

Exa-DoST: Exascale Data-oriented Software and Tools



PROGRAMME
DE RECHERCHE
NUMÉRIQUE
POUR L'EXASCALE

PEPR NumPEx project ANR-23-PECL-0007
1 Jan 2023 – 31 Oct. 2029 (82 months)
PI: Gabriel Antoniu (Inria), Co-PI: Julien Bigot (CEA)

The Exa-DoST “core” team

CEA/DAM - DPTA, SISR, SANL

CEA/DRF - MdIS, IRFM, IRFU

CNRS/INSU, Observatoire de Paris, Observatoire de la Côte d’Azur

Inria - DataMove, KerData, MIND, TADaaM, +SODA, +STATIF, +THOTH

DDN

The NumPEx Program

Co-directors: Dr J. Bobin (CEA), Pr M. Krajecki (CNRS), Dr J-Y. Berthou (INRIA)

Project leaders and co-leaders:

Exa-Ma - C. Prudhomme, U. de Strasbourg– Hélène Barucq, Inria

Exa-Soft - R. Namyst, Inria/U. de Bordeaux - Alfredo Buttari, IRIT

Exa-Dost - G. Antoniu, INRIA - Julien Bigot, CEA

Exa-AtoW - F. Bodin, U. de Rennes - Mark Asch, U. Picardie - Thierry Deutsch, CEA

Exa-DI - J-P. Vilotte, DR CNRS - Valérie Brenner, CEA

The French NumPEX Program

Context and motivations




*A technological
breakthrough
Hybrid scalar/acc.
fewer memory/node
more concurrency
Mixed precision, ...*




*In the digital
continuum
Increased flux/volume from
the edge to the HPC system*

**Convergence
HPC/HPDA/IA**



Traditional
HPC
Systems



**Capable
Exascale system**


The French NumPEx Program Objectives

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

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

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


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



Help aggregate the French HPC/HPDA/IA community

The French NumPEX Program


Objectives




European Pre-Exascale system



Co-develop the exascale software stack
Preparing the applications for the Exascale era



NumPEX by numbers



6 Years
41 M€*

2023-2028

* Funding 41M€=500 person.year non permanent staff

+ 170 person.year permanent staff

Total cost : 81 M€

**Core
Research
Institutions**

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

**3
Focus
Area**

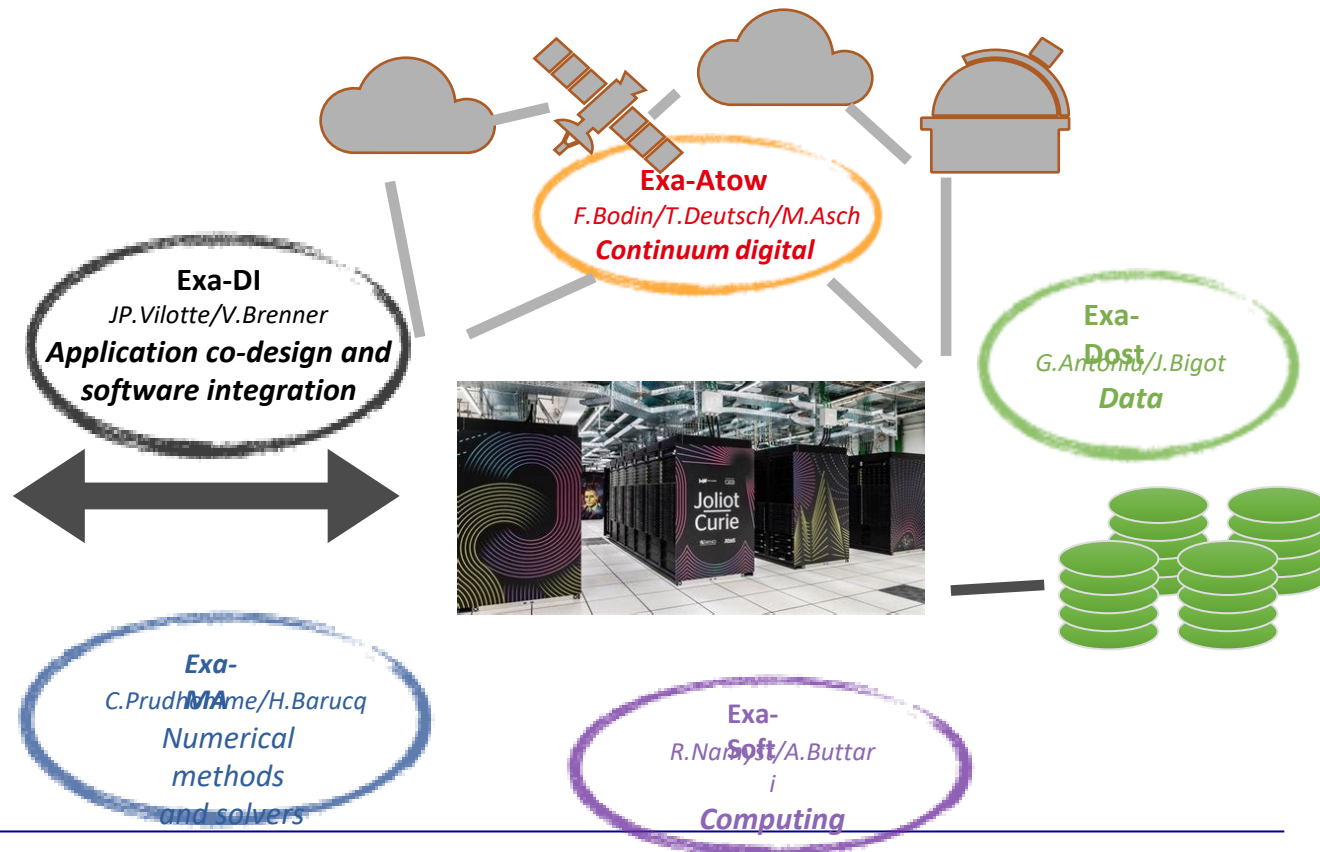
Software stack development (PC 1-3)
Wide-area workflows and architecture (PC 4)
Integration and application development (PC 5)

**80
R&D teams
500
Researchers**

The French NumPEX Program workplan



Applications



The French NumPEX Tranversal groups

**Accelerated
architectures and
programming models**

S.Thibault/M.Pérache

AI

T.Moreau/E.Franck/J.Bobin

Computing centers

F.Bodin/N.Lardjanne

**Energy management
and optimization**

A.Guermouche/G Da Costa

**Gender/Equity/Diversit
y**

A-L Pelé/V. Grandgirard

Training

M.Krajecki

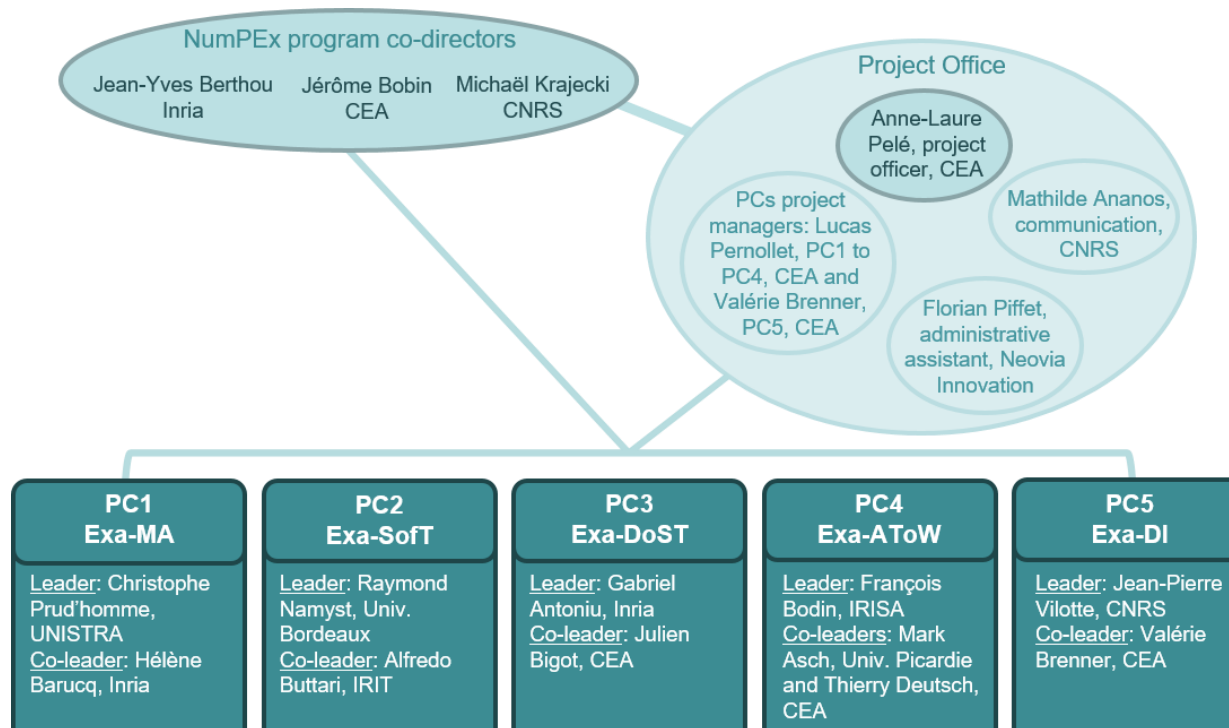
**Software production
and integration**

B.Raffin

**International
collaborations**

J-Y Berthou

Organisation



Create a network of young recruits in NumPEX (internships, theses, postdocs on fixed-term contracts, and permanent young employees).

Develop a cross-disciplinary community through PCs, thereby enriching exchanges and even fostering collaborations.

Propose original initiatives to the NumPEX management, such as:

- Actions during General Assemblies (AGs)
- Seminars
- Actions for the integration of young recruits into NumPEX
- Communication, training, careers
- ...

Don't miss the YoungPEX workshop !

The International Post-Exascale (InPEX) Project

InPEX expected outcomes

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

Actions:

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

Participants:

Researchers, engineers, industry, funding bodies



The International Post-Exascale (InPEX) Project Inpex.science

Date	(10/2023)	11/2023	06/2024	04/2025
Location	Preparatory phase EU (France)	SC'23 - BOF	Workshop1 EU/BSC	Workshop2 Japan
Date	03/2026	09/2026	06/2027	09/2027
Location	Workshop3 US	Workshop4 EU	Workshop5 Japan	Workshop6 US

Latest InPEX meeting, Kanagawa /Japan, 14-17 April 2025



Four parallel sessions:

- AI and HPC
 - Sharing AI-centric benchmarks of hybrid workflows
 - Generative AI for Science
- Software production and management
- Digital Continuum and Data management

Data at Exascale: Challenge at Hardware Level

- Increasing **gap between compute and I/O** performance on large-scale systems
 - Ratio of I/O to computing power divided by ~10 over the last 10 years on the top 3 supercomputers
- ... and data deluge!
 - At NERSC, **data volume x41** in 10 years
- New storage tiers and advanced architectures to try to mitigate this increasing bottleneck**
 - More complex on-node memory layout
 - Emerging complex applications and workflows have to adapt

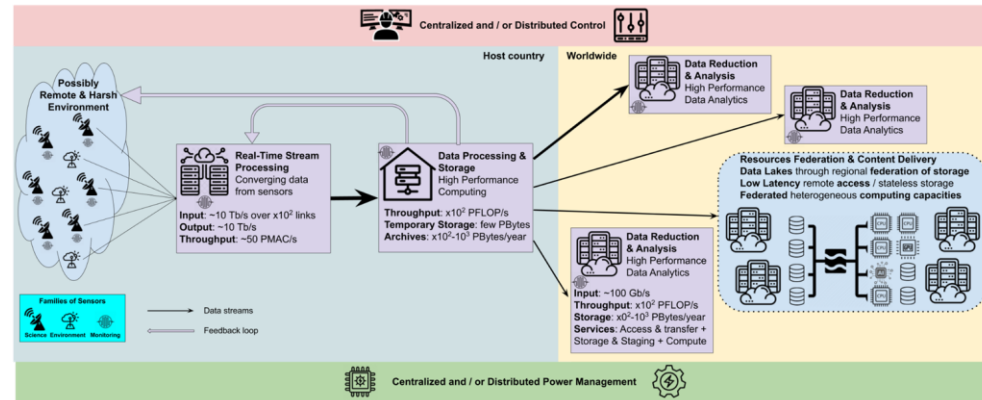
Ratio of I/O bandwidth (GBps) / TFlops of the top 3 of the Top500



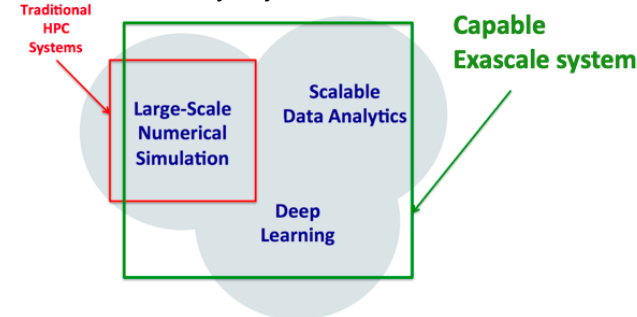
Trend in storage technologies available on extreme-scale systems

Data at Exascale: Challenge at Usage level

- HPC centers do not live in isolation anymore
 - Edge - Cloud - HPC continuum
- Emerging workloads are hybrid
 - High-performance simulation
 - High-performance data analytics
 - Machine learning and artificial intelligence
- Interaction with data from the outside world sensors
 - Large scientific instruments



SKA data workflow from sensors to HPC centers



Exa-DoST: Expected Outcomes

Approach:

- **Research** on data-oriented tools for HPC
- Transverse, **re-usable tools**
- Usable **in production** at exascale

Fill the gaps in the existing software stack designed by previous projects (e.g. ECP)

Take into account French & European specificities

Ensure French & European needs are taken into account in roadmaps

⇒ ExaDoST will produce:

- **New approaches** to handle the data challenge at exascale
- Transverse **libraries & tools** that implement these approaches

Fully application agnostic

Fully open-source

Validated in illustrators at full scale

Work Packages in Exa-DoST

WP1: Exascale
I/O and
storage

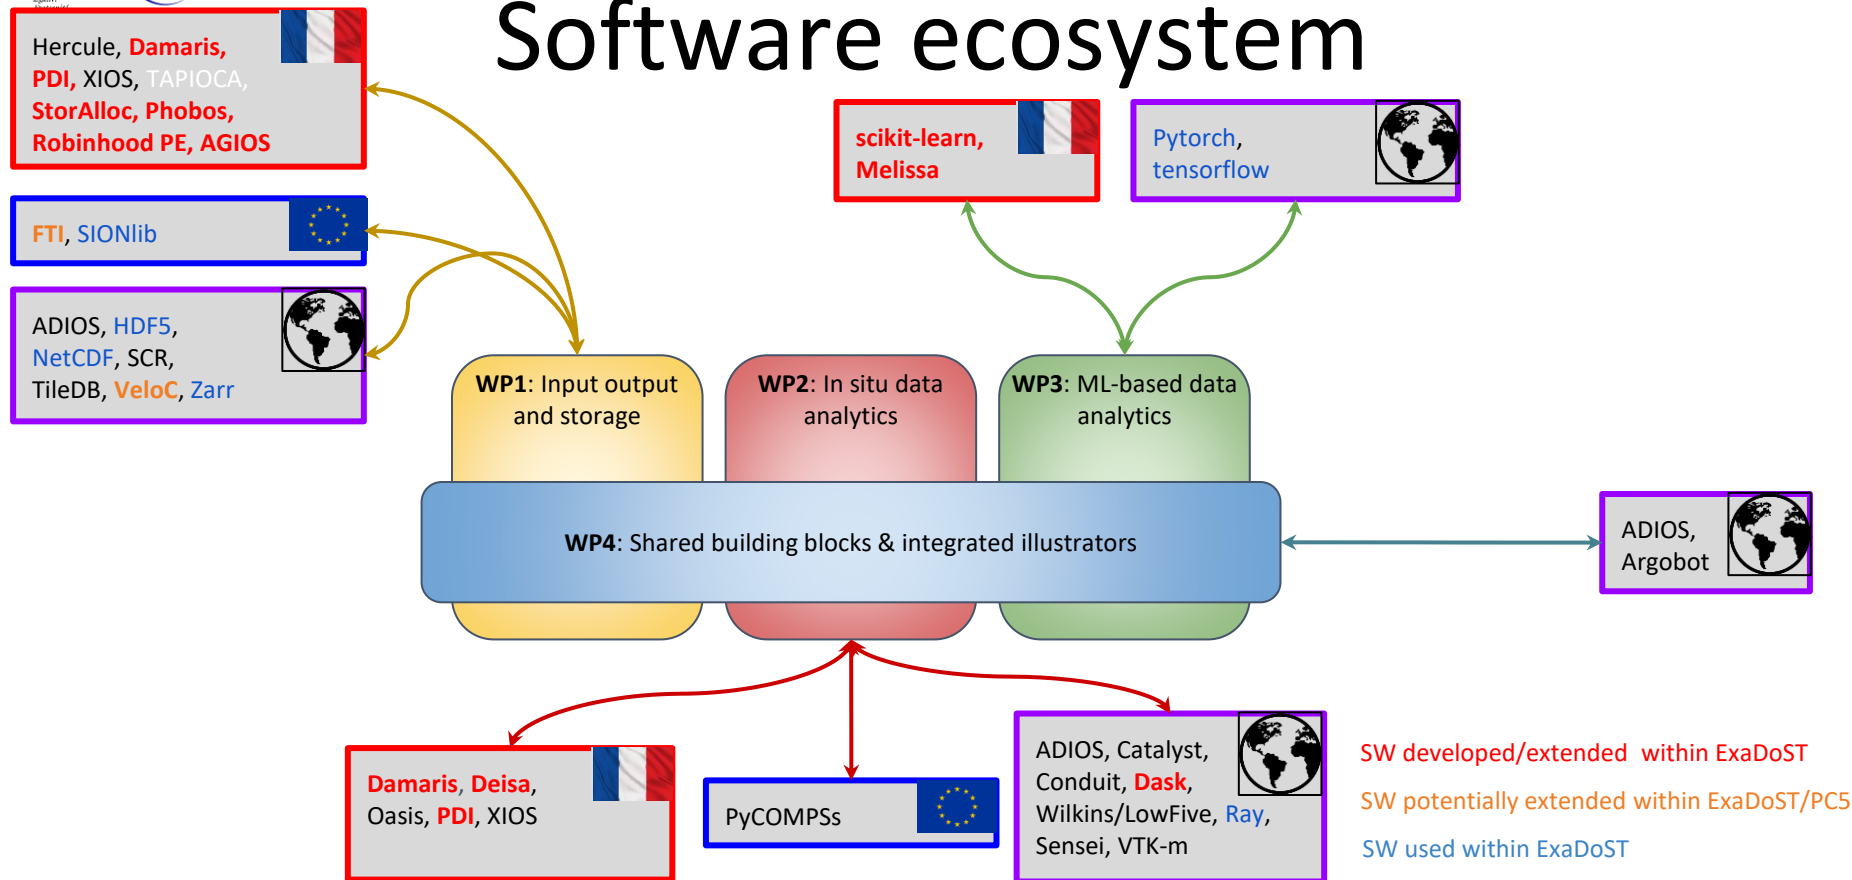
WP2: Exascale
In-situ data
processing

WP3: Exascale
ML-based data
analytics

WP4: Shared building blocks
& integrated illustrators

WP5: Management, dissemination and training

Software ecosystem





Illustrator 1: GYSELA towards exascale

Main challenge : optimized management of huge amount of data with in-Situ AI-based diagnostics



- GYSELA a non-linear 5D gyrokinetic code developed for 25 years at CEA/IRFM to simulate plasma turbulence in tokamaks.
- Optimized up to 730k CPU -> Intensive use of petascale resources (~150 millions of CPU h / year)



Relative efficiency of 85%
on more than 500k cores

Typical simulation:

- **100 billion points** (5D mesh: 3D space + 2D velocity)
- **~7 million of CPU hours** (3.5 days / 65k cores)

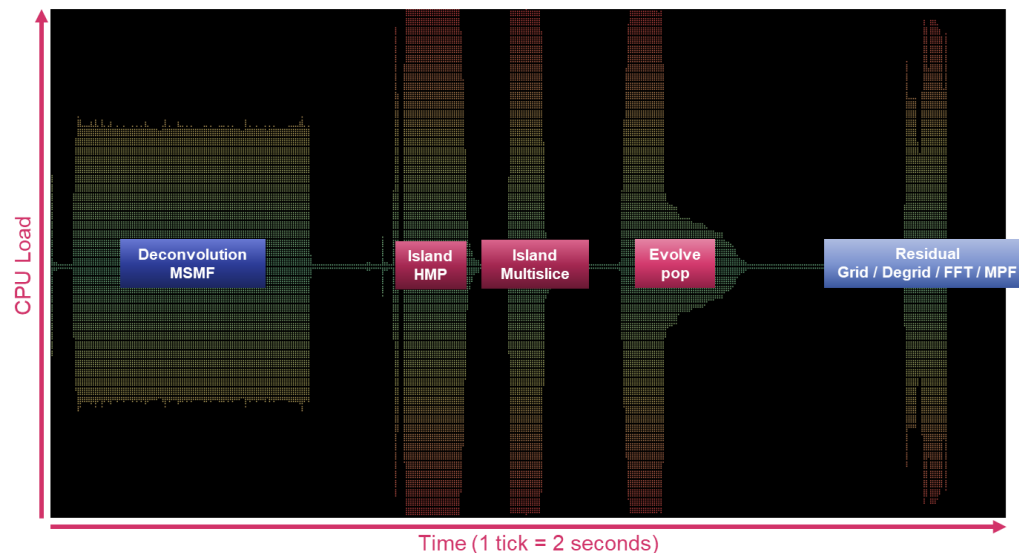
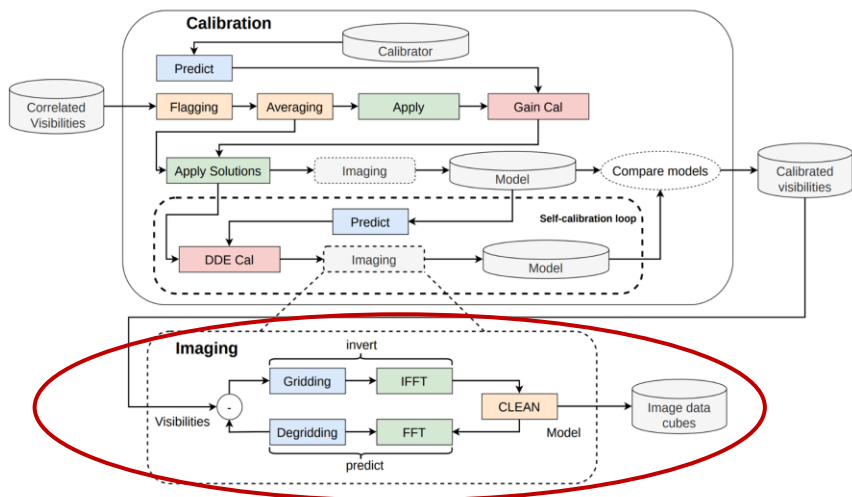
[V. Grandgirard et al., PASC 2022]

- **I/O scalability is an issue:** ~50% for 3072 nodes and ~38% for 4096 nodes. Crash on 5696 nodes
- **Need to be solved for exascale ITER-like simulations**

Illustrator 2: Square Kilometer Array (SKA)

Based on SotA software used to process large surveys (DDFacet)

- Complex iterative pipeline: optimize I/O footprint and upscale
- Ongoing work: definition of benchmarks and motifs identification



Illustrators overview


Integrated cross-WP illustrators

- Gysela (CEA/DRF/IRFM)
- SKA (CNRS + ...)

Motif-specific motivators

- Coddex (CEA/DAM)
- CROCO (CNRS + Inria + ...)
- Dyablo (CEA/IRFU)

Updated Gantt Chart



Hiring: Status

WP	PhD students	Engineers	Post-doc	Total
WP1	1	2	0	3
WP2	1	4	0	5
WP3	1	1	1	3
WP4	0	2	0	2
WP5	0	0	0	0
Total	3 (6 prévus)	9 (12 prévus)	1 (2 prévus)	13 (19 prévus)

Collaborations with Other Projects

Within NumPEx

- Exa-AtoW: collaboration within the ECLAT joint lab on SKA (WP1, WP4), InPEx (WP1-WP5)
- Exa-DI: deployment of tools (particularly from WP2) for SKA, GYSELA), joint meetings
- Exa-MA: joint ANR project with the team of Emmanuel Franck on turbulence models (WP4)

With other PEPR programs (France 2030)

- PEPR TRACCS: HPC/IA coupling (WP3)
- PEPR Cloud : scalable storage (WP1), SKA demonstrator (WP1, WP4)

At European level

- EoCoE-III Gysela-X++ exascale demonstrator (<https://www.eocoe.eu/>)



PROGRAMME
DE RECHERCHE

NUMÉRIQUE
POUR L'EXASCALE

Retrouvez toutes nos actualités

 NumPEX