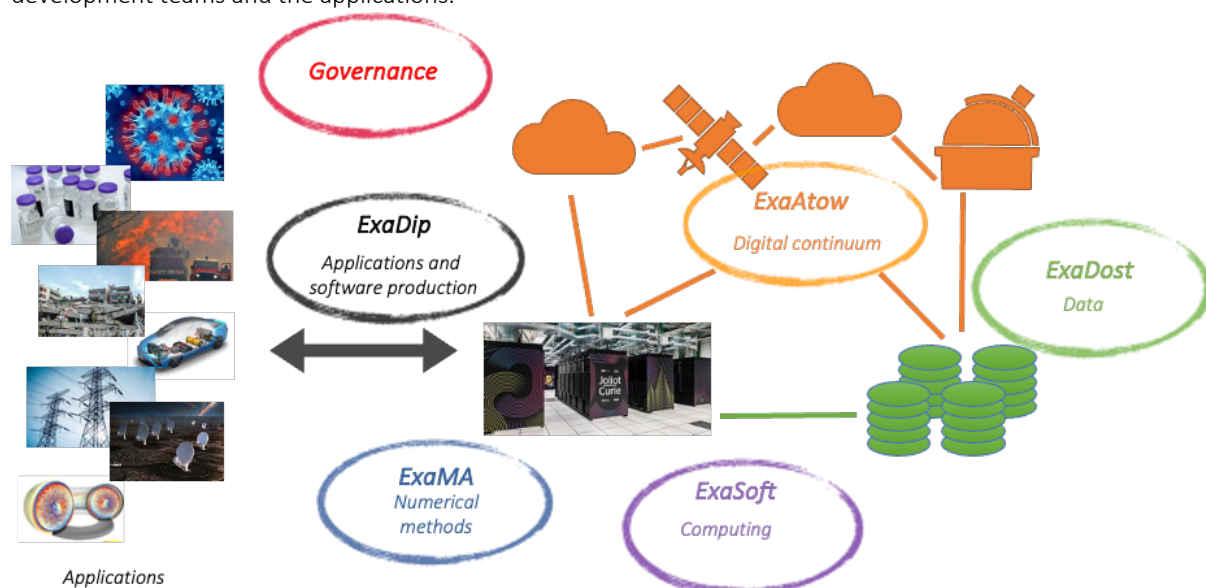


Figure 1 - Sketch of NumPEX's organisation

II – Detailed scientific programme

1. Exa-MA : Methods and Methods for Exascale

i. Context and challenges

There is a growing number of problems where experiments are impossible, hazardous, or extremely expensive. Extreme-scale computing enables the solution of vastly more accurate predictive models and the analysis of massive quantities of data. More than that, it enables the possibility to create a digital copy of a physical asset that can be fed with data to understand, improve or fix the latter. These challenges include: **(C1)** Reduce carbon footprint (GHG) in transportation, buildings and cities; **(C2)** Design, control, and manufacture of advanced materials; **(C3)** Understand and simulate the human brain; **(C4)** Understand fission and fusion reactions and design advanced experiment facilities for fusion; **(C5)** Monitor the health of our planet (climate prediction, impact assessment of environmental policies, etc.); **(C6)** Monitor and personalize the health of human beings; **(C7)** Design better drugs; **(C8)** Design cost-effective renewable energy resources (batteries, biofuels, solar photovoltaics, etc.); or more generally **(C9)** Understand the Universe. These challenges require tremendous computing power to understand them and help decision makers.

Exascale computing is the next frontier to unlock new discoveries. We face, however, new bottlenecks as we reach these computing facilities including **(B1) energy efficiency**: develop energy efficient technologies to meet the at most 20 MW target. **(B2) interconnect technology**: improve vertical (intra-node) and horizontal (inter-node) data movement in terms of energy efficiency and performance. **(B3) Memory technology**: integrate new memory technologies (e.g., PCRAM, NOR Flash, ReRAM, memristor) to improve capacity, bandwidth, resiliency, and energy efficiency. **(B4) Scalable system software**: Increase the scalability, power sensitivity, and resiliency of system software (e.g., operating systems, runtime systems, monitoring systems). **(B5) Programming systems**: develop new programming paradigms to express fine-grained concurrency, locality, and resilience. **(B6) Data Management**: develop software that can handle massive amounts of data—this concerns both offensive I/O (e.g., data analysis & compression) and defensive I/O (e.g., fault tolerance). **(B7) Exascale Algorithms**: redesign algorithms to improve scalability (e.g., reduce communication, avoid/hide synchronization) and computational efficiency on accelerators. **(B8) Discovery, design, and decision algorithms**: Research should focus not only on "single heroic simulations" but also on ensembles of many small runs (e.g., common for uncertainty quantification or parameter optimization). **(B9) Resilience, robustness and accuracy**: Computations must be correct, reproducible and verifiable, even in the presence of software and hardware errors (hard and/or soft error). **(B10) Scientific productivity**: scientists must have the tools to use exascale systems productively (e.g., develop programs, run applications, prepare inputs, collect outputs, analyze results). **(B11) Reproducibility, replicability of computation**: reproducibility is an essential ingredient of the scientific enterprise. The ability to reproduce results builds trust so that we can rely on the results as foundations for future scientific exploration. Presently, the fields of computational and computing sciences provide two opposing definitions of *reproducible* and *replicable*. In computational sciences, reproducible research means authors provide all necessary data and computer codes to run analyses again, so others can re-obtain the results. The concept was adopted and extended by several communities, where it was distinguished from replication: collecting new data to address the same question, and arriving at consistent findings. **(B12) Pre/Post processing**: visualization, in situ processing. **(B13) Opportunity to integrate uncertainties directly into the core of the calculation** (unseen).

If the bottlenecks **(B1-B6;B12)** are to be tackled at the methods and algorithms level through transverse collaborations within the PEPR NumPEX, Exa-MA will directly address **(B7-B11;B13)** and thus its **main objectives** are **(O1) to develop methods, algorithms, and implementations that, taking advantage of the exascale architectures**, empower modeling, solving, assimilating model and data, optimizing and quantifying uncertainty, at levels that are unreachable at present; **(O2) to develop or contribute to software libraries** allowing to assemble specific critical reusable components, hiding the hardware complexity and exposing only the specific methodological interface; **(O3) to identify and co-design Methodological and Algorithmic Patterns** at exascale that can be reused efficiently in large scale applications (e.g., in weather forecast); **(O4) to enable AI algorithms** to attain performances at exascale, exploiting the methods **(O1)** and the libraries **(O2)** developed; and **(O5) to provide demonstrators** through mini-apps and proxy-apps that will be openly available and benchmarked. The figure below provides an overview of Exa-MA and its connections to the other PCs.

ii. [Project workplan](#)

Exa-MA is a research project aimed at advancing scientific simulations and modeling capabilities to reach and surpass the exascale barrier. The project is organized into several work packages (WP) that focus on different aspects of the research objectives. The following is an updated overview of the workplan for each work package in Exa-MA.

Work Package 0 (WP0): Project Management

Responsible for project management, technical and scientific coordination, and administration. Ensures effective governance, coordination, and communication within the project, establishing a solid foundation and facilitating collaboration.

Work Package 1 (WP1): Geometric and Physics-based Modeling

Focuses on revisiting methods and algorithms for large-scale models, incorporating multiple phenomena and process couplings. Tasks include mesh generation, adaptive refinement, finite element frameworks, and efficient parallelization for time integration and multiphysics coupling.

Work Package 2 (WP2): Surrogate Models and Machine Learning

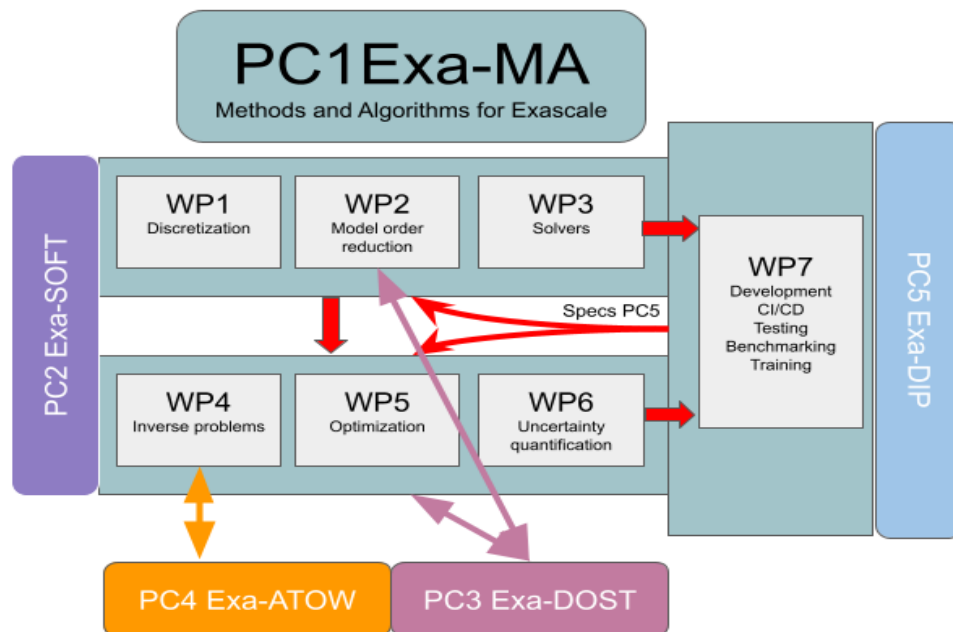
Aims to develop reduced order models using physics-driven deep learning ANNs and surrogate-based optimization. Tasks involve novel algorithms, neural operators, data-driven model reduction, and multi-fidelity modeling for improved accuracy and efficiency.

Work Package 3 (WP3): Numerical Kernels and Coupled Solvers

Focuses on designing and implementing efficient and possibly provable numerical kernels and solvers for large-scale problems. Tasks include domain decomposition methods, data sparsity techniques, multiple precision, adaptive solution strategies, and efficient coupling of multiphysics simulations.

Work Package 4 (WP4): Inverse Problems and Data Assimilation

Addresses formulation and solution of inverse problems and data assimilation. Tasks involve deterministic and stochastic methods, modeling and reducing uncertainties, multi-fidelity models, and improved observation



strategies.

Work Package 5 (WP5): Optimization

Focuses on exascale optimization algorithms for solving large-scale problems. Tasks include exact and approximate algorithms, surrogate-based optimization, shape optimization, and optimization for AutoML. Addresses challenges of optimization in AutoML.

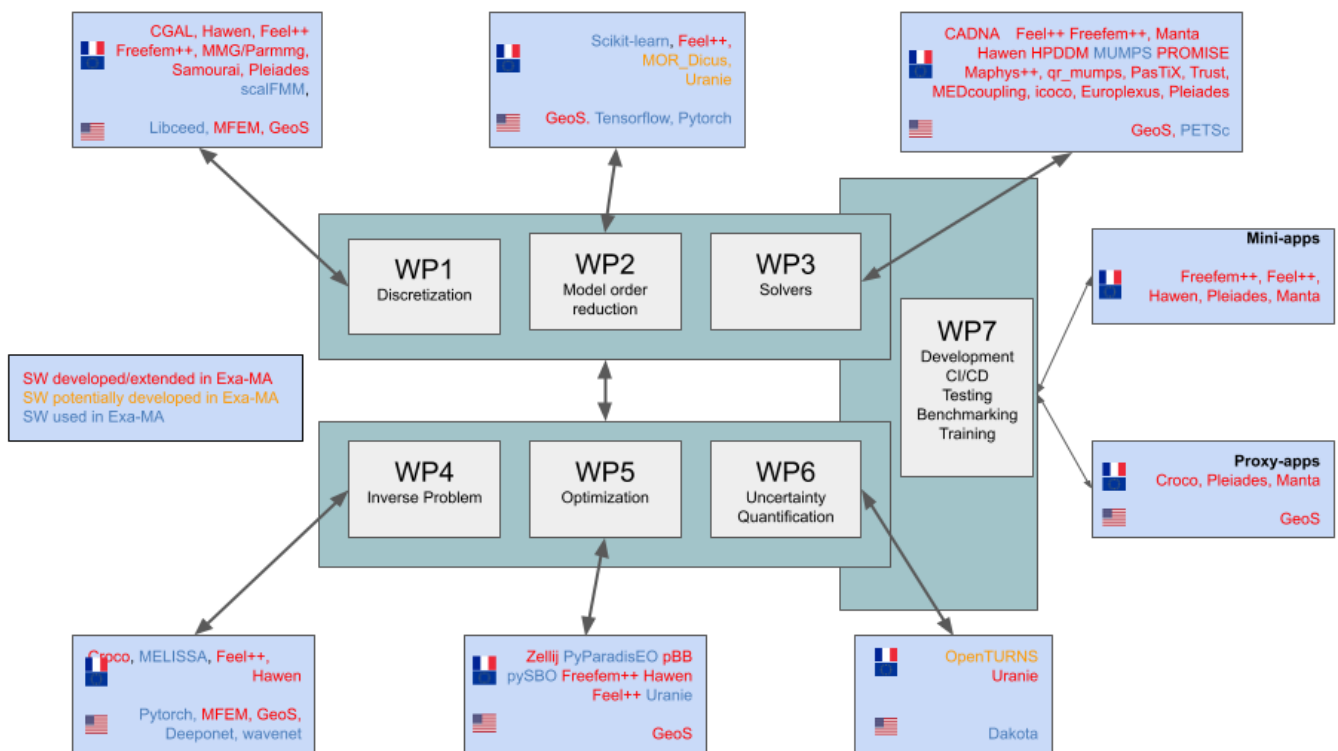
Work Package 6 (WP6): Uncertainty Quantification

Dedicated to uncertainty quantification (UQ). Tasks include sensitivity analysis, UQ in PDE solving frameworks, surrogate modeling for complex multi-physics problems, and accelerating UQ with exascale computing. Develops tractable UQ methodologies for high-dimensional integrals and modeling uncertainties in multiscale systems.

Work Package 7 (WP7): Software Development, Coordination and Co-Design

Focuses on software development, coordinating co-design activities, showcasing results, and creating training materials. Tasks include testing, benchmarking, coordinating activities with other projects, establishing a showroom for presenting results, and creating training materials on exascale toolboxes, mini-apps and proxy-apps. Objective is to ensure quality, integration, and dissemination of software and project results.

Overall, the scientific workplan of Exa-MA covers a wide range of research areas, including geometric and physics-based modeling, surrogate models, numerical kernels, inverse problems, optimization algorithms, uncertainty quantification, and software development. The work packages are designed to address short-term challenges and pave the way for mid- and long-term advancements in scientific simulations and modeling capabilities, ultimately pushing the boundaries of exascale computing.



iii. [Software ecosystem and contributions](#)

The Exa-MA project involves the utilization of various software tools and libraries across different work packages to address the research objectives. The figure above provides an overview of the software ecosystem associated with each work package.

iv. [Post-PEPR scientific prospective](#)

While the Exa-MA project is focused on addressing current challenges in exascale computing, it is also laying the groundwork for solutions to future challenges in the field. The development of efficient and robust algorithms, the integration of AI methods, and the focus on resilience and robustness are all areas that will become

increasingly important as we move into the era of zettascale computing and beyond. Here are some points regarding post-Exa-MA perspectives with respect to what Exa-MA builds.

1. **Next-level Scalability:** Even as the project aims to tackle exascale computing, the next frontier of zettascale computing is on the horizon. This represents a significant step up in terms of complexity and computational power, and it will necessitate the creation of even more efficient and robust algorithms.

2. **Quantum Computing Integration:** Quantum computing represents a significant shift in computational paradigm and holds the potential to revolutionize many fields, including scientific computing. While this project does not directly work with quantum computing, the development of algorithms and computational models capable of harnessing the potential of exascale computing could provide valuable insights for future integration with quantum systems.

3. **Artificial Intelligence (AI) Integration:** AI and machine learning methods are increasingly being integrated into scientific computing for tasks such as data analysis, prediction, and even computational model design. The Exa-MA project's work on AI algorithms (WP2, WP4, WP5, WP7) is directly contributing to the integration of AI and exascale computing, addressing one of the key future challenges in the field.

4. **Resilience and Error Handling:** As systems scale, the likelihood of errors increases. New methods for ensuring system resilience and efficient error handling will be needed. The work in this project on improving system resilience and robustness (WP3, WP7) will be directly relevant to future challenges in these areas.

5. **Complexity Management:** As models and simulations become more complex, efficient and effective ways to manage this complexity are required. This involves the development of more sophisticated modeling and solving techniques, optimization strategies, and uncertainty quantification methods. The Exa-MA project's focus on geometric and physics-based modeling (WP1), surrogate models and machine learning (WP2), numerical kernels and solvers (WP3), inverse problems and data assimilation (WP4), optimization (WP5), and uncertainty quantification (WP6) all contribute to managing this increasing complexity.

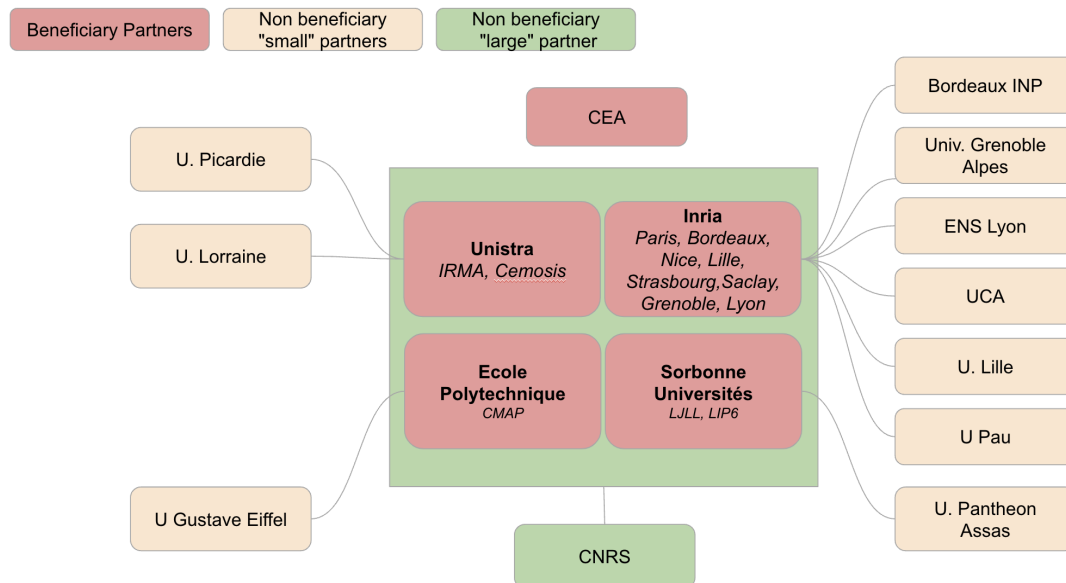
6. **Software Infrastructure:** With the growth in complexity and scale, there is a critical need for robust, efficient, and user-friendly software infrastructure to support exascale computing. The Exa-MA project, through WP7, aims to contribute to this by developing a strong framework for non-regression benchmarking tests, continuous integration and continuous delivery (CI/CD), and containerization. These activities ensure the reliability and robustness of the software tools developed, facilitating their integration into the broader European exascale software stack.

Furthermore, WP7 aims to contribute to Software Development Kits (SDKs) and demonstrate the project's results via a showroom. It will also provide training materials on exascale toolboxes, mini-apps and proxy-apps to enable wider use of the project's outputs. These activities will ensure the project's outputs are accessible and useful to the broader scientific and technical community, driving the adoption and impact of exascale computing solutions.

In this way, the Exa-MA project is not only addressing immediate challenges in exascale computing but is also strategically positioning its activities to meet future challenges in the field, paving the way for advancements in zettascale computing and beyond.

v. [Involved research teams](#)

The figure below displays the consortium that will develop Exa-MA.



vi. [Budget](#)

The requested budget for Exa-MA is 6.255 M€ which funds 91 Person Year.
The Total budget for Exa-MA is 24,417 M€.

2. [Exa-Soft - HPC softwares and tools](#)

i. [Context and challenges](#)

The year 2022 marks a milestone in the history of high performance computing as the exaflop performance target was finally hit with the installation of the Frontier supercomputer at the Oak Ridge Leadership Computing Facility, USA. Reaching such an extreme performance is the result of a hardware trend that was established by the end of Dennard's scaling and the slow down of Moore's law that increasingly relies on heterogeneity (i.e., the use of accelerators such as Graphical Processing Units, or GPUs) and unprecedented levels of parallelism. As a consequence of this increasing complexity, exploiting supercomputers poses daunting scalability and heterogeneity challenges to the HPC practitioners. The NumPEX PEPR project aims at addressing these challenges at multiple levels of the HPC functional stack through five Focus Projects (FP). FP2, entitled "Exa-Soft: HPC software and tools" is concerned with the software and programming environment.

Supercomputer architectures are becoming more and more heterogeneous as a consequence of the increasing use of specialization to improve performance under stringent energy consumption constraints: hardware components are optimized for specific workloads or operations. This obviously affects processing units; one example being the extensive use of GPUs, but other type of accelerators exist such as tensor units or low-precision units. This trend towards heterogeneity also affects the memory system (with the use of novel technologies such as high bandwidth or non-volatile memories) and interconnects (having different bandwidths and latencies for node-to-node, CPU-to-CPU, CPU-to-GPU etc. communications). At the same time, applications are also becoming increasingly complex and heterogeneous as they mix different workflows, programming paradigms, precisions, etc. In such a context, developing numerical libraries and applications with high efficiency and performance and functional portability is a challenging task which must not be discarded as a simple software engineering problem. Rather, this issue must be addressed with disruptive innovations from research. This is, precisely, the objective of the Exa-Soft project.

Multiple efforts in the academic and industrial HPC community aim at tackling this challenge. One notable example is the American Exascale Computing Project (ECP) whose “Software Technologies” focus area is closely related to this FP. In Europe, these topics are being addressed by multiple projects funded by the EuroHPC Joint Union whose activity is guided by the ETP4HPC Strategic Research Agenda which places the “programming environment” as one of the key challenges towards the production of scalable and portable exascale applications. Despite all these efforts, most of the software which is installed and run on supercomputers still relies on parallel programming models and tools which have been mostly designed for homogeneous architectures. Instead of producing incremental improvements of these models and tools, Exa-Soft aims at providing disruptive innovations in the programming environment to support the development of a novel generation high performance computing software that are scalable, portable and composable and where the expression of algorithms is not bound to low-level architectural details. To do so, Exa-Soft relies on a holistic approach that targets multiple levels of the software stack.

ii. [Project workplan](#)

To tackle the aforementioned challenges, Exa-Soft will be structured in six strongly interconnected work-packages. A heading WP0 work-package will be devoted to project management and coordination with other NumPEX projects, and will most notably synchronize all software developments with the FP5 “Development and integration project - ExaDI” project.

Cross-cutting activities will be established between all the work-packages to favor collaborations and exchanges and to ensure the global coherence of the FP activity and, ultimately, of the resulting software stack.

WP1: High-level approaches for developing efficient and composable parallel software. HPC systems have experienced significant growth over the past years, with modern machines having hundreds of thousands of computing units. These powerful machines come with several challenges for the HPC community.

The first one consists in efficiently using all the processing units with a clear and straightforward programming language. This language must express enough parallelism to deal with heterogeneous architectures. Computer scientists, applied mathematicians, and physicists must be able to use this language.

The second challenge is to be able to debug these applications on complex architectures. For that, we need dedicated tools adapted to both programming languages and architectures that allow easy validation of numerical methods at scale.

A third challenge is to ease code composability to improve separation of concerns and code reuse. To obtain composable parallel codes, two main questions have to be answered: how to compose codes and how to implement a composable code in a programming language. Therefore, this workpackage aims at providing models and software to define and implement composable building blocks. Task 1.3 focuses on defining composable elements while Task 1.4 deals with the data description and partitioning manipulated by those elements and by runtimes. This WP includes the following tasks:

Task 1.1 C++ complexity disambiguation for advanced optimizing and parallelizing code transformations

Task 1.2 Tools for parallel heterogeneous scientific application at scale

Task 1.3 Foundation of an HPC Composition Model

Task 1.4 High level data description and partitioning for reusable parallel building blocks

WP2: Just-in-Time code optimization with continuous feedback loop. It is now generally accepted that the scheduling and assignment of parallel tasks must be managed by a runtime system, which assigns tasks to computing units following the dependencies among tasks, as soon as these latter are resolved. The system usually also handles the heterogeneity of the nodes, by managing several versions of tasks adapted to a given hardware (*e.g.* CPU and GPU versions).

However, many decisions are usually made statically, *i.e.* before launching the application. Our goal is to make the runtime system take advantage of more advanced mechanisms devoted to improve its reactivity relative to the current runtime context, by implementing just-in-time task optimizations in three main directions that are described below. The general objective is obviously to make some big steps in runtime performance, but also in efficiency regarding energy consumption, with an emphasis on code optimizing transformations that require dedicated additional resources, only if they are largely compensated by the obtained gains. This WP includes the following tasks:

- Task 2.1 Runtime multi-versioning of parallel tasks
- Task 2.2 Resource-aware task generation
- Task 2.3 Specialization-based Dynamic Parallelization of Sparse Codes
- Task 2.4 Integration and unification of the runtime mechanisms

WP3: Runtime Systems at Exascale. HPC software stacks are more and more relying on runtime layers to achieve advanced optimizations during the execution of applications on HPC platforms. This indeed allows to hide the complexity of the hardware, and seamlessly leverage the various resources of the platform: network, GPUs, memory, disk, ... The coming exascale architectures however raise unprecedented concerns to these runtime systems, which have to automatically adapt to newer types of resources and sustain extreme scalability at all levels, and cope with hardware faults.

We thus propose a software stack that exposes a task-based programming paradigm and aims at executing at exascale levels by using fully asynchronous execution. It will benefit from high-level expression of parallelism and information from the compiler and analysis tools, to optimize the execution of the targeted application cases. This WP includes the following tasks:

- Task 3.1 Integration of asynchronous network communications scheduling and local task scheduling
- Task 3.2 High-level data description and partitioning mechanisms
- Task 3.3 Data placement in heterogeneous memory levels
- Task 3.4 Fault Tolerance for large-scale systems

WP4: Portable, scalable numerical building blocks and software. High-Performance Computing Applications targeting exascale systems require many building blocks to be able to fully exploit the platform. Among these building blocks, linear and multilinear (tensor) algebra libraries represent the core (and often the most costly) component for such applications. In this WP we aim at designing high-performance linear/multilinear algebra libraries for exascale systems. These libraries will strongly rely on the task-based programming paradigm (in collaboration with WP3) to be able to fully exploit exascale parallelism and the heterogeneity of the platform thanks to asynchronous execution models, efficient data management and scalable communication mechanisms. To be able to efficiently run on such complex computing nodes and to achieve better scalability at extreme scale, our libraries will also rely on approximate computing algorithms. Finally, from the point of view of the usability of such libraries, it is important to consider high-level programming interfaces and think about how to expose high-level programming models allowing for numerical composability and re-usability of such codes. The outcome of this WP will be the design of linear algebra and tensor software stacks on top of modern task-based runtime systems going from basic building blocks (e.g., dense linear algebra kernels) to high level sparse solvers like Maphys++. This WP includes the following tasks:

- Task 4.1 Composability of numerical libraries
- Task 4.2 Scalable dense and sparse linear algebra using task-based programming models
- Task 4.3 Efficient implementation of approximate computing algorithms
- Task 4.4 Sparse and dense tensor computations using task-based algorithms
- Task 4.5 Extension of Chameleon to small dimension tensors for large distributed systems with applications to deep neural networks

WP5: Performance analysis and prediction. Running traditional performance analysis tools on a parallel application either collects too little information for detecting problems, or too much information, which significantly degrades the performance of the program being analyzed. Moreover, analyzing the generated traces that consist of GBs of data requires considerable computing power.

We propose to design a new tracing tool for exascale applications and platforms that would be able to log data from the whole software stack without altering the performance of the application or generating terabytes of data. The tool suite will rely on a hierarchical trace format that groups sequences of events in order to limit the trace size, and to allow fast processing of millions of events.

This tool suite consist of several pieces of software: a tracing tool (designed in Task 5.1), and a post-mortem analysis tool (designed in Task 5.2). Tasks 5.3 and 5.4 (that are shared with WP 6 and 3) will extend these tools in order to make them collaborate with the other components of the Exa-SoFT software stack. This WP includes the following tasks:

- Task 5.1 Scalable tracing tool

Task 5.2 System-wide post-mortem trace analysis

Task 5.3 Fine-grained energy measurements

Task 5.4 On-the-fly performance analysis that guides the runtime system

WP6: Energy profiling and control. It is now a common knowledge that the power consumption of supercomputers is and will be a major concern. As a matter of fact, Frontier, the fastest super computer in the world consumes around 20 MW. As a consequence, finding solutions to reduce the power consumption of the architectures is mandatory. In this work package, we aim at reducing the power consumption of the target architecture when running the applications. For that purpose, we target two different levels: the node level and the cluster level.

On the node level, we aim at combining computation scheduling with fine control of the processor leverages (like the core and the uncore frequency). Understanding and using power measurements is not an easy task. It starts with accurate measurements : On the one hand, one cannot measure power consumption as often as other counters (because power meters are updated less often). On the other hand, to provide accurate feedback, measurements cannot be done at the application level, but at a lower granularity. As a consequence, we have to develop a fine-grain technique to measure and report the power consumption or decide the granularity of the measurements depending on the application. This work will be done in conjunction with WP-5.

Based on that, a feedback can be provided, either to the user, so he/she knows how much power and energy is consumed by the application, or to the runtime system to allow for better scheduling. The latter will be done in conjunction with WP-3. The runtime can use several leverages to change the power consumption (DVFS, application reconfiguration...). Depending on the task profile, these reconfigurations can impact power consumption while having nearly no impact on performance. The runtime system needs these higher-level models to decide the best mode for each application and hardware. Tasks 1 to 3 are those involved in the node-level optimization.

On the cluster level, the runtime can adapt the scheduling to the environment (scheduling only low power consuming tasks when electricity is scarce) but also allocate resources in an efficient way. Tasks 4 and 5 are at the cluster level. This WP includes the following tasks:

Task 6.1 Fine-grained energy measurements and models

Task 6.2 Energy-aware scheduling algorithms

Task 6.3 Cluster-level power measurements

Task 6.4 Energy-aware job scheduling and feedback

iii. [Software ecosystem and contributions](#)

The research and development activity of the Exa-Soft project is expected to achieve substantial advances at multiple levels of the software stack of exascale supercomputers. Although functional and performance portability are one of the main objectives of this FP, special attention will be devoted to achieving high performance and scalability on the possible future French exascale system. At the bottom of the software stack, contributions are expected to modern parallel programming models and languages and JIT code optimization with integrations into C++ compilers such as CLANG/LLVM, the COMET component-based model and framework, the APOLLO polyhedral analysis tool. Task-based parallel programming models and runtime systems will play a central role in Exa-Soft. Substantial developments will be integrated in the StarPU and XKaapi runtime systems targeting performance optimization, portability and novel features to achieve high efficiency and scalability on large scale supercomputers. Numerical libraries for dense and sparse linear and multilinear algebra will be improved or newly developed based on the advances made on the lower levels of the software stack. These include CHAMELEON, Celeste, Maphys++/COMPOSE, HPDDM and potential contributions can be provided to PasTiX, qr_mumps, MUMPS, PETSc. Finally, tools for performance and energy consumption analysis such as EZtrace, Mojito/S, DUF will be improved to efficiently provide useful feedback to high performance parallel libraries developers.

iv. [Post-PEPR scientific prospective](#)

The main challenges that Exa-Soft plans to tackle are related to the ever evolving landscape of supercomputing infrastructures. Because the architectures of supercomputers will continue to evolve past the end of NumPEX,

most of the questions addressed by this project will still be relevant for many years. Furthermore, some evolutions of the supercomputing industry may be so disruptive (e.g., massive use of AI methods, Quantum Computing) that fundamental changes will be made in the way we approach high performance computing.

Exa-Soft aims at proposing solutions whose use and applicability extend beyond the duration of the NumPEX project. In particular, it aims at producing programming models and tools that allow for expressing algorithms in a way that is as much as possible agnostic of the underlying architectural features. This will enable great performance and portability on current and future computing platforms, provided that the models and tools are updated and extended to support novel architectural features. This actually represents the main hurdle towards the long-lasting and widespread adoption of the solutions produced by Exa-Soft. Standardization of the proposed parallel programming languages and models can be seen as a possible mean to ensure the sustainability of the developed approaches.

v. [Involved research teams](#)

Project partners and teams:

Inria: Avalon, Diverse, Storm, Polaris, Topal, Concace, Camus, Cash, TADAAM, ROMA

CNRS: IRIT-APO, IRIT-SEPIA, LAPP, LIP6-PEQUAN

UPSaclay: PARSYS

CEA

Telecom SudParis: ACMES

vi. [Budget](#)

The overall budget allocated to the Exa-Soft project is 6.225 M€. Most of this will be used for the recruitment of research (PhDs and post-docs) and engineering personnel. The distribution of the budget among WPs is roughly uniform (WPs 5 and 6 are slightly smaller than the others) although, as explained above, a considerable part of the budget was used to support cross-WP activities.

3. *Exa-DOST - Data-oriented Software and Tools for the Exascale*

i. [Context and challenges](#)

The emergence of Exascale supercomputers poses significant challenges in terms of data management and utilization. One major challenge is the widening gap between the exponential increase in compute performance and the slower rate of improvement in storage bandwidth. This discrepancy threatens to result in the loss of valuable information generated by computational codes. Additionally, the integration of simulations into workflows that involve data from sources external to the supercomputer introduces new types of workloads and data access patterns, which may render existing software assumptions invalid.

On the other hand, new storage technologies such as CXL NVMe disks are being developed, leading to a more complex storage hierarchy. This complexity requires software libraries and middleware to adapt and provide applications with an abstracted and portable approach to data input and output. Furthermore, to efficiently store simulation results, it is necessary to increase the amount of information per byte before moving it to storage, necessitating in situ transformations and analytics that extract relevant information as data is generated. However, the adoption of this approach has been hindered by development and flexibility issues.

Shifting from post-hoc to in situ analytics reduces user involvement and offers performance advantages, but it requires the design of smart analytics that do not rely on interactive user input. The advent of machine learning presents an opportunity to achieve this objective, but it requires adaptation to handle HPC data in real-time and optimization for exascale execution patterns.

The Exa-DoST project aims to tackle these challenges by focusing on three key areas: 1) **scalable storage and I/O**, 2) **scalable in situ processing**, and 3) **scalable smart analytics**. This project builds upon previous initiatives, such as the Exascale Computing Project in the US, by targeting a higher technology readiness level in the HPC software

stack. The project will develop operational solutions co-designed and validated with French and European applications, filling the gap left by international projects and ensuring that the specific needs of France and Europe are considered in the **development of a data-oriented Exascale software stack aimed to be deployed on the future Jules Vernes Exascale supercomputer expected to start operating in France in 2025.**

ii. [Project workplan](#)

The Exa-DoST project specifically addresses three major data-related challenges that must be targeted in order to achieve Exascale performance:

Enable storage and I/O at Exascale. Tomorrow's application workflows generated by projects like SKA¹ will generate unprecedented volumes of geographically distributed data that need to be stored, transferred and processed on a potentially distributed infrastructure (HPC/cloud/edge). This requires innovative models and architectures for I/O, storage and processing.

Enable in situ visualization, processing and analytics at Exascale. Whereas ensuring the ability to perform such operations at the next scale is a natural challenge, the increasingly complex workflows expected to run on Exascale machines equipped with accelerators require innovative features including: on-demand in situ analytics, accelerator-enabled analytics, automatic in situ anomaly detection; high-performance advanced visualization.

Build advanced data-oriented tools to support AI-based analytics and smart simulation sampling at Exascale. With the advances of AI tools to model complex data, there is an opportunity to build novel workflows to analyze the outputs of large simulation runs and empower them with advanced steering capabilities. These novel tools will require advances on the statistical modeling of very large scale data including: out-of-core and distributed ML-tools; online simulation-based inference; and rare-event detection in large statistical signals.

Exa-DoST is organized around four technical work packages (WPs). The first three WPs correspond to the 3 major challenges described above. The fourth WP is dedicated to transversal topics including shared software libraries and components and integrated cross-WP illustrators, and WP5 is dedicated to management, dissemination and training activities, as illustrated in Figure 2.

¹ Scaife AMM. 2020 Big telescope, big data: towards exascale with the Square Kilometer Array. *Phil. Trans. R. Soc. A* **378**: 20190060.

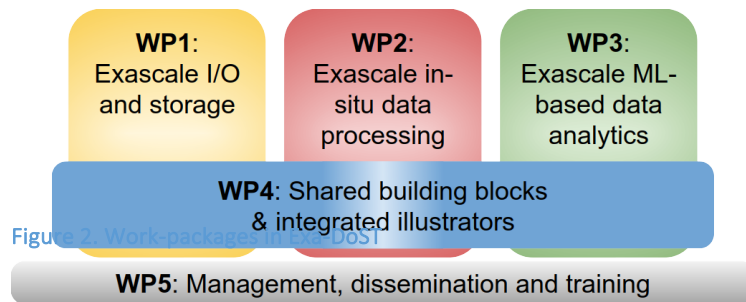


Figure 2. Work packages in ExaDoST

iii. [Software ecosystem and contributions](#)

In the process of building the Exa-DoST project we have built an inventory of existing relevant software projects developed by the project partners and of existing software use by the Exascale community in France, in Europe and worldwide. The picture below provides an overview of the planned software contributions of the participating teams which are planned to be used in the Exa-DoST. A detailed description is provided in the scientific annex of the Exa-DoST project.

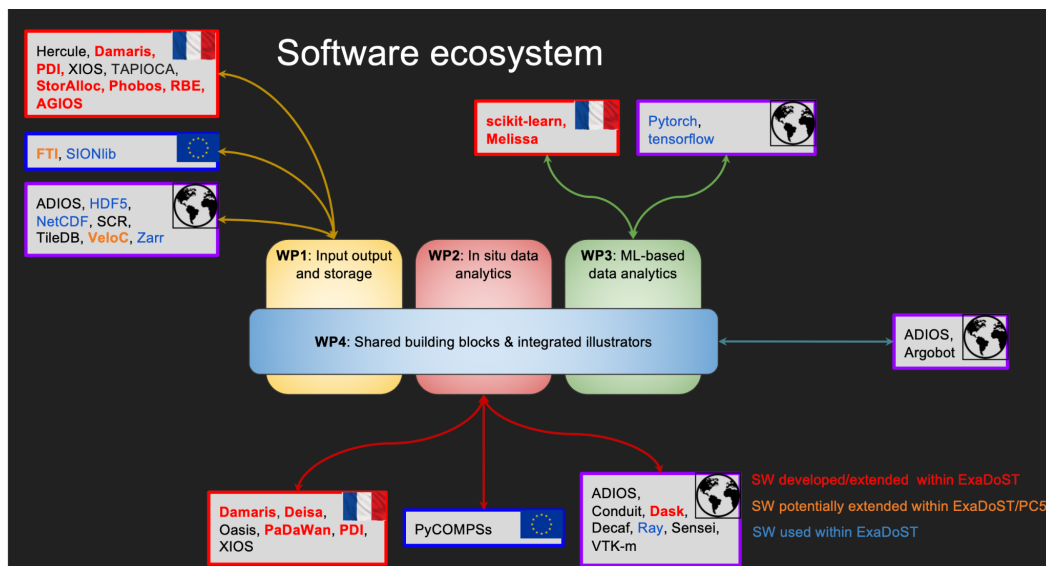


Figure 3. Software ecosystem. Red framed boxes correspond to French software; blue framed boxes correspond to European software; purple framed boxes correspond to international software. The software libraries that will be developed/extended within the framework of Exa-DoST are listed with red text.

iv. [Post-PEPR scientific prospective](#)

Exa-DoST will design, develop and distribute using an open-source model, a full, integrated software stack handling specific French and European needs in terms of data storage, IO, in situ processing and ML-based analytics. This software stack will be packaged and deployed on French and European supercomputers (in particular, a major target is the future Exascale machine expected to be hosted in France in 2025). This is an important milestone as we enter the post-Exascale era, where we expect most of the identified challenges to keep being valid from a high-level perspective, as it was the case in the past, when moving from Petascale to Exascale.

Ensuring scalable I/O and storage beyond the Exascale - as bottlenecks at I/O and storage level have always limited the scalability potential of increasingly complex HPC codes; Exa-Dost activities in WP1 are a first step in this direction.

Inventing disruptive mechanisms for real-time data processing at scale in increasingly complex infrastructure settings will be necessary, as the computing continuum (HPC-Cloud-edge) is becoming more and more relevant in support of an increasingly larger application spectrum; Exa-Dost activities in WP2 are a step towards this goal. We expect a strong, disruptive use of AI, that will impact the need for high-performance resources, thus the need for complex, integrated HPC-AI data analytics approaches. Exa-Dost activities in WP3 will contribute to progress in this direction.

v. Involved research teams

The consortium gathers **11 core research teams, 6 associate research teams and one industrial partner**. These teams represent 12 of the major French establishments involved in the field of data handling at Exascale. Below are listed the core teams:

DataMove (CNRS, Grenoble-INP, Inria, Université Grenoble Alpes) DataMove

DPTA (CEA) - Department of Theoretical and Applied Physics (DPTA) of the DIF,

IRFM (CEA) The Institute for Magnetic Fusion Research (CEA/IRFM)

JLL (CNRS, Observatoire de la Côte d'Azur, Université Côte d'Azur) - JL Lagrange Laboratory leading the French contribution to the SKA² global observatory project

KerData (Inria, INSA Rennes) KERDATA (<https://team.inria.fr/kerdata/>)

LESIA (Observatoire de Paris, CNRS, Sorbonne Université, Université Paris Cité) - Laboratoire d'Études Spatiales en Astrophysique

LAB (Université de Bordeaux, CNRS) - Laboratoire d'astrophysique de Bordeaux

MdIS (Université Paris-Saclay, UVSQ, CNRS, CEA) - Maison de la Simulation (<https://mdls.fr/>)

MIND (Inria, CEA) (<https://team.inria.fr/mind/>)

SANL (CEA) - a team from the computer science department (DSSI) of CEA

SISR (CEA) - a team at CEA is in charge of every data management activity inside the massive HPC centers hosted and managed by CEA/DIF.

TADaM (Inria, CNRS, Université de Bordeaux, Bordeaux INP) TADaM (<https://team.inria.fr/tadaam/>)

In addition, the following associated teams will also directly participate in the project, without receiving any direct funding: **CMAF** (CNRS, Ecole Polytechnique), **IRFU (CEA)**, **M2P2** laboratory (CNRS, Université Aix-Marseille), **Soda** (Inria), **Stratify** (Inria, CNRS, Grenoble-INP, Université Grenoble-Alpes), **Thoth** (Inria, CNRS, Grenoble-INP, Université Grenoble-Alpes) and the **DataDirect Network (DDN)** company.

vi. Budget

The requested budget for Exa-DoST is 6 125 000€.

The Total budget for Exa-DoST is 14 249 940,45 €.

4. *Exa-ATOW - Architectures and Tools for Large-Scale Workflows*

i. Context and challenges

The deployment of exascale supercomputers is a response to the growing demand for computing capacity for numerical models. In parallel with this growth, the volume of data produced by large scientific instruments, the IoT, and numerical simulations (e.g. CMIP 6) is also increasing exponentially. In this context, the design of scientific applications requires consideration of distributed processing chains between data centers, processing/calculation infrastructures, and scientists' laboratories (through the implementation of large-scale workflows, for example). This issue is illustrated in Figure 5.

² <https://www.skao.int/en>

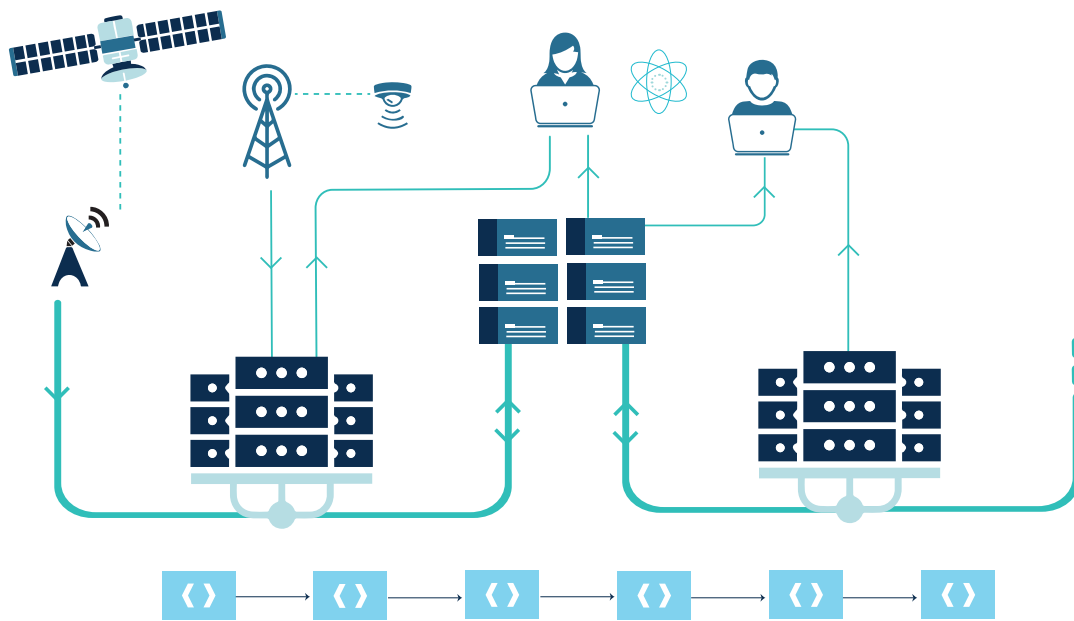


Figure 4. Exa-AToW aims to integrate data centers, cloud infrastructures and supercomputers into an overall vision to facilitate the development and deployment of scientific applications based on complex processing chains.

These infrastructures are not generally under the control of the same administrative authority, and are heterogeneous in their technical deployment. This raises the question of how to implement data logistics and orchestrate processing/calculation on this scale.

Data logistics must meet a number of constraints, including

- 1) Respect infrastructure storage constraints.
- 2) Take into account the capacities (as well as topology, latency and transfer protocols) of communications networks and computing infrastructures.
- 3) Ensure overall security.
- 4) Minimize energy expenditure.
- 5) Simplify application writing, deployment and debugging.

As we can see, the implementation of these logistics is part of a multi-criteria optimization problem in a context of collaborative systems-of-systems. Among the fundamental points addressed by these logistics are the traceability of operations and the FAIR policy, which must be applied to all scientific data. Taking into account metadata - associated with data or processes in the analysis of constraints and logistical solutions - appears to be an essential lever.

ii. [Project workplan](#)

The Exa-AToW project tackles the problem of large-scale workflows along 6 axes:

1) Design of a system architecture for collaborative systems.

This involves defining the set of services and their articulation to enable the management of data, metadata and the deployment of workflows.

2) Development of cybersecurity schemes for large-scale workflow deployment.

The federation of resources required to implement processing chains must comply with the site security policy. This WP aims to define the strategies and tools to be put in place for cybersecurity.

3) Design of data logistics and FAIR data lifecycle management solutions.

This work package aims to design data management schemas, metadata and specifications for the implementation of Machine Actionable Data Project Plans.

4) Creation of tools for resource federation management and workflow deployment.

This involves setting up tools for deploying the software needed to run workflows, storage and data flow management.

5) Design of energy management and environmental impact measurement models for processing chains.

The aim of this part of the project is to design models for estimating the energy consumption and environmental impact of applications, given a choice of implementation. The models should make it possible to compare deployment alternatives and thus reduce the energy footprint where possible.

6) Implementation of use cases involving data centers and computing centers

Exa-AToW will initially be based on 4 use cases from Data-Terra and PEPR Diadem. These use cases combine issues linked to data (e.g. produced by the "Sentinels" satellites) or data analysis and numerical models. They also require a combination of functionalities found in data centers and HPC centers.

In addition, to ensure the coherence of its work with the Jules Verne initiative and all Numpex projects, Exa-AToW is in charge of 2 cross-disciplinary workgroups:

- 1) The cross-disciplinary resilience axis (coordinator: L. Morin, University of Rennes)
- 2) Synergies with computing centers (coordinators: F. Bodin, Université de Rennes, N. Lardjane, TGCC).

The first cross-cutting theme will examine issues of resilience and fault tolerance across all NumPEx activities. The aim is to share and coordinate strategies and methods for implementing resilience. This topic will be addressed in collaboration with the Exa-DI project.

The "synergies with computing centers" working group aims to establish the operational elements, ideally shared, between computing centers and Numpex's targeted projects. Although this concerns all centers in the long term, the priority short-term objective is the future Exascale machine (2025). In particular, the group aims to build compromises or consensus on the use of Numpex infrastructures, software development processes and deployment technologies.

iii. [Post-PEPR scientific prospective](#)

By the end of the PEPR the Zettascale race will be at full speed. Emerging issues such as energy management, sustainability, data management, scalability, complexity management, and human resources to name a few will just be more acute challenges compared to the today's ones. Furthermore, the vision of the Zettascale as one big centralized system is unlikely to happen for eco-systemic and technical reasons. Indeed, the need for reinforced EU collaborations and open science will remain. But also more data sources for feeding scientific applications (for instance IoT, new large scientific instruments), more constraints on energy consumption, and a shifting HPC market toward cloud technology will very likely increase the underlying collaborative system-of-systems logic of the Zettascale infrastructure. Exa-AToW vision on large scale workflow is preparing for this new landscape. The tools we will be providing will help the community to develop the new generation of scientific applications capable of handling huge data, IA/HPDA, and numerical model in the distributed and federated infrastructure we will call the Zettascale.

iv. [Involved research teams](#)

Rennes I University, CEA, INRIA, CNRS

v. [Budget](#)

The requested budget for Exa-AToW is 5.05 M€.

The Total budget for Exa-AToW is 8.57 M€.

5. [Exa-DI - Development and integration project](#)

i. Context and challenges

Capable exascale computing ecosystems need to integrate exascale applications, software technologies and hardware innovations with an expanded software stack that can be easily deployed onto exascale facilities. The exascale Computational Science and Engineering (CSE) applications are themselves built on underlying software technologies that play an essential supporting role in application efficacy on a broad range of computing systems. They are a foundational element and an important vehicle for the delivery of consequential solutions and insight from exascale systems.

The Exa-DI project recognises that productivity, performance portability and sustainability of CSE applications require improvements in scientific software development methodologies, adopting new mathematical approaches, algorithmic or model improvements, and software developments leveraging optimised libraries and frameworks, extreme-scale programming and execution environments and data logistics, thereby improving exascale capability and thus ensuring new science and engineering insights.

However, tangible development progress requires close coordination among exascale application, algorithm and software development to address key CSE 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.

To this end, **the central objective of Exa-DI is to develop an iterative and incremental participatory process across NumPEX** — referred as a co-design and co-development process

— for science- driven improvement and new capability planning, by coordinating research and software developments (from PC1-PC4 in the NumPEX project) and representative, strategic exascale CSE application demonstrators (ADs) that are recognised as requiring a much higher level of inherent exascale effectiveness to achieve their challenge problem goals and impact. Every AD must be focused on targeted developments addressing a scientific and engineering challenge problem, i.e., one that possesses a solution amenable to computational insight and that is intractable without

exascale capabilities and/or capacities and strategic exascale ADs span a wide range of CSE domains and each addressing an exascale challenge problem — see Figure 6.

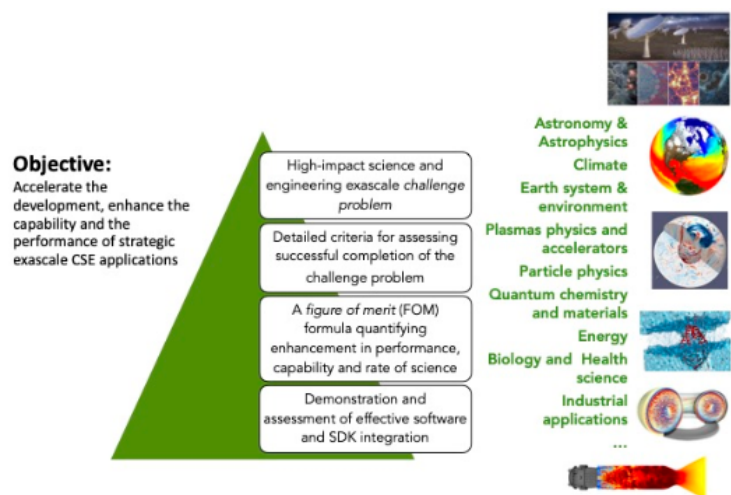


Figure 5. Exa-DI application demonstrators from a range of strategic domains address unique exascale challenge with identified capability and capacity issues.

The above Exa-DI co-design and co-development process will be enacted by recruiting and networking a Computational and Data science Team (CDT), composed of researchers, postdocs and engineers—see Figure 7. They will have complementary skills and expertise, at the interface between the CSE ADs and the NumPEX research and software development teams (PC1-PC4), that Exa-DI will allocate and manage across the co-design and co-development activities. The CDT's goal is to leverage common understanding and practices among the CSE ADs and the NumPEX research and software development teams, thereby enabling a push-and-pull relationship: pushing ADs' requirements into research and software developments; pulling these developments into the ADs for productivity and sustainability improvements.

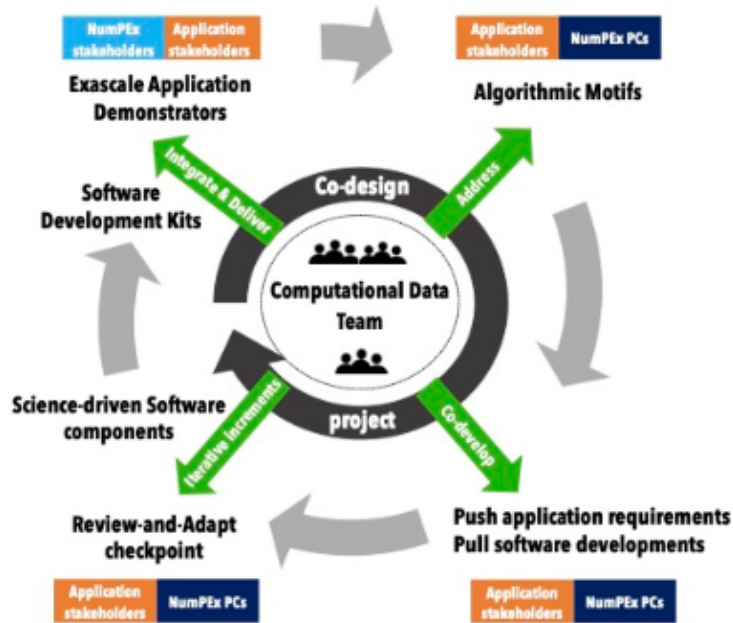


Figure 6. The CDT through a push-and-pull iterative relationship eases a co-development process between the CSE applications and the research and software development teams of the NumPEx PCs

The Exa-DI co-design strategy thus defines an important application-driven aggregation layer that will steer the coordination between initially loosely coupled research and software developments across NumPEx, while fostering longer term research. This research can then address disruptive shifts in extreme-scale computing, can advance science-driven software capabilities. Co-design will lower the complexity of application software by co-developing common efficient software building blocks or motifs, as well it will improve software development methodologies in the CSE community. This obviously mandates strong actions towards to CSE community, including communication, mediation and training.

Finally, Exa-DI strategy builds upon and is in synergy with several European and international initiatives, in which the project partners have previously participated. Among these, we mention BDEC (www.exascale.org), EXDCI (www.exdci.eu), EuroHPC Centres of Excellence (<https://www.hpccoe.eu/>), France Exascale Project. It is also in line and tap into other related Exascale projects with which the project partners have been and are collaborating, e.g., ECP project (<https://www.exascaleproject.org/>), E4S (<https://e4s-project.github.io/>) in the US, Fugaku Flagship 2020 project (<https://www.r-ccs.riken.jp/en/fugaku/fs2020/>) in Japan, and EUPEX flagship project (<https://eupex.eu/applications/>) in Europe.

ii. Project workplan

The Exa-DI project is organised around four work packages (see Figures 6, 8):

1. **Work Package 1 (WP1):** Select, steer and manage the co-design activities across Exa-DI.
2. **Work Package 2 (WP2):** Define, plan, execute and track the Agile Release Trains addressing co-designed algorithmic computational and communication motifs.
3. **Work Package 3 (WP3):** Define, plan, execute and track the Agile Release Trains addressing software integration and delivery as SDKs.
4. **Work Package 4 (WP4):** Engage in training and outreach activities.

Exa-DI is an aggressive project that will develop and implement a synergistic, iterative and incremental co-design and co-development strategy across NumPEx taking steps toward an augmented, integrated science-oriented exascale software stack driven by the needs to accelerate the development, the productivity and the sustainability of strategic exascale ADs. Each strategic Exa-DI AD is focused on targeted development to address a unique exascale challenge problem and represents a strategic problem important to the NumPEx stakeholders, and is currently intractable without the computational power of exascale systems. The strategic ADs are carefully co-analysed and co-prioritised with the different PCs (PC1-PC4) of NumPEx.

High-priority, co-designed algorithmic (computation, communication) motifs cross-cutting several ADs, and their most urgent development challenges, are co-identified and prioritised to be collaboratively addressed by co-design projects driven by the CDT. This applies as well to proxy-applications efforts whose mission is to develop proxy tools to explore algorithms, data structures/layouts, optimisations, etc., together with the associated trade-offs on different architectures. The goal of co-designed algorithmic motifs in Exa-DI is to integrate a rapidly developing software stack while developing software components that embody the most common patterns of application computation and communication.

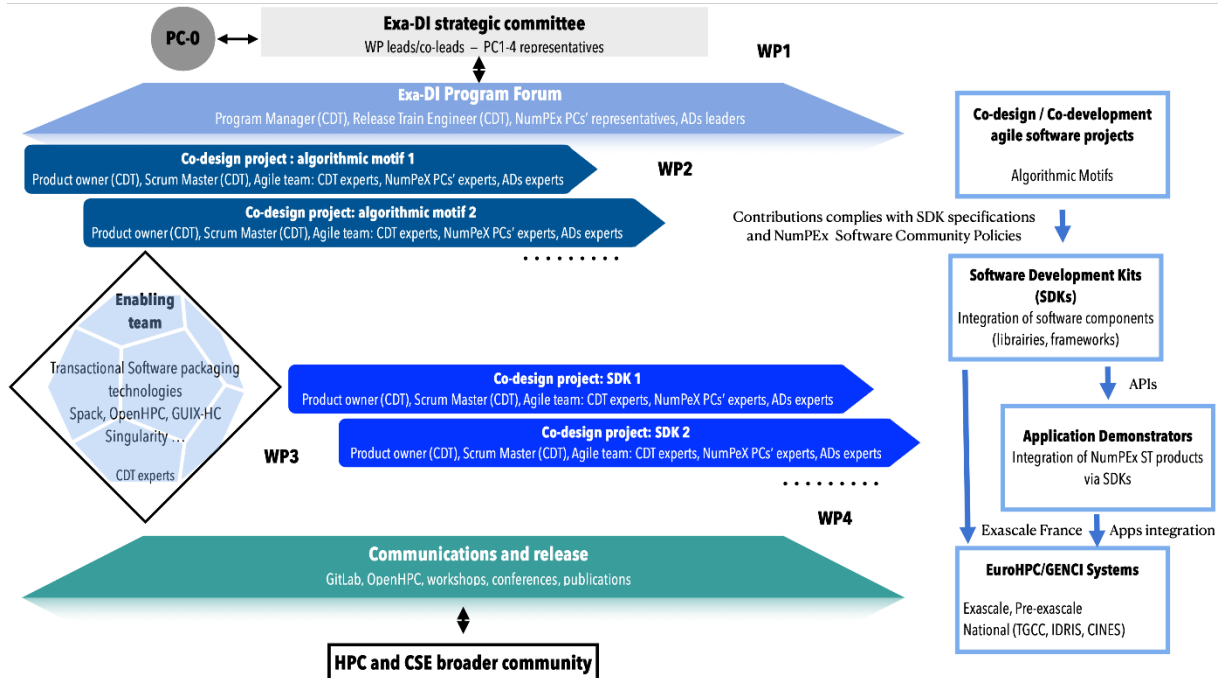


Figure 7. Exa-DI flow of product delivery and deployment

Science-driven logical collections of co-developed software components (libraries, frameworks and tools) are identified and prioritised for integration within the NumPeX software stack portfolio, possibly with external software components, by collaborative co-design projects driven by the CDT. They are delivered in the form of SDKs –see Figure 4, easily deployable on national computing facilities so that ADs can instantiate and compose them into their respective application software environments to accelerate their exascale development, improve productivity, sustainability and performance portability. This is supported by an enabling team focused on extending and hardening common NumPeX exascale software packaging and containerisation technologies, and by defining NumEx software community policies and providing a common software integration hub.

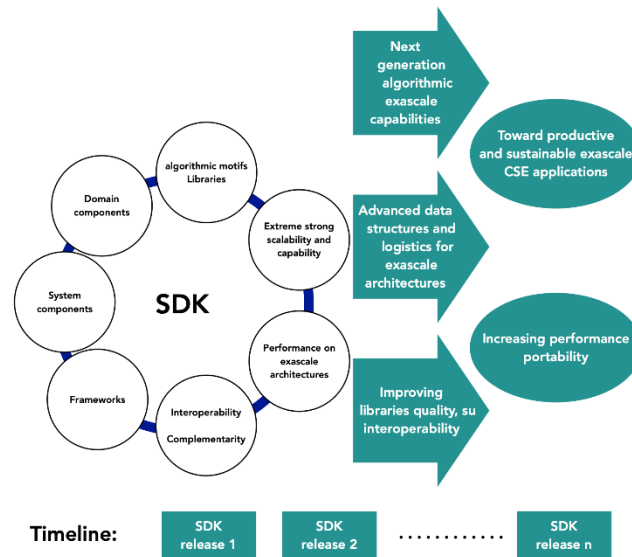


Figure 8. Software Development Kits (SDKs) are logical collections of co-developed science-driven software components and other components that are integrated, packaged and delivered to be easily deployed onto national computing facilities and instantiated by the CSE applications.

iii. [Co-design, a transversal project within NumPEX](#)

The co-design projects addressing cross-cutting algorithmic motifs and SDKs are developed by self-organising and self-managing Agile teams, driven and managed by the CDT, that aggregate cross-functional resources committed from the research and development (PC1-PC4) and ADs teams, together with additional resources from the CDT. The Exa-DI adopts a Lean and Agile development methodology in which co-design projects are executed as Agile Release Trains (ARTs). Aligning, steering and coordinating the projects during the ARTs and identifying cross-dependencies is under the responsibility of the Exa-DI program forum involving a program manager and a release train engineer from the CDT together with representatives from the NumPEX PCs (PC1-PC4) and ADs.

This Exa-DI co-design and co-development strategy must balance CSE application requirements with what is feasible in the NumPEX PCs' research and software development program.

The Exa-DI strategies thereby articulate four complementary and interconnected areas.

1. ***Curating science-driven software development methodologies to improve coordination across NumPEX:*** Providing information to improve scientific software development practices and processes targeting software productivity and sustainability — for better planning, development, interoperability, integration, performance portability, reliability, collaboration, and skills. The CDT promotes community collaboration across NumPEX and ensures that information is tailored to address the needs of exascale CSE applications.
2. ***Upgrading Incrementally and iteratively software practices:*** Providing a lightweight iterative development workflow driven by the Agile teams addressing well-identified ADs-driven co-designed cross-cutting algorithmic motifs and their most urgent software development bottlenecks. The CDT enables Agile teams to work collaboratively across NumPEX to overcome them. Through improvements in software co-development practices such as integrating (SDKs), testing, refactoring and on-boarding, the iterative process enables Agile teams to mitigate technical risk and advance.
3. ***Establishing software communities:*** Agile teams foster the growth of topical software communities whose members work together to leverage common understanding and practices to advance the quality, interoperability and performance portability of software technologies. Key strategies determine community software policies and create topical SDKs, enabling CSE ADs to more easily compose and instantiate complementary software products.
4. ***Engaging in community outreach:*** A central information hub sharing practices, processes, tools, and SDKs—see Figure 9— will emphasize high-quality software components and methodologies to improve software productivity and sustainability. Additional initiatives

include materials for tutorials and webinar series targeted toward exascale CSE software developers; organising technical and training meetings to nurture communities of practice.

This will enable:

- Improving current methodologies to deliver productive science-driven software stack at exascale where possible;
- Performing the research to conceive of new approaches when current approaches will not suffice;
- Integrating and delivering SDKs addressing the needs of the CSE applications and that can be deployed to the national computing facilities and instantiated to accelerate their exascale development.

iv - Involved research teams

CEA, CNRS, INRIA

v - Budget

Budget: 9,31 M€

Total budget: 13,58 M€

III – Transverse actions across the programme

Leveraging the challenges of the Exascale scientific and technological shift requires an interdisciplinary effort from fundamental research to software developers. The success of the NumPEX programme will largely depend on the ability to foster collaborations between these communities within and outside NumPEX. Each of these groups are composed of representatives of each of the focused projects PC1-5.

To that purpose, NumPEX programme will build upon internal transverse actions animated with dedicated working groups:

- **Software productivity working group:** the goal of this working group is to co-build a software community policy, define what a NumPEX software contribution is, and evaluate and control software production across the programme from the software development at the level of the scientific project to their final integration in the application and software integration project.
- **National computing facilities working group:** The NCF WG be in charge of co-build with the operational committee a roadmap describing the synergies with the NCF, facilitate the implementation of the programme roadmap with respect to the NCF, and coordinate communications and activities related to the NCF across the programme.
- **Equity and Gender balance:** It will aim at establishing the NumPEX code of conduct, proposing and performing precise actions to promote gender balance and equity within the programme, propose a roadmap of actions to prioritize gender equality and favor a gender-balanced recruitment of the personnel hired by the projects, promote the participation of women in research professional careers
- **Transverse technical initiatives working group:** Identify cross-projects technical and scientific topics of interest. These cross-programme topics will be analysed and re-evaluated all along the programme, organize technical and scientific actions to leverage collaborations. One of their goals will be to identify potential gaps in the scientific and technical program to prepare the call for proposals. Current active working include: accelerator programming, resilience, energy and AI.

NumPEX Training:

The NumPEX programme will impact the French HPC/HPDA/IA ecosystem at all levels, from the development of new tools, contributions to well-established softwares, their integration in applications to their deployment in present and future computing facilities. As it requires a significant shift in software use, development and integration, training actions are central in the NumPEX programme. Within NumPEX, training events will be

proposed and organized from actions targeting the HPC community to early training students. Interactions will be fostered with training actors in the French and international ecosystem, such as French training actors (PRACE training centers), engineering schools, University graduate schools. Training activities will be coordinated and organized with a dedicated working group, which is composed of University professors and expert in HPC training. In order to reach a large audience, training activities will take various forms, beyond standard PRACE-like events: participation to Master-level Exascale-oriented courses, hackathon and coding session, tutorials and a web training hub, etc.

National and international collaborations:

Within NumPEX we aim at fostering collaboration at the national level, with similar initiatives in AI or cloud technologies or application-focused PEPR programmes. It will also be in charge of organizing the national and international scientific events (e.g. workshops, sessions in conferences, etc.). Collaborations will target two main communities:

- at the national and European level, initiate collaborations with similar PEPR programmes such as PEPR cloud and PEPR IA, which are complementary to the NumPEX programme. Foster interactions and collaborations with PEPR programmes related to application domains, such as TRACCS, DIADEM or ORIGINS to name a few. At the European level, establish interactions with our European colleagues in projects including EuPEX, EPI, or the Center of Excellence (CoE). These collaborations will allow to answer to call for proposal from the European commission and EuroHPC JU.
- At the international level, we will (co)-organize of international workshops, as part of an international initiative with European, American and Japanese collaborators. Six workshops are scheduled that will alternately take place in the three major regions of the (Europe, US, Asia). During the duration of the programme, NumPEX will fully organize one 3 -days international in France, targeting an audience of about 100 attendees. These series of workshops, coined International Post-Exascale initiative (InPEX), will foster active international collaborations on keys topics for preparing the post-Exascale era: benchmarks, AI, hardware/software/application co-design, software production and sustainability to name a few.

Call for proposals:

Call for proposals will be open in 2024 and 2025 to fund cutting-edge scientific and engineering projects. The topics to be addressed will be discussed and established with representatives of all the focused project, and based on gap analysis produced by the transverse technical initiatives working groups. This process will allow to complement the initial scientific workplan for a better adaptation to the current and foreseen trends in HPC, HPDA and AI.

IV – Post-NumPEX prospective

By the end of the NumPEX programme, the HPC/HPDA/AI community will be will largely on its road to Zettascale. As detailed in section II, the scientific and technological challenges will likely lie in the increasingly more complex HPC architectures based on more heterogenous technologies, the overwhelming use of data-oriented and AI solutions at all levels of the HPC system, the development of hardly energy-constrained workflows, to only name a few. More generally, the post-NumPEX prospective has to be discussed with respect to the general objectives we detailed at the beginning of this document:

- **Preparing the building blocks for post-Exascale software solutions:** it is expected that future trends in HPC/HPDA/AI will follow the same path with highly heterogeneous architectures (e.g. accelerators, chiplets, FPGA, etc.) and data-intensive use involving advanced AI/ML tools. The initial objectives of NumPEX are twofold: i) a short-term goal for the production of Exascale-grade software components to contribute to a European Exascale software stack and ii) the exploration of new approaches for Exascale and post-Exascale systems (e.g. advanced numerical solvers, resilience, energy-efficient tools, advanced data analysis methods/software, hardware-agnostic programming models, etc.). Within NumPEX, transverse scientific actions will be launched to tackle challenges (e.g. AI, energy, resilience, accelerators, etc.) that can only by addressed through inter-disciplinary collaborations.

- **Preparing the academic and research applications to enter the Exascale era:** the Exascale and post-Exascale shift mandates a significant re-design of the academic and industrial application codes. A re-design of HPC codes generally requires years of development with large engineer and researcher teams, which is out of reach for most applications. The spirit of the co-design activity of the NumPEX program starts a new path for code development: i) mutualize the development of Exascale-grade software bricks that will be common to several applications, ii) deploy modern agile software development methodologies, which are essential to develop scalable and efficient scientific or industrial application code. In the long-term this will largely build upon an active Exascale community, gathering the French experts from fundamental mathematics to computer science and people from the applications domains. To that end, training will be central to build the community we need to cope with the future post-Exascale challenges.
- **Software sustainability and evolution.** The NumPEX programme is the *software* contribution the Jules Verne project that will yield the second European Exascale system in 2025. To that end, and as a *production* project, NumPEX will deliver software components for a European software stack to be used on the forthcoming HPC systems. To that regard, a central question relates to the sustainability of the NumPEX software production. Sustainability will build upon three pillars:
 - o The development of adaptable and efficient software components (e.g. scalable, resilient, energy-efficient, etc.) that can be maintained and adapted to post-Exascale architectures and usage. To that end, and this is the basis of the NumPEX program, we need to build collaborations between the different communities of the HPC/HPDA/AI ecosystem to build an integrated Exascale community. This will allow converging to common software components and a large-scale use of the NumPEX production.
 - o Sustainable softwares need to be maintained and evolve in the long-term, which means they are living softwares. The larger the users community the more sustainable such softwares will be. A key element for sustainability is to create and animate a large community of users for the software productions, which will be software engineers and more generally users from the application domains. To that purpose, the software/application co-design as well as the training actions will play a central role to enlarge the community of users. Furthermore, the NumPEX software production and integration activity will be important to deliver usable and easily deployable softwares.
 - o Software production sustainability is a common concern for all our international partners. To get prepared for future challenges, it is paramount to exchange with the international community, forge new collaborations to converge to a common view of the future evolution of the HPC software stack and their deployment in applications. Shared or common productions could include softwares as well as benchmarks. Forging national (e.g. with other PEPR on strategic topics such as the cloud or AI, or applications such as digital Health and climate for instance) and international collaborations is a central activity in the NumPEX program.
- **AI for Science.** The very fast development and deployment of Artificial Intelligence has a huge impact on the whole chain of scientific HPC, from the hardware and software developments as well as the applications needs. AI is expected to play a central role to increase the scalability of traditional HPC applications (e.g. multi-physics code coupling, surrogate models for sub-grid simulations, to only name two). As well, this requires a significant evolution of the software stack to better integrate and interface HPC components in AI-centric workflows and vice versa. International initiatives are starting, focusing on extreme-scale models for Science and engineering applications (e.g. [Trillion Parameter Consortium](#)), as well as European and French AI strategies. NumPEX is central for the HPC/AI ecosystem to bridge the gaps between the HPC and AI communities. NumPEX will push the development of SW components answering to the needs in AI for Science and engineering, in coordination with these national and international initiatives.

More information and news can be found at <https://numpex.org> or on [LinkedIn](#).