MAISON DE LA SIMULATION

La Maison de la Simulation

A joint lab between

Maximize the scientific output of supercomputers.

Build multidisciplinary teams with all the skills to meet the great HPC challenges

Three axes for development :

- Multidisciplinary **research** lab around numerical simulation
- **Service** unit to help communities develop cutting edge numerical tools
- A **training** hub for HPC

Develop new numerical methods

- Méthodes stables à «grands pas de temps» pour les écoulements compressibles (Chalons, Girardin, Kokh)
  - Propriété AP avec terme source raide (publication SISC)
  - Écoulements multi-régime de Mach
- anti-diffusive scheme for multi-fluid flow (Billaud-Friess, Kokh)
- « all Mach » hydrodynamical scheme (Chalons, Kokh, Stauffert,…)
- *New* numerical schemes : DG , MOOD,….

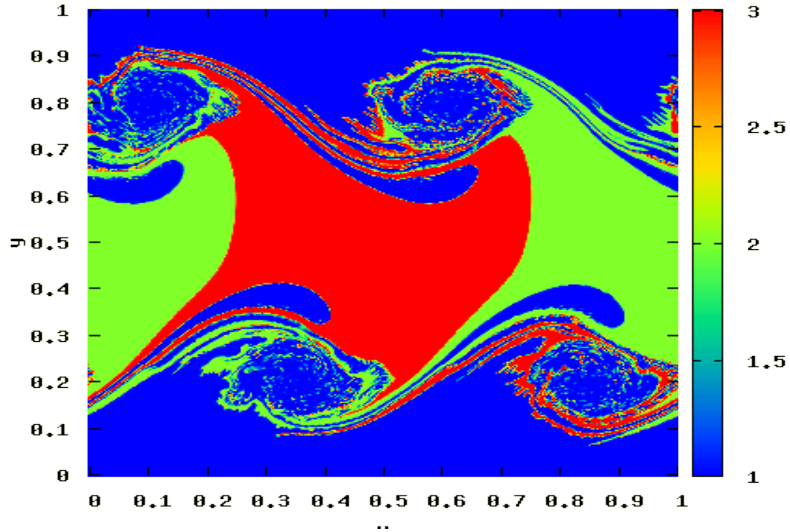

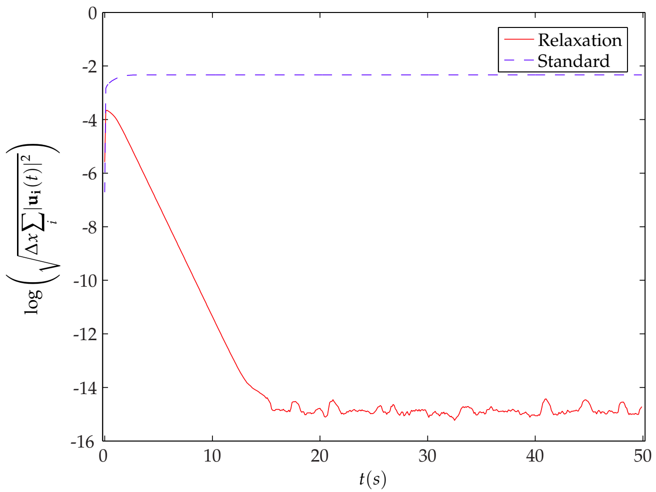

- Simulation d'écoulements magnéto-hydrodynamiques (Audit, Nkonga, Vides)
- Solveur de relaxation pour un modèle Euler-Poisson en écoulement gravitationnel (Audit, Berthon, Braconnier, Nkonga, Vides)

Paradigmes de Programmation

Arrivée de l'Exascale : un changement de paradigme *La forte hiérarchisation des machines à venir et la croissance du nombre de niveaux, associés à la fin de la loi de Moore telle que l'on la connaît, conduisent à introduire de nouveaux langages et paradigmes de programmation*

- Programmation *multi-niveaux :* Composants, parallélisme distribué, multi-threading
- Assurer la ré-utilisabilité du code a long terme
- Séparer logique applicative et parallélisation

Quelques résultats :

- Un modèle basé *workflow/dataflow* pour le HPC (avec A.L. Drummond – LBNL)
  Support efficace du parallélisme multi-niveaux mis en œuvre dans **YML**
- Intégration des paradigmes *wokflow* et *PGAS* (avec U. Tsukuba, Japon)
  YML (MdlS) + **XMP** (U. Tsukuba, Japon,testé sur *K* et Hooper à LBNL)
  Utilisé pour l'implémentation de méthodes MERAM
- Un modèle pour la modularisation de code (avec équipe Avalon - Inria / ENS Lyon)
  Séparation logique applicative / parallélisation (1 code, 3 parallélisations)
  Mis en œuvre dans **L²C**
- Une approche multi-niveau pour la tolérance aux pannes
  Mis en œuvre dans **FTI** (F. Capello - Argonne National Lab.)

Algèbre linéaire

- **Aujourd'hui :**
  - Utilisation de l'état de l'art pour aborder des problèmes
    - d'astrophysique (doctorat)
    - de modélisation électrique du cœur (CDD INRIA)
  - Former à l'utilisation des grandes bibliothèques

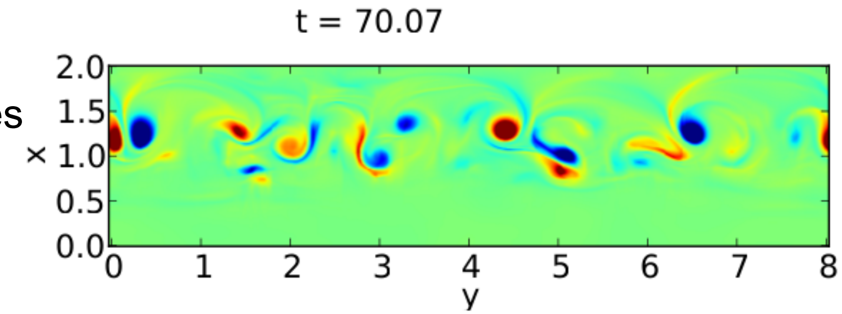

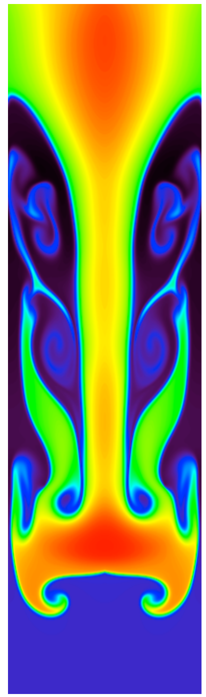

- **Tester les limites :**
  - Etudier la portée des méthodes actuelles : projet GENCI sur CURIE

- **Travailler sur les algorithmes de demain :***De nouvelles méthodes numériques doivent être introduites pour le calcul exascale (produits scalaires bloquants pénalisants):* ***Éviter les communications, minimiser l'énergie, être tolérant aux pannes****,...*

  → *« algèbre linéaire intelligente ».*

Principaux résultats de recherche :

- Méthodes numériques hybrides pour la résolution de grands problèmes « difficiles »
- Stratégies d'auto-tunning pour les méthodes de Krylov
- Faits marquants: application à la diffusion épidémique, expérimentation sur plusieurs machines : CURIE, Hooper (LBNL), Poincarré (MDLS),….algorithmes d'auto-tunning inclus dans une bibliothèque, accélération des performances avec des MIC ou GPU,..

Plateforme CanoP

CanoP : Adaptive Mesh Refinement applications - 2017

- **At MDLS:**
  - [done] **Poisson solver** (O. Iffrig, PK) for self-gravitating accretion disk studies, M2 internship in 2018 with IRFU/SAp (P. Hennebelle)
  - [on-going], **time-adaptivity** (O. Iffrig), prototyping in python before integration into canoP
  - [on-going] **C++ refactoring towards static AMR on accelerators** (PK): use Kokkos perf. portability library, expected by end of March
- **Collaboration with A. Allou** (CEA/DEN, Cadarache) + PhD Fang Chen (2nd year): **two-phase flow application, nitrogen / sodium jet study**
- **Collaboration with M. Massot, S. Kokh, A. Larat**, F. Drui, PhD defended in July 2017, two-phase flows
- **Collaboration with M. Massot, S. de Chaise Martin**, M. Essadki, PhD to be defended in 2018, kinetic models for spray / polydisperse flow, integrated into canoP

F. Drui, PhD

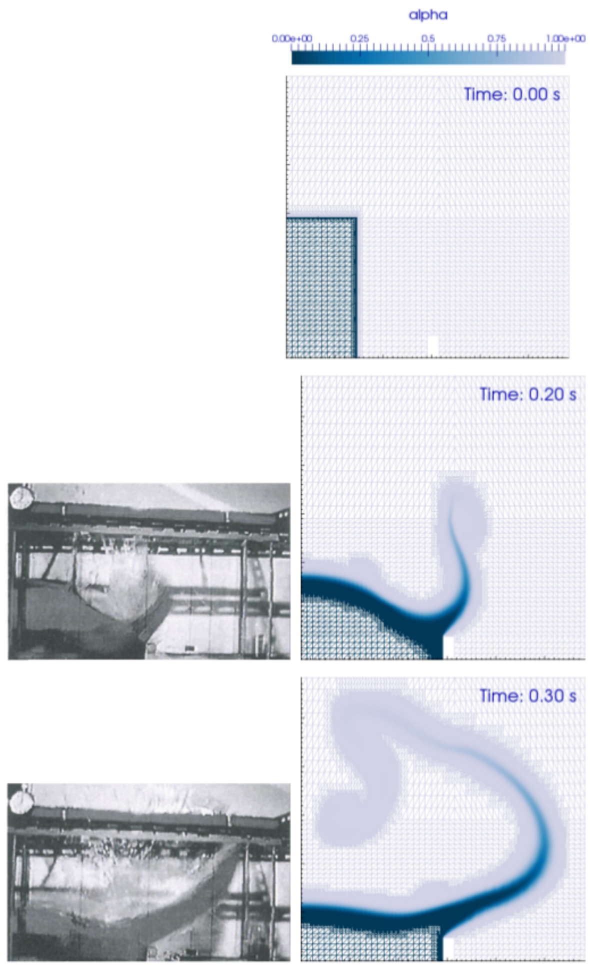

Plateforme CanoP

Kokkos R&D for CFD applications on cartesian and AMR grids

- **R&D performance portability: Kokkos**
- Code `euler_kokkos`: it's a rewrite of RamsesGPU (70 kSLOC) with `kokkos`,
  - Only 15 kSLOC, **compressible Hydro and MHD solver** (MUSCL-Hancock + Constraint Transport), no-AMR, cartesian grid; tested on both **K80 and new P100 GPU (IDRIS/Ouessant prototype)**
  - Already used by T. Padioleau, 1st y. PhD (All-Mach regime scheme adaptation), and C. Tenaud (CNRS/LIMSI)
- Code `ppkMHD`: same as `euler_kokkos`, with **high-order schemes**
  - **MOOD schemes** (not so efficient)
  - **Spectral Difference Method** (very efficient, similiar to DG): tested up yo order 6 in space, and Runge-Kutta RK54 (4^{th} order) in time
  - Test on **IBM P8 CPU** and **Nvidia Pascal P100**; Results presented at GTC2017 conference

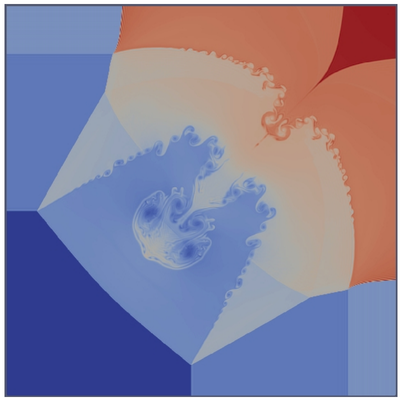

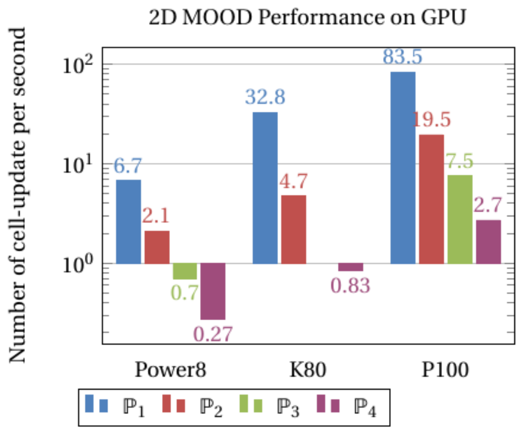

- Code `khamr`: R&D, standalone (no `p4est`) AMR rewriting using `kokkos`

Matériaux pour l'Énergie (ERC de M. Salanne)

Mathieu Salanne, chaire d'excellence de la Maison de la Simulation

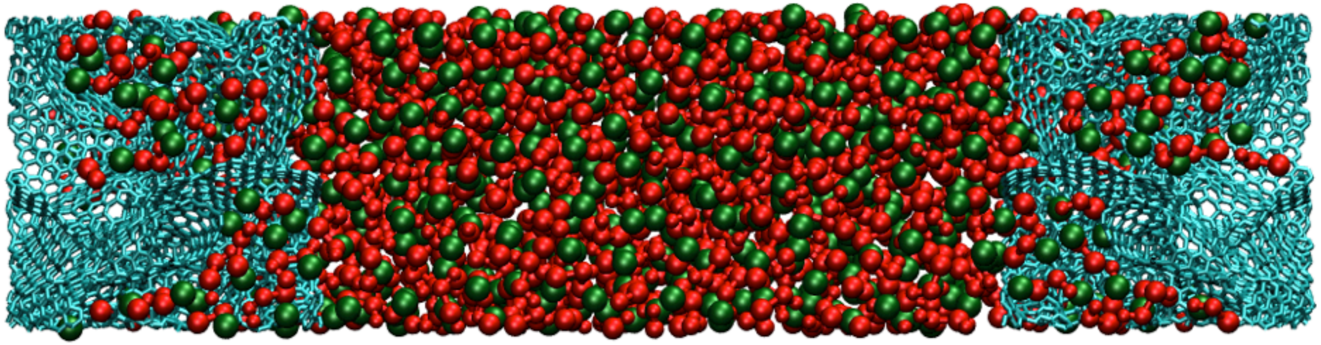

Merlet et al., Nature Mater., 2012; Nature Comm., 2013

Supercondensateurs: stockage électrochimique de l´ énergie

Carbones nanoporeux: performances exacerbées

Objectif: criblage haut-débit de supercondensateurs carbone-carbone

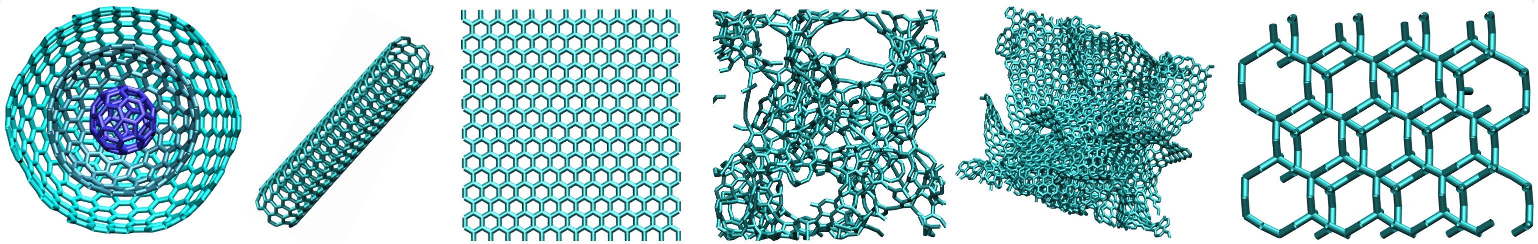

Nouveaux matériaux: carbones dopés, nanotubes, graphène… (G. Bidan, ANR SUGGEST)

Objectif: extension du code de simulation Metalwalls à ces types de matériaux

Collaboration avec T. Deutsch, P. Pochet (INAC) —> *Calculs ab initio* (bigDFT)

Extension aux matériaux d'anodes de batteries (Li-ion, Na-ion)

Matériaux pour l'Énergie

- Développement de Metalwalls (A. Marin-Laflèche et M. Haefele)
  - Initialement : F90 / pure MPI
  - Portage sur Intel Xeon (vectorisation + OpenMP) → perf x 3,5 / version initiale
  - Portage sur Intel Xeon Phi KNL → perf +30% comparée à Xeon (5 noeuds vs. 5 noeuds)
  - Portage sur IBM OpenPower → perf x 2 comparée à Xeon (1 noeud vs. 5 noeuds)
- Travail en cours : implémentation FPGA avec la technologie Maxeler (A. Marin-Laflèche et M. Haefele)
  - M. Haefele a passé 2 semaines à Londres chez Maxeler
  - Implémentation de 50% des noyaux de la boucle en temps
  - Résultats corrects dans l'émulateur
  - Premiers tests sur le vrai matériel en cours
- Développement de nouveaux algorithmes (M. Burbano)
  - Nouvelle méthode d'optimisation des charges → perf x 2 à 3
  - Benchmarking vs. LAMMPS
- Simulation de nouveaux matériaux pour supercondensateurs (M. Burbano, Z. Li et T. Mendez-Morales)
  - Électrolytes redox (Li et al., J. Phys. Chem. Lett., 8, 1925, 2017)
  - Electrodes en graphène (Mendez-Morales et al., J. Chem. Phys., sous presse + 1 article en cours)
  - Désalinisation de l'eau de mer (Simoncelli et al., soumis)

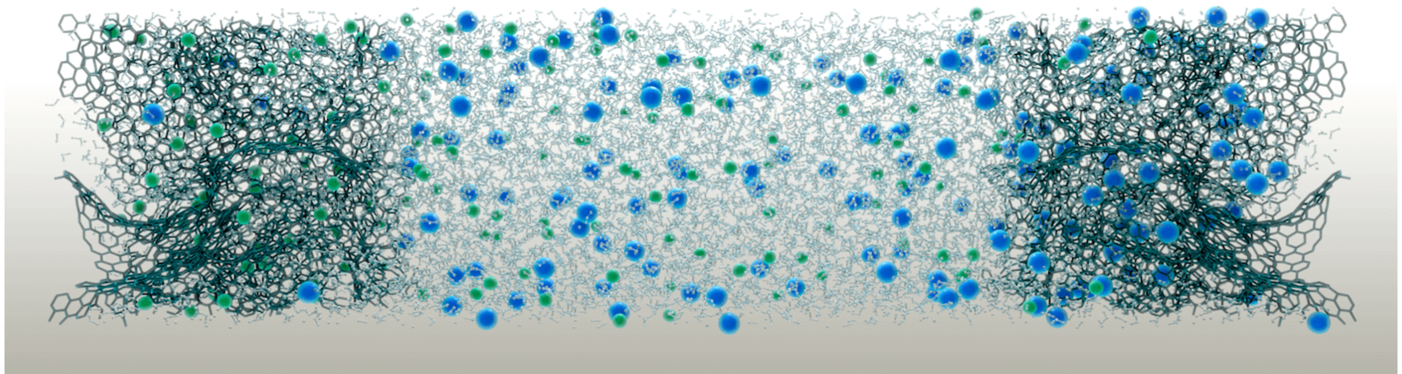

ATMO: a ERC Starting Grant project for HPC numerical simulations of physical processes in the atmosphere of exoplanets (PI: P. Tremblin)

- Goal: Model with 3D startified turbulent simulations the process of fingering convection
- Mean: by using efficient and innovative numerical methods combined with the computation power of new HPC architectures (GPU/XeonPhi)
- Impact: unable to interpret spectroscopic observations of exoplanets with future space mission such as the James Webb Space Telescope (JWST)

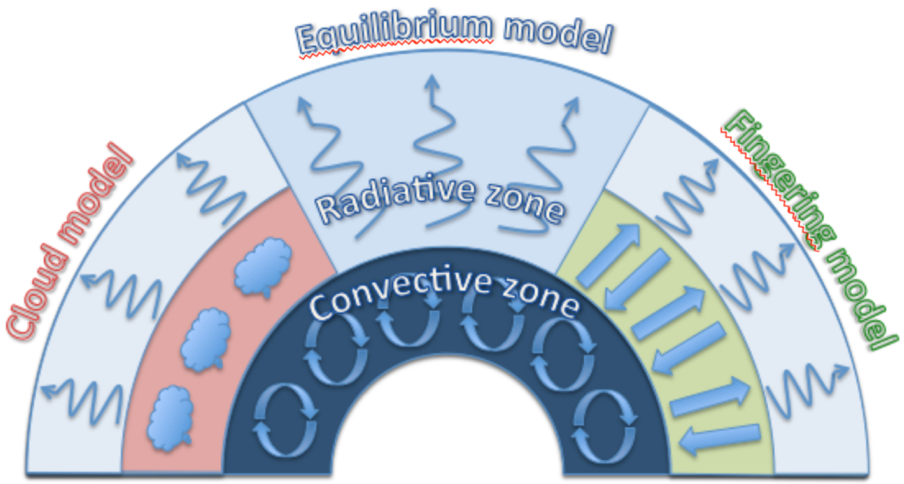

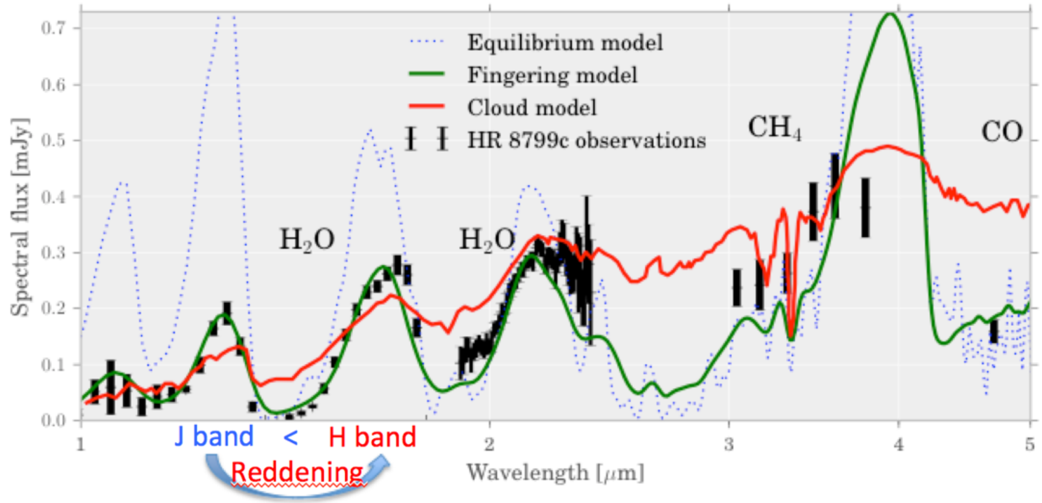

Co-Développement de code : SMILEI

Software for Matter Interacting with Laser at Extreme Intensities

Développement d'un code ouvert pour l'exploitation scientifique du laser Apollon

- Développement d'un code PIC collaboratif orienté Hautes Intensités
  - Objectif principal : Apollon

- Développé par la communauté PIC du Plateau de Saclay :
  - CEA : IRAMIS/SPAM/PHI, IRAMIS/LSI, IRFU/SACM, INSTN
  - CNRS : LLR, LOA, LPGP, LULI
  - Maison de la Simulation

- Maison de la Simulation :
  - Architecture modulaire et performante (adaptabilité aux différentes applications physiques)
  - Parallélisme massif, nouvelles architectures
  - Gestion des outils connexes au projet

- La MdS centralise et intègre les contributions des physiciens, maintien le code

Co-Développement de code : SMILEI

Software for Matter Interacting with Laser at Extreme Intensities

Développement d'un code ouvert pour l'exploitation scientifique du laser Apollon

- Finalisation des *grand challenges* sur Occigen et Turing
- Nouveau parallélisme hybride MPI/OpenMP
  → Communication au **European Advanced Accelerator Concepts Workshop**
- Extension de la communauté des utilisateurs. Démarrage de 4 thèse
- Demande ANR
- 14 millions d'heures demandées sur Turing/Curie

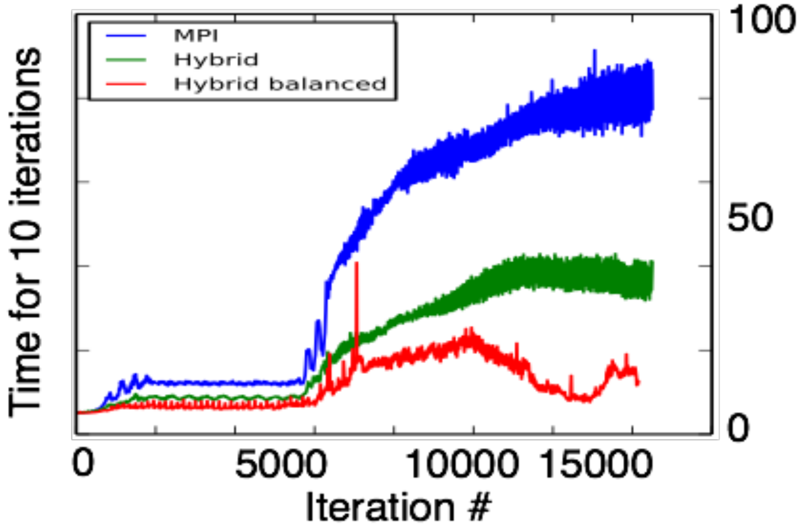

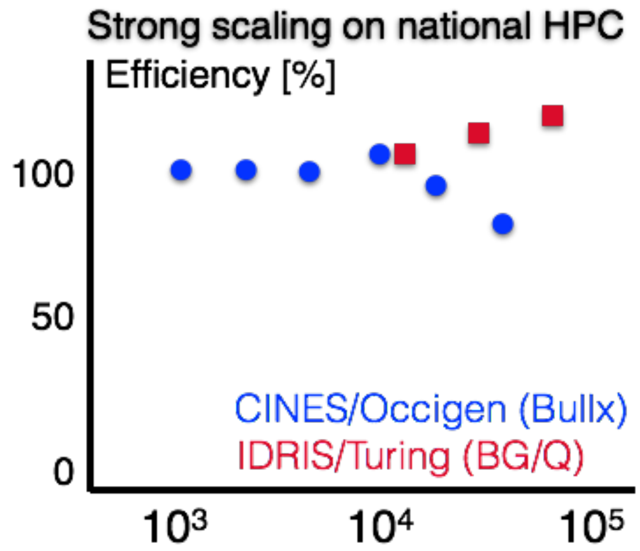

Support a Gysela5D

- Gysela5D : un code CEA/IRFM
  - but : comprendre le développement des turbulences dans ITER
- Support a la maison de la simulation : 1 post-doc, 1 thésard
  - Portage sur Blue Gene/Q
    - I/O retravaillées: résultats et checkpointing
    - Collaboration avec Argonne National Lab.
    - Schémas de communication MPI réécrits
    - Parallélisme intra-nœud amélioré
    - **Résultat**: scalabilité **91% efficacité @ 1,8 M.threads**
  - Amélioration de la scalabilité mémoire
    - Prediction de consomation
    - **Conso -50% @ 32k.coeurs**

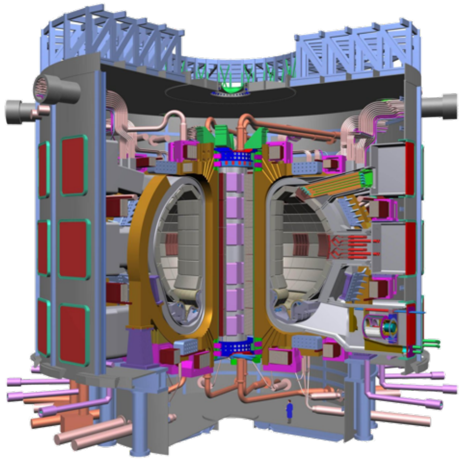

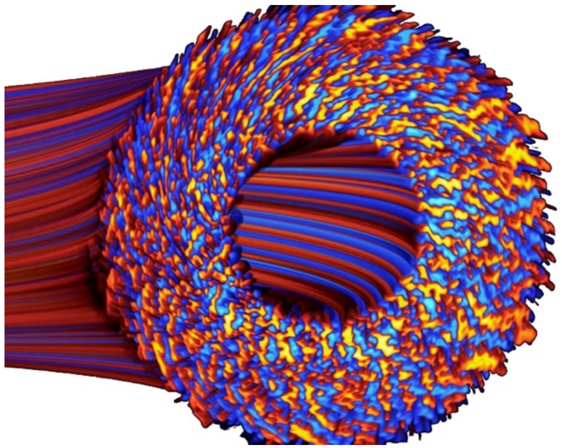

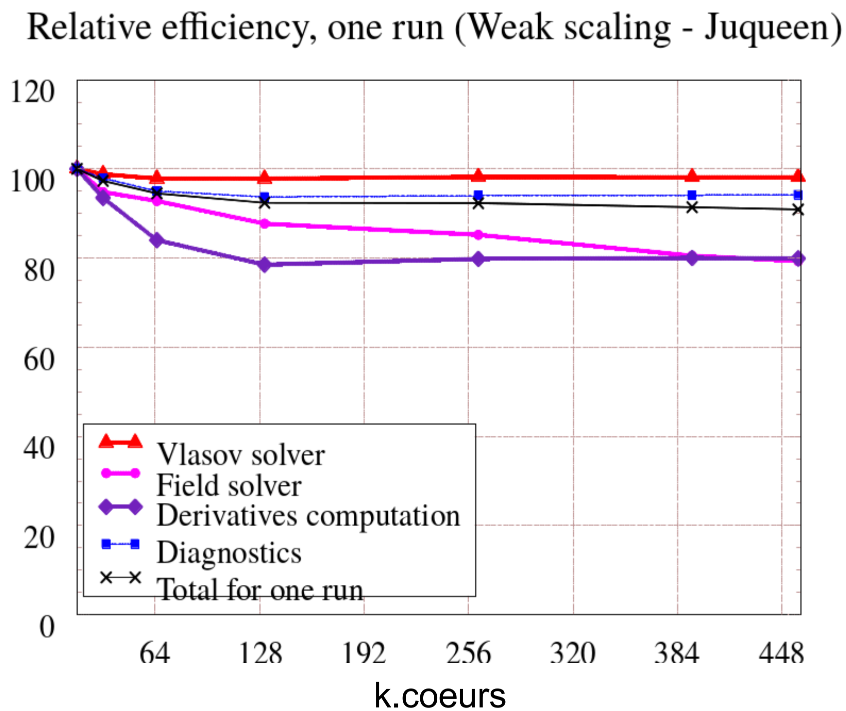

European HPC strategy

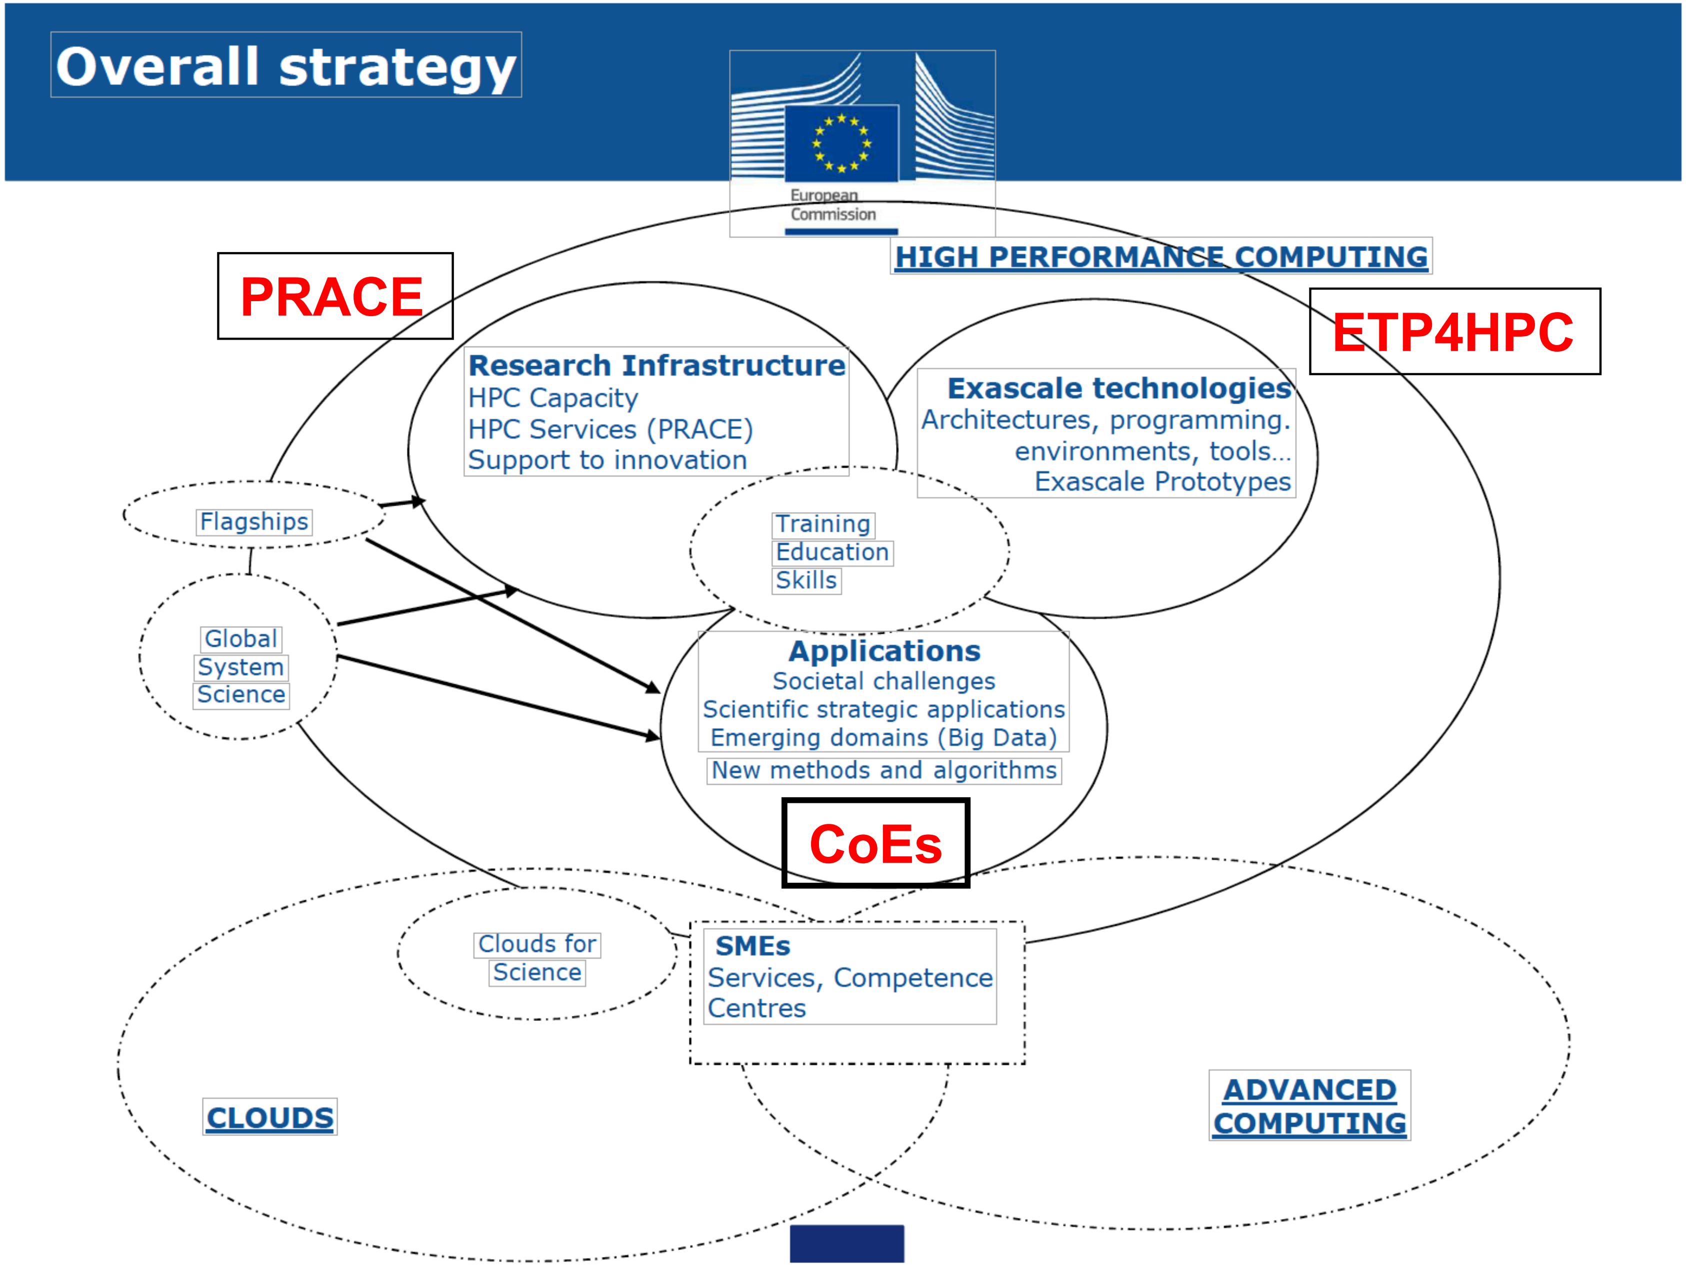

PATC – Course in 2017-2018

Titre	Organisation	Date	Participants
Performance portability for GPU application using high-level programming approaches	MdlS	Jan	20
Parallel linear algebra	Inria / VSB Rep. Tchèque	Fév.	40
Parallel filesystems and parallel IO libraries	MdlS (+ centres)	Mars	16
C- C++ Multicore application programming	MdlS	Mars	7
Advanced usage on Curie : Best practice for current and future HPC architectures	TGCC	Avril	8
Uncertainty quantitfication	EdF / EADS / Phimeca	Mai	9
Mastering GPU-Acceleration on OpenPOWER Platform	Idris	Novembre	18
Debugging & Optimization	Cines	Décembre	8

E-CAM Centre of Excellence

- **Molecular modelling for materials and biology**
  - *Quantum dynamics*
  - *Electronic structure*
  - *Molecular dynamics*
  - *Multi-scale modeling*

- **Creating modular open source codes, libraries and interfaces**
- **Training junior scientists and personnel from industry**
- **Establishing a network of experts to advise European industry**

Projet 2016-2020

cecam - MOSER

Centre Européen de Calcul Atomique et Moléculaire

- Délivrable « Hardware developments II », 06/17
- Nouvelle politique HPC/HTC
- Assemblée Générale, 29-30 Nov. 2017

A venir:

- Départ du développeur logiciel, Mai 2018
- Embauche de 2 post-doc (QMC et Quantum Dynamics)
- Organisation d'un SAW sur QMC et d'un ESDW
  sur Quantum Dynamics

EoCoE

<u>Main objective :</u> Using the prodigious potential offered by the ever-growing computing infrastructure to foster and accelerate the European transition to a reliable and low carbon energy supply.

EoCoE is at the crossroad of the numerical and energy revolution

The present revolution in hardware technology calls for a similar paradigm change in the way application codes are designed.

EoCoE assists the energy transition via targeted support to four renewable energy pillars: Meteo, Materials, Water and Fusion.

These four pillars are anchored within a strong transversal multidisciplinary basis providing high-end expertise in applied mathematics and HPC.

EoCoE consortium

Foster the European transition to a reliable low carbon energy supply using HPC

8 countries, 22 partners, 5.5 M€, lead by MdlS
3 years project, started in October 2015, ends in September 2018

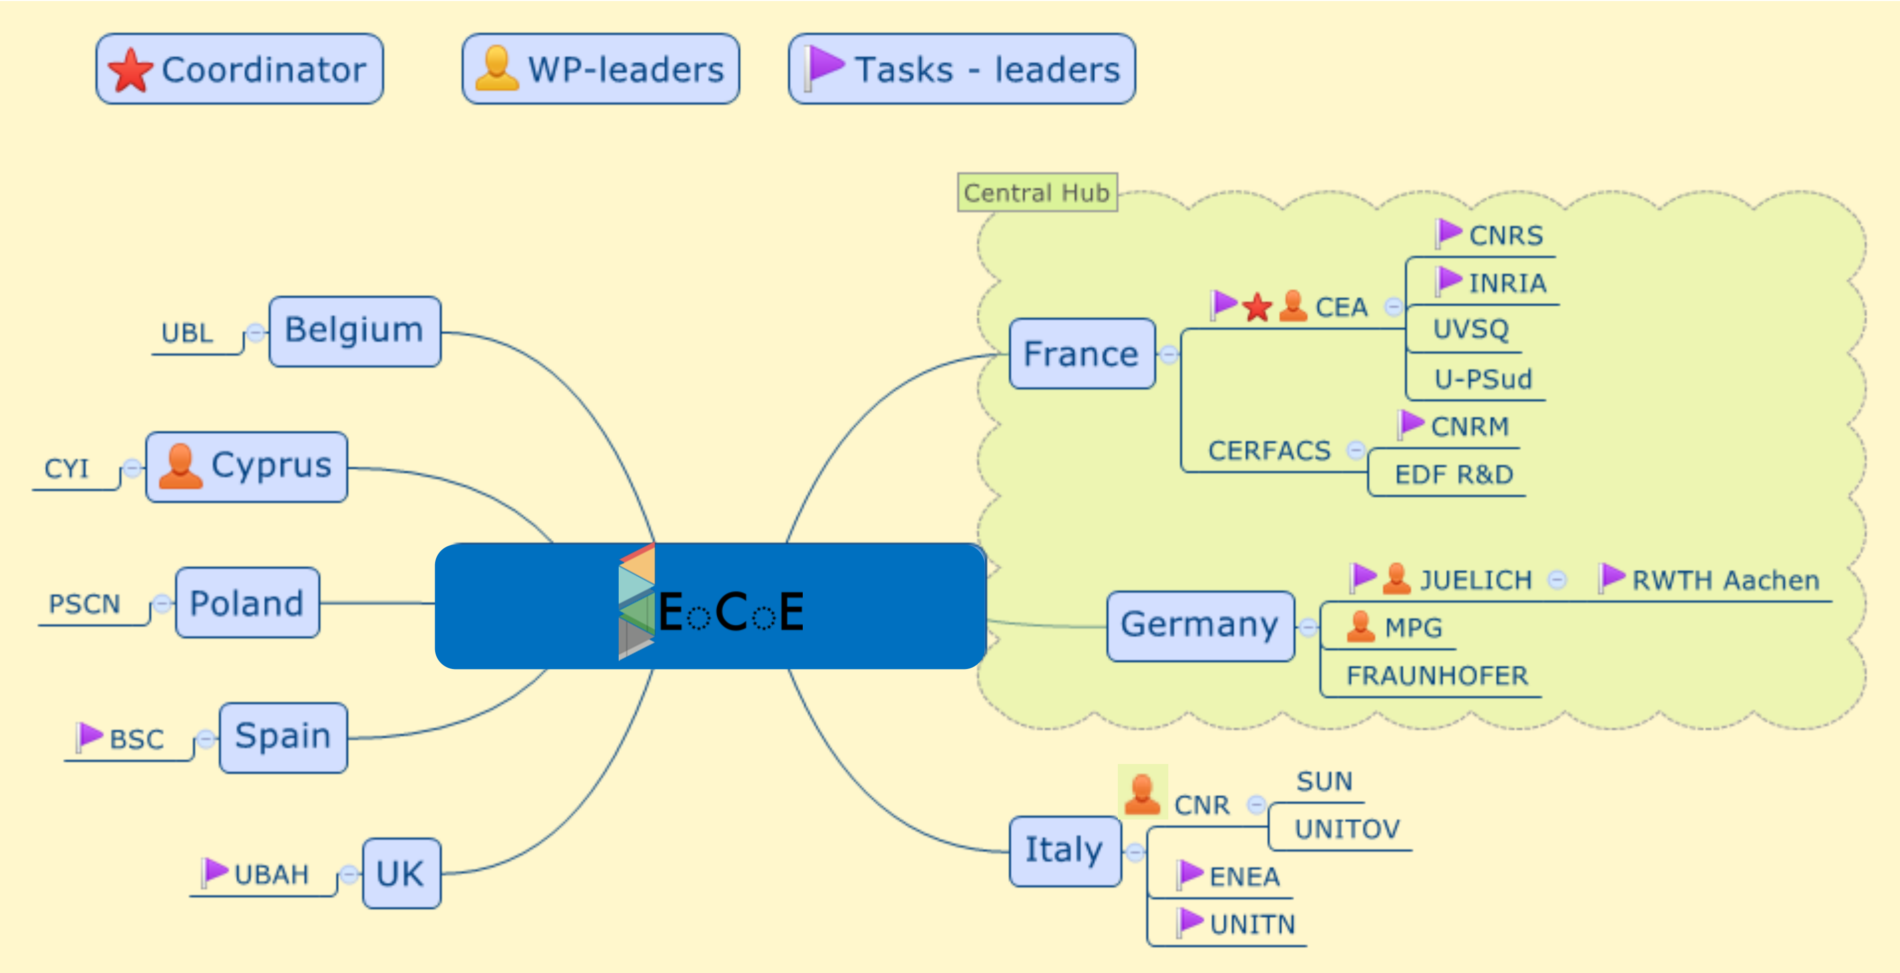

EoCoE overview

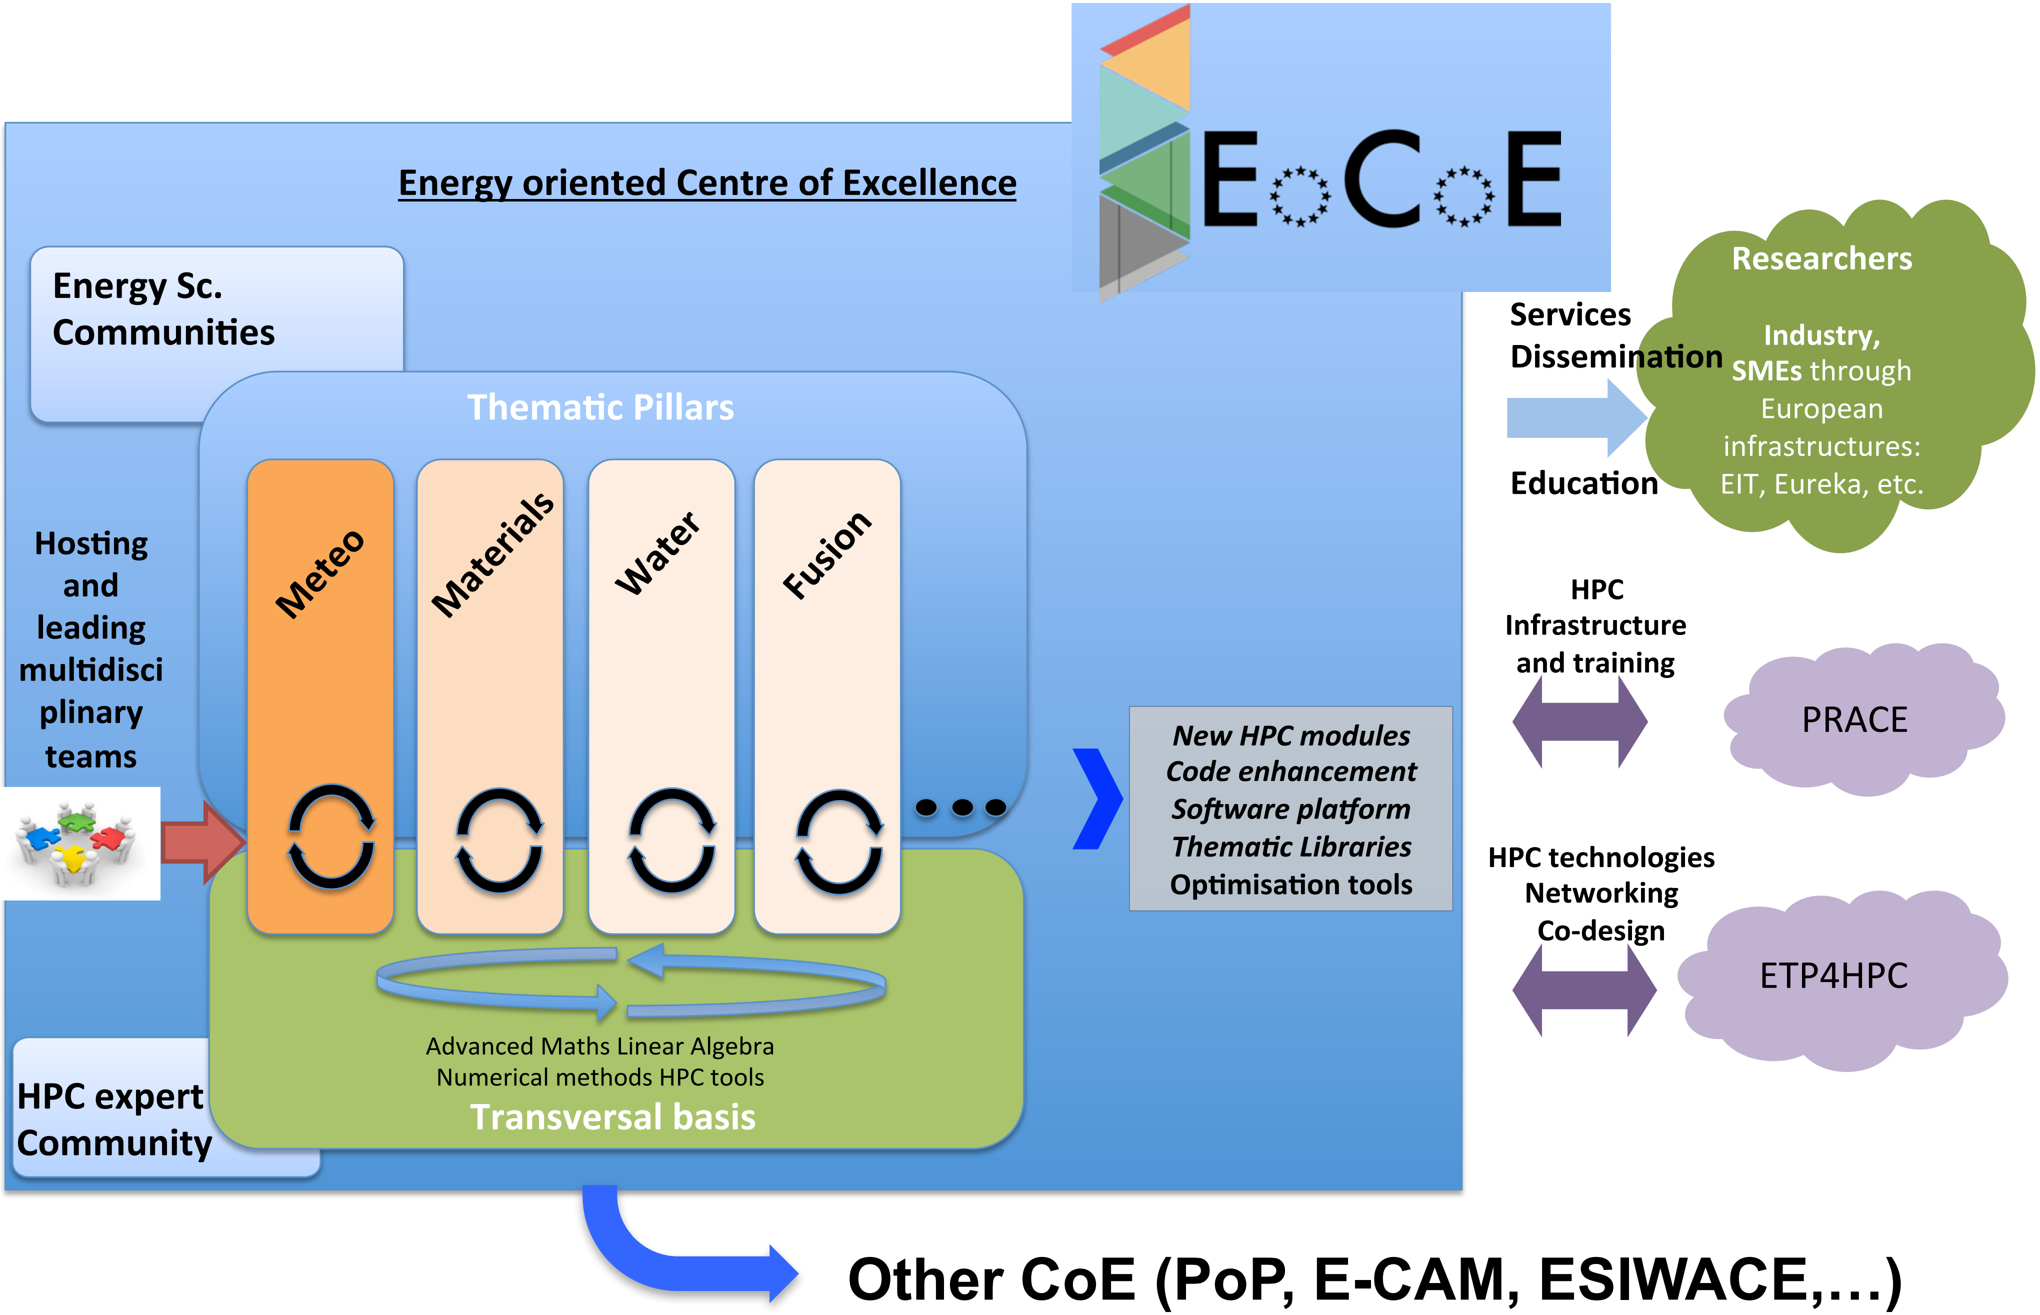

Objectives - I

Develop or optimise high end tools and software for all the communities. Get ready for exascale computers

- Applied math and numerical methods
- Linear algebra
- System tools for HPC
- New programing models

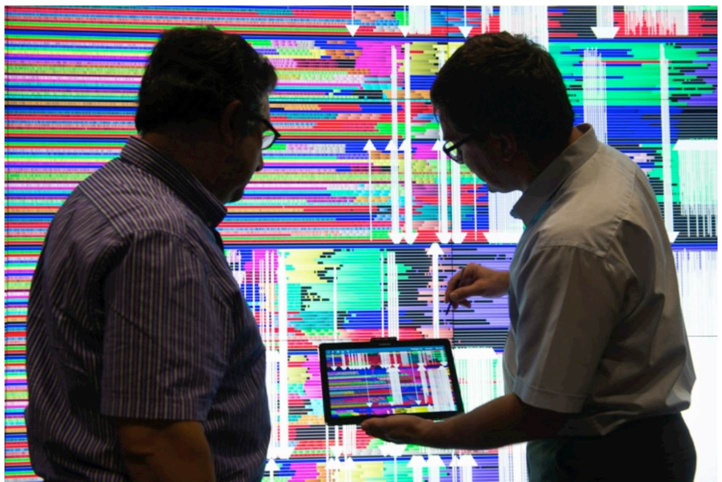

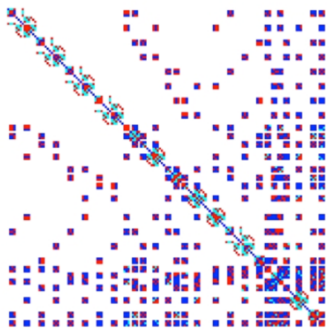

Objectives - II

*Help accelerate transition to **low-carbon electricity***

Meteorology for energy : Very short term forecast to predict the production of solar and wind farm – Efficient coupling to the grid, energy trading, prediction of extrem events

Fusion for energy : Coupling kinetic and fluid codes, mesh aligned with equilibrium configurations.

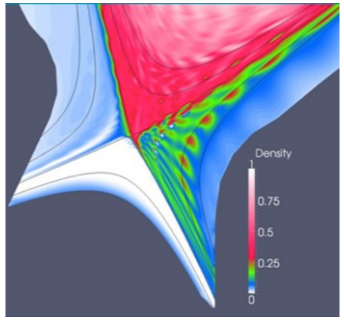

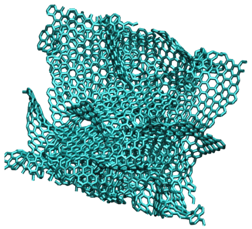

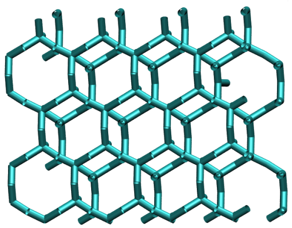

Material for energy : Photovoltaic cells, batteries and super-capacitors

Water for energy : Geothermal and hydro-power – management of resources, strategy of usage - influence of climate change.

HPC/Exascale approach - I

A systematic approach for code monitoring and performance analysis

- Compute
- Memory
- I/O
- Communications
- …..

Detailed Performance metrics

Code Diaries

- Development of a complete metric sheet and associated tools to evaluate codes in a systematic way.
- Fully automated evaluation workflow
- Code diaries and follow-up

- ❑ **Make the communities performance aware**
- ❑ **Establish a clear road-map for optimisation**
- ❑ **Monitor the progress**
- ❑ **Identify the main bottleneck for porting to exascale architecture**

Code performance monitoring

TOOLS

Code	WP	JSC Account	Data server account	Gitlab account	JUBE integration	benchmarks defined in JUBE	Tools integrated in JUBE	Allinea report	Score-P profile	Score-P trace	Scalasca analysis	Vampir analysis	Extrae measurem	Paraver analysis	Darshan results	VTune analysis	Advisor analysis	Performance rep	Total Progress (%)	Handled b	Support ac	Code	
ALYA	WP 2	2	2	2	2	2	2	2	2	2	2	2	2	2	0	2	1	2	100	MdlS	done	ALYA	1
ESIAS	WP 2	2	2	2	2	2	2	0	2	2	2	0	0	0	2	0	0	2	100	JSC	no	ESIAS	2
Metalwalls	WP 3	2	2	2	2	2	2	2	2	2	2	2	2	2	2	0	2	2	100	MdlS	done	Metalwall	3
PVnegf	WP 3	2	2	2	2	2	2	2	2	2	2	1	0	0	2	0	0	0	90	JSC	done	PVnegf	4
SHEMAT	WP 4	2	2	2	2	2	1	2	1	1	1	0	2	2	2	0	0	1	90	JSC	done	SHEMAT	5
ParFlow	WP 4	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	100	JSC	on-going	ParFlow	6
GYSELA	WP 5	2	2	2	2	2	2	1	1	1	1	0	2	2	2	0	0	2	100	MdlS	done	GYSELA	7
nowcast system	WP 2	2	2	2	2	2	2	2	0	0	0	0	0	0	2	2	2	2	100	MdlS	done	nowcast s	8
CP2K	WP 3	2	2	2	2	2	2	0	2	2	1	2	1	0	2	0	0	2	100	JSC	no	CP2K	9
MDFT	WP 3	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	0	2	100	MdlS	done	MDFT	10
TELEMAC	WP 4	2	2	2	2	2	1	2	1	1	1	1	1	1	0	0	0	0	86	MdlS	starting A	TELEMAC	11
COMPASS	ext	2	2	2	2	2	1	2	1	1	1	1	1	1	0	0	0	0	86	MdlS	no	COMPASS	12
EIRENE	WP 5	2	2	2	2	2	2	2	2	1	1	0	1	0	0	0	0	0	86	JSC	no	EIRENE	13
TOKAM3X	WP 5	2	0	2	2	1	0	2	2	2	2	2	2	0	2	1	1	0	90	MdlS	on-going	TOKAM3X	14
PARCOMB	ext	0	0	0	2	2	1	0	2	2	2	0	2	2	0	0	0	0	68	JSC	no	PARCOMB	15
OpenFOAM	ext	2	0	0	2	2	2	0	2	2	2	2	2	2	0	0	0	0	90	JSC	no	OpenFOAM	16
ALYA	WP 2	0	0	2	1	0	0	0	2	2	2	0	2	0	0	0	0	2	100	MdlS	no	ALYA	17
MUMPS	WP1	2	2	2	2	2	2	0	2	2	2	2	2	2	0	0	0	0	90	MdlS	no	MUMPS	18
Maphys	WP1	2	2	2	2	2	2	0	2	2	2	1	2	2	0	0	0	0	90	MdlS	no	Maphys	19
CP2K	ext	2	0	0	0	0	0	0	2	0	1	0	0	0	0	0	0	0	23	JSC	no	CP2K	20
DL_MESO	ext	2	0	0	2	2	1	2	2	2	2	0	2	2	2	0	0	0	90	JSC	no	DL_MESO	21
Compass	ext	2	0	0	2	2	2	0	2	2	2	0	2	2	0	0	0	0	82	MdlS	no	Compass	22
DIVA	ext	2	0	0	2	2	2	2	2	2	2	2	2	2	0	0	0	0	90	MdlS	no	DIVA	23
WRF-Solar	WP 2	2	0	0	2	2	1	0	2	2	2	0	2	2	2	0	0	0	86	??	no	WRF-Solar	24

1

2

3

CODES

t.b.c.

Automated performance metrics

	Metric name	03/01/2016
	Test-case	case1
Golbal	Total Time (s)	43.2
	Time IO (s)	0.0
	Time MPI (s)	12.4
	Memory vs Compute Bound	1.1
	Load Imbalance MPI	24.8
IO	IO Volume (MB)	35.8
	Calls (nb)	384000
	Throughput (MB/s)	105.0
	Individual IO Access (kB)	0.1
MPI	P2P Calls (nb)	0
	P2P Calls (s)	0.0
	P2P Message Size (kB)	0.0
	Collective Calls (nb)	2721
	Collective Calls (s)	0.1
	Collective Message Size (kB)	908.4
	Synchro / Wait MPI (s)	11.7
	Ratio Synchro / Wait MPI	94.8
Node	Time OpenMP (s)	0.0
	Ratio OpenMP	0.0
	Time Synchro / Wait OpenMP	0.0
	Ratio Synchro / Wait OpenMP	0.0
Mem	Memory Footprint (B)	66 mB
	Cache Usage Intensity	N.A.
Core	IPC	N.A.
	Runtime without vectorisation (s)	46.5
	Vectorisation eff ciency	1.1
	Runtime without FMA (s)	44.6
	FMA eff ciency	1.0

Total execution time

Time spent in MPI

Memory Footprint

HPC/Exascale approach - II

Close collaborations between WP1 and WP2-5

→ Very significant results on real **production** applications

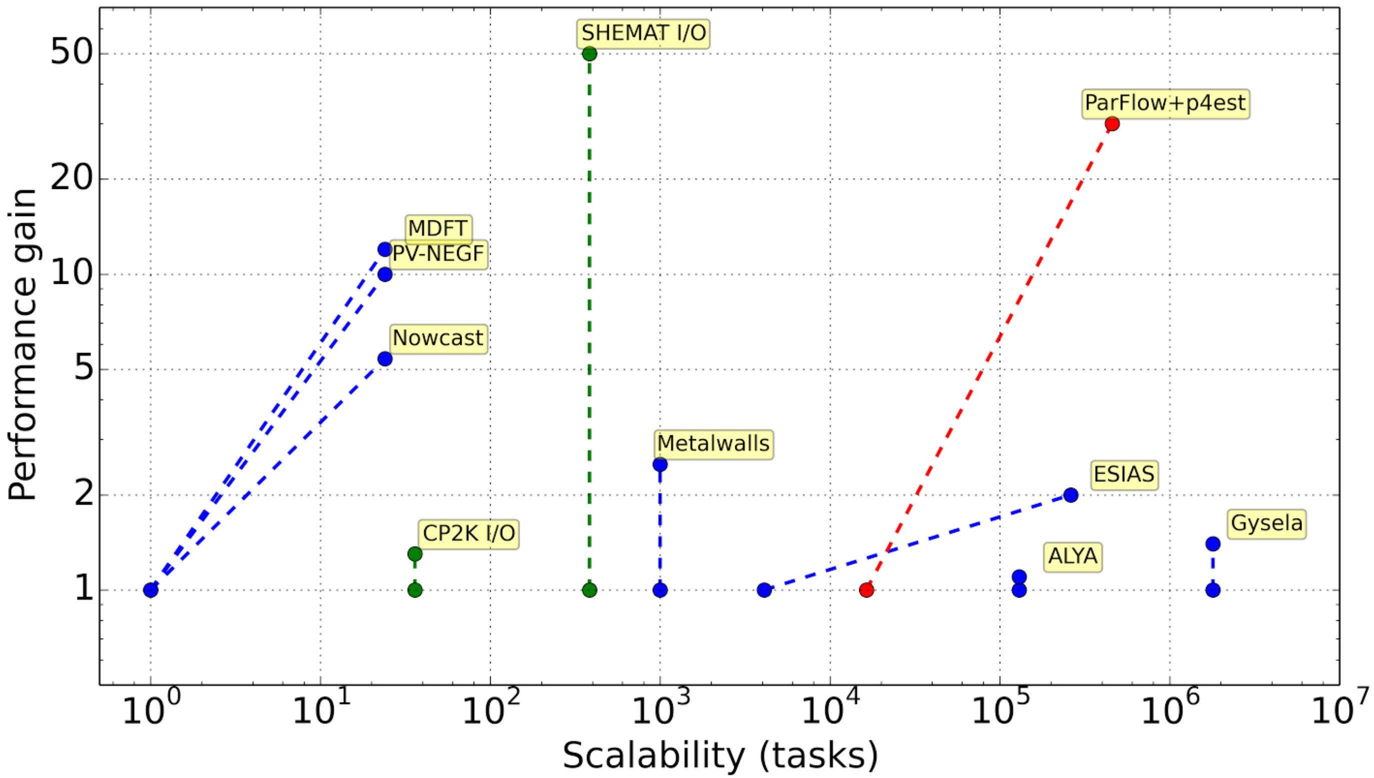

optimisation and parallelisation (blue), parallel I/O library (green), kernel refactoring (red).

- ❑ **Turning Exascale into benefit**

- ❑ Pushing relevant application toward exascale…
- ❑ … and bridging the skill gap for less advanced users.

Turning Exascale into benefit

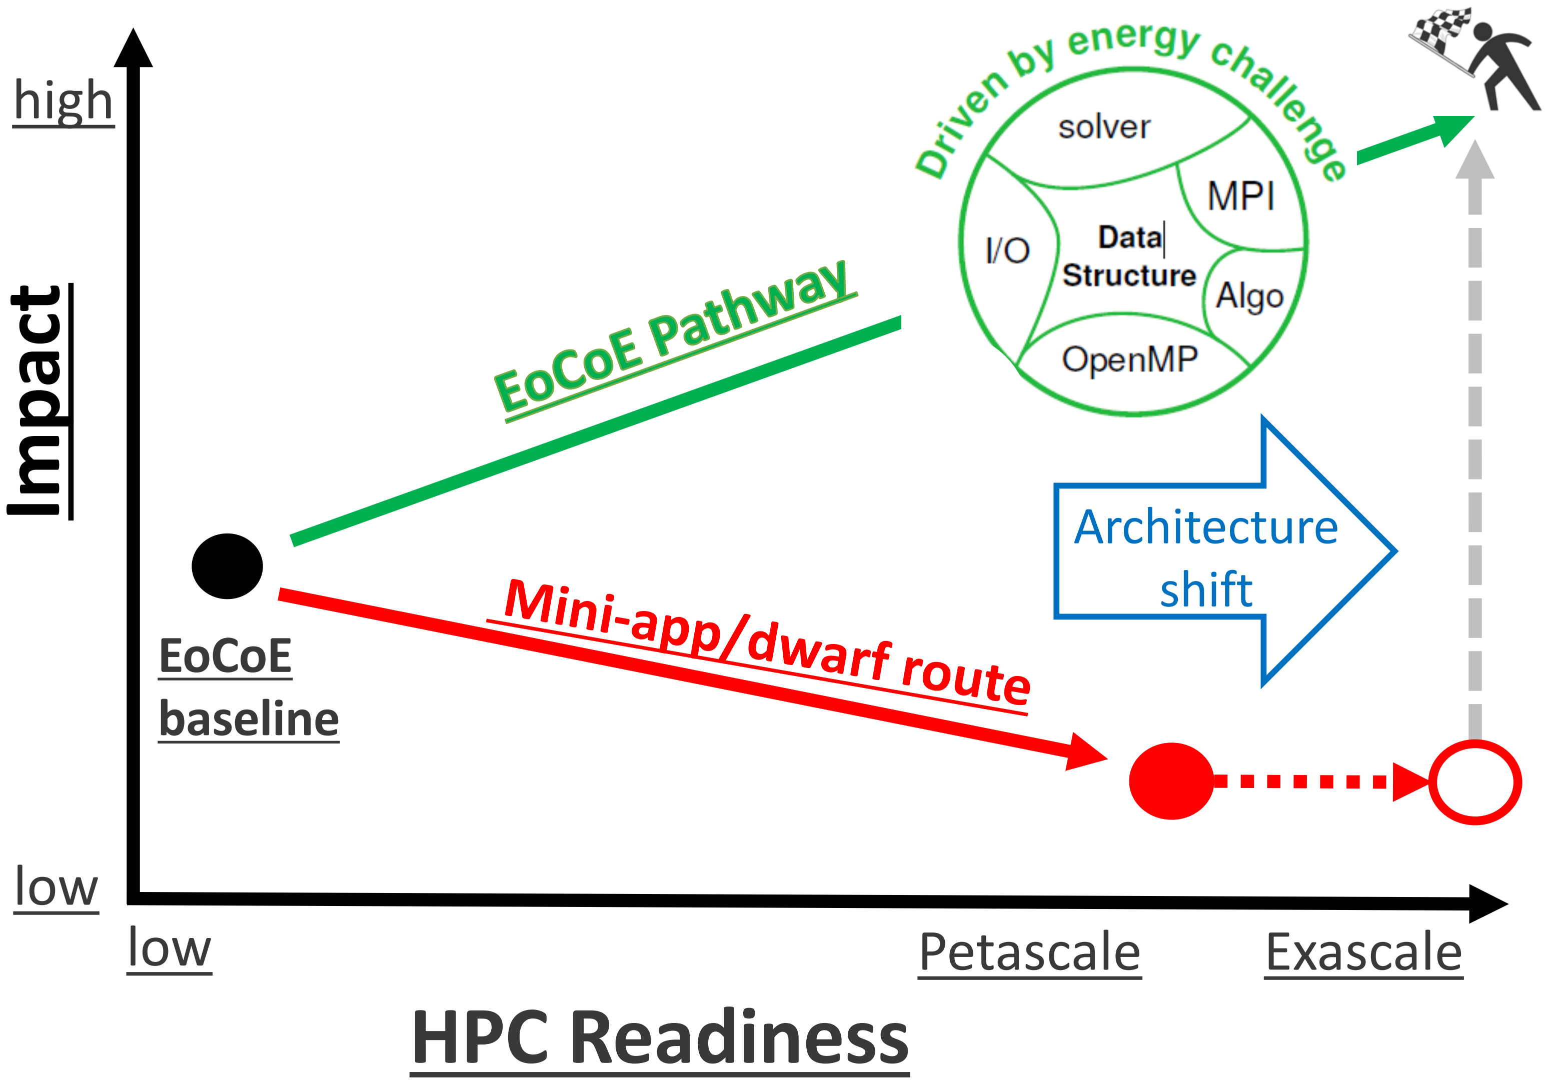

European HPC strategy

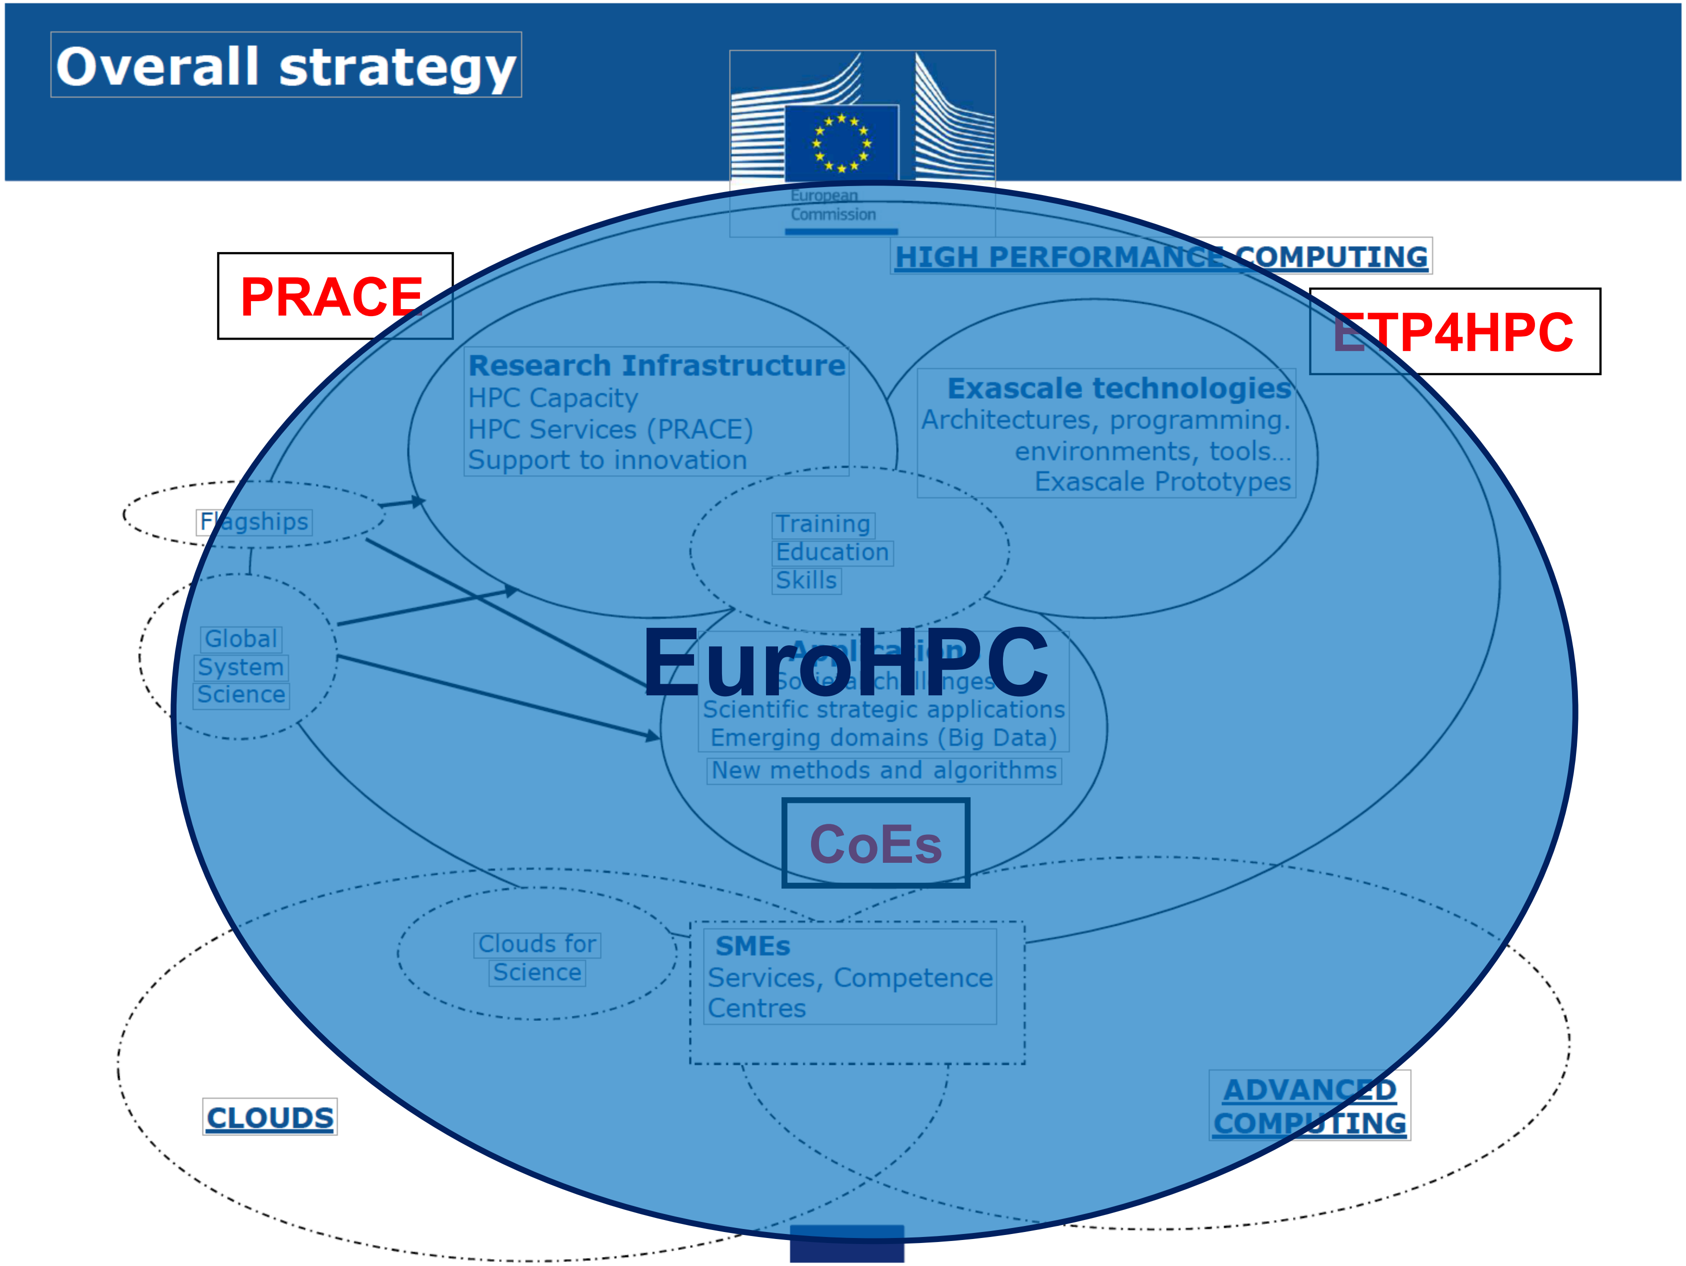

Global Picture HPC

- **USA**, 4 pre-exa and 3 exascale systems in 2018-2022
- **China**, exascale in 2021?
- **Japan**, exascale in 2022

2 pre-exascale by 2020 and two exascale systems by 2022/2023

Hybrid HPC/Quantum infrastructure

emerging "computing architectures" (quantum/neuromorphic)

novel applications in key areas (Cybersecurity, AI)

#EuroHPC (High Performance Computing) **Declaration**

Signatory European countries

Seven countries – France, Germany, Italy, Luxembourg, Netherlands, Portugal and Spain – signed the declaration in March 2017.

Since then, another nine countries – Belgium, Slovenia, Bulgaria, Switzerland, Greece, Croatia, Czech Republic, Cyprus and Poland – have also signed.

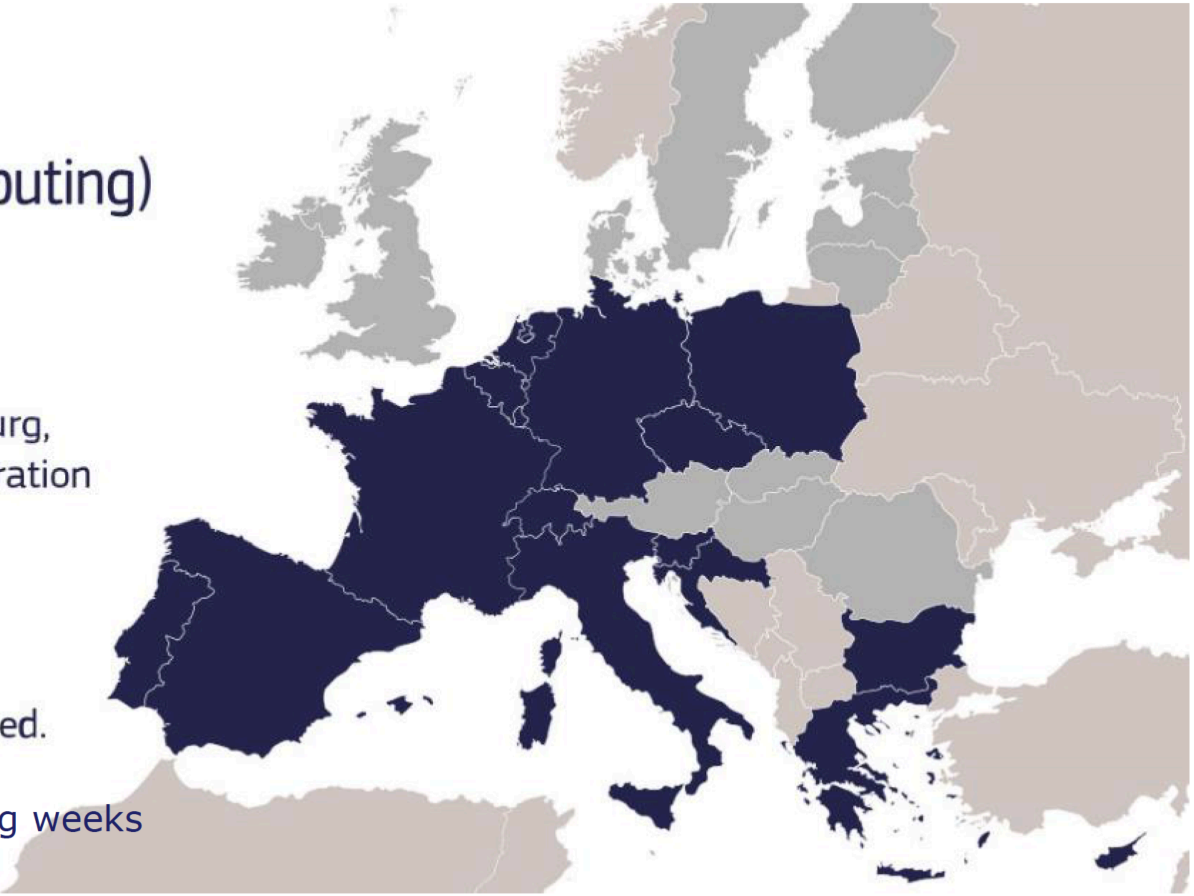

More countries are expected to join in the coming weeks

European Commission | Communications Networks, Content & Technology

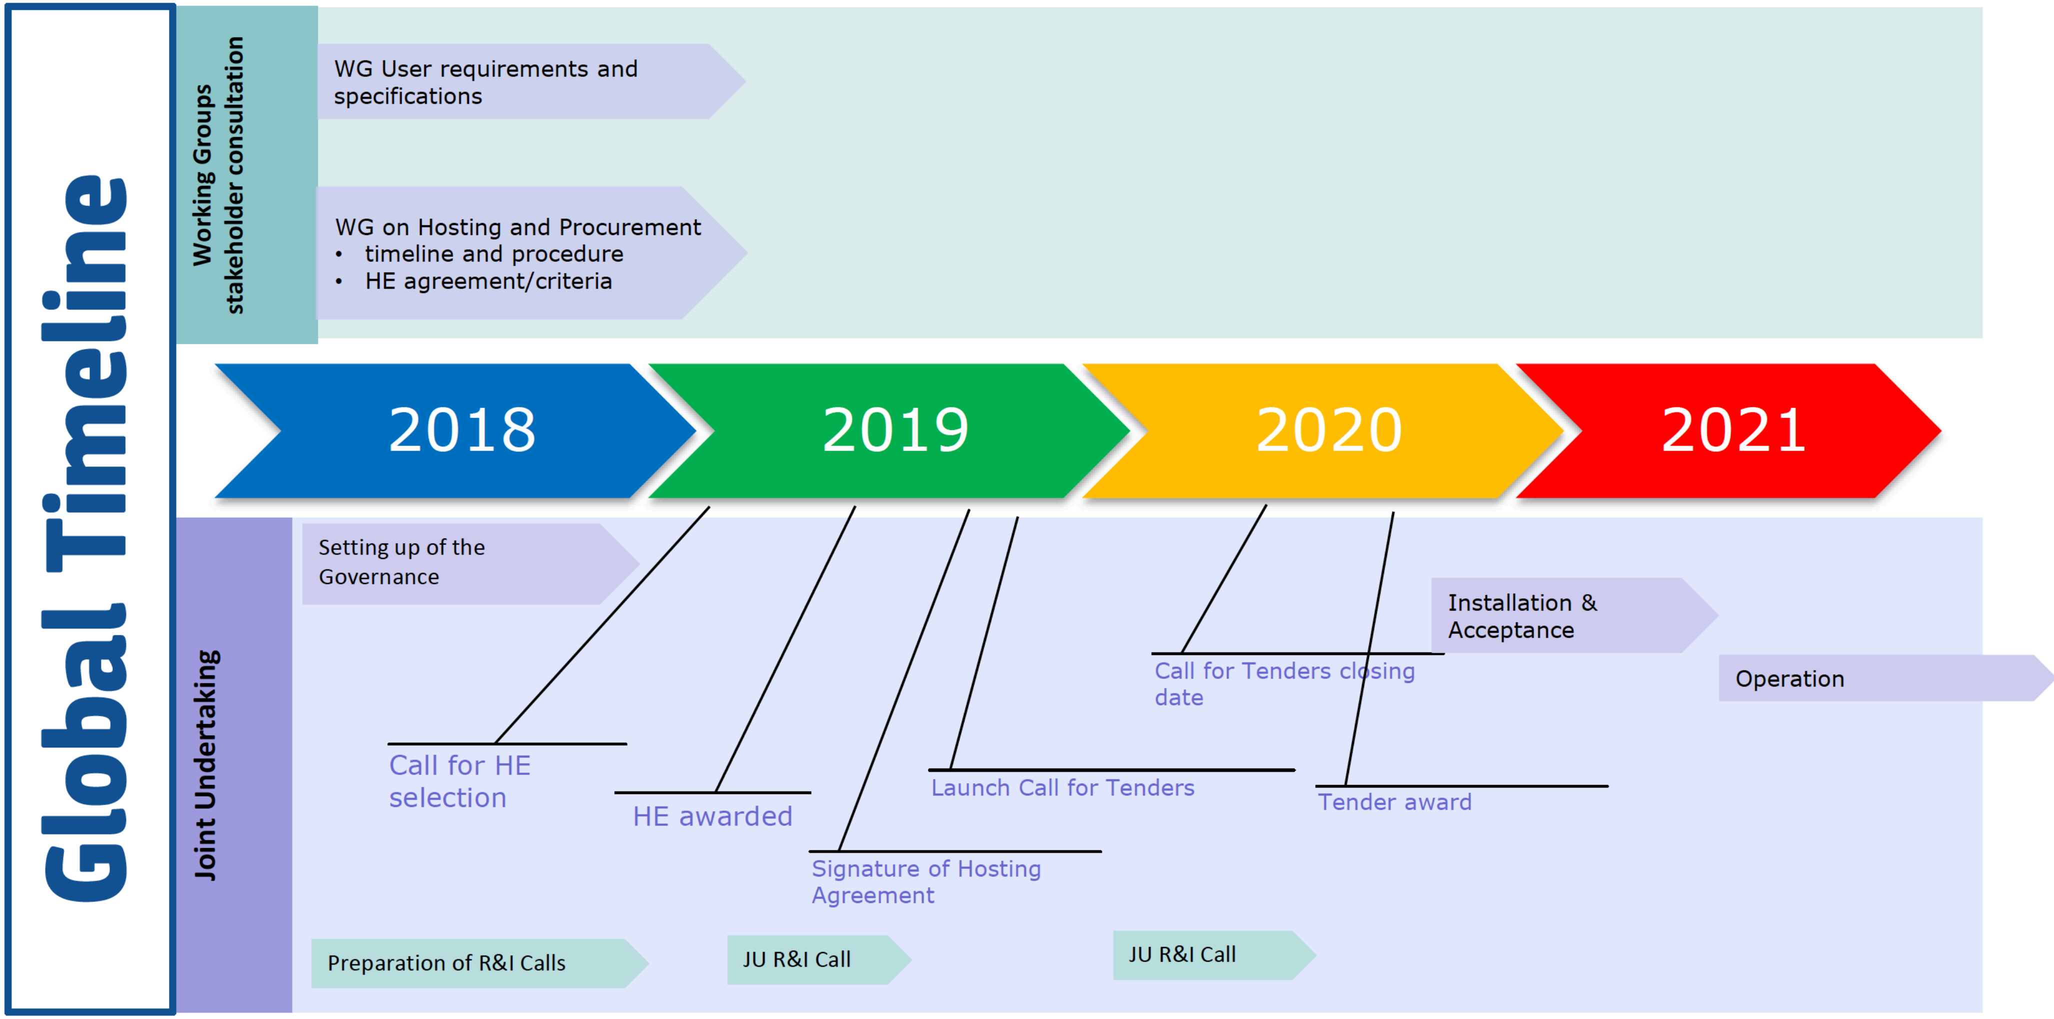
Global Timeline
Working Groups
stakeholder consultation
WG User requirements and specifications
WG on Hosting and Procurement
• timeline and procedure
• HE agreement/criteria
2018
2019
2020
2021
Joint Undertaking
Setting up of the Governance
Call for HE selection
HE awarded
Signature of Hosting Agreement
Launch Call for Tenders
Call for Tenders closing date
Tender award
Installation & Acceptance
Operation
Preparation of R&I Calls
JU R&I Call
JU R&I Call

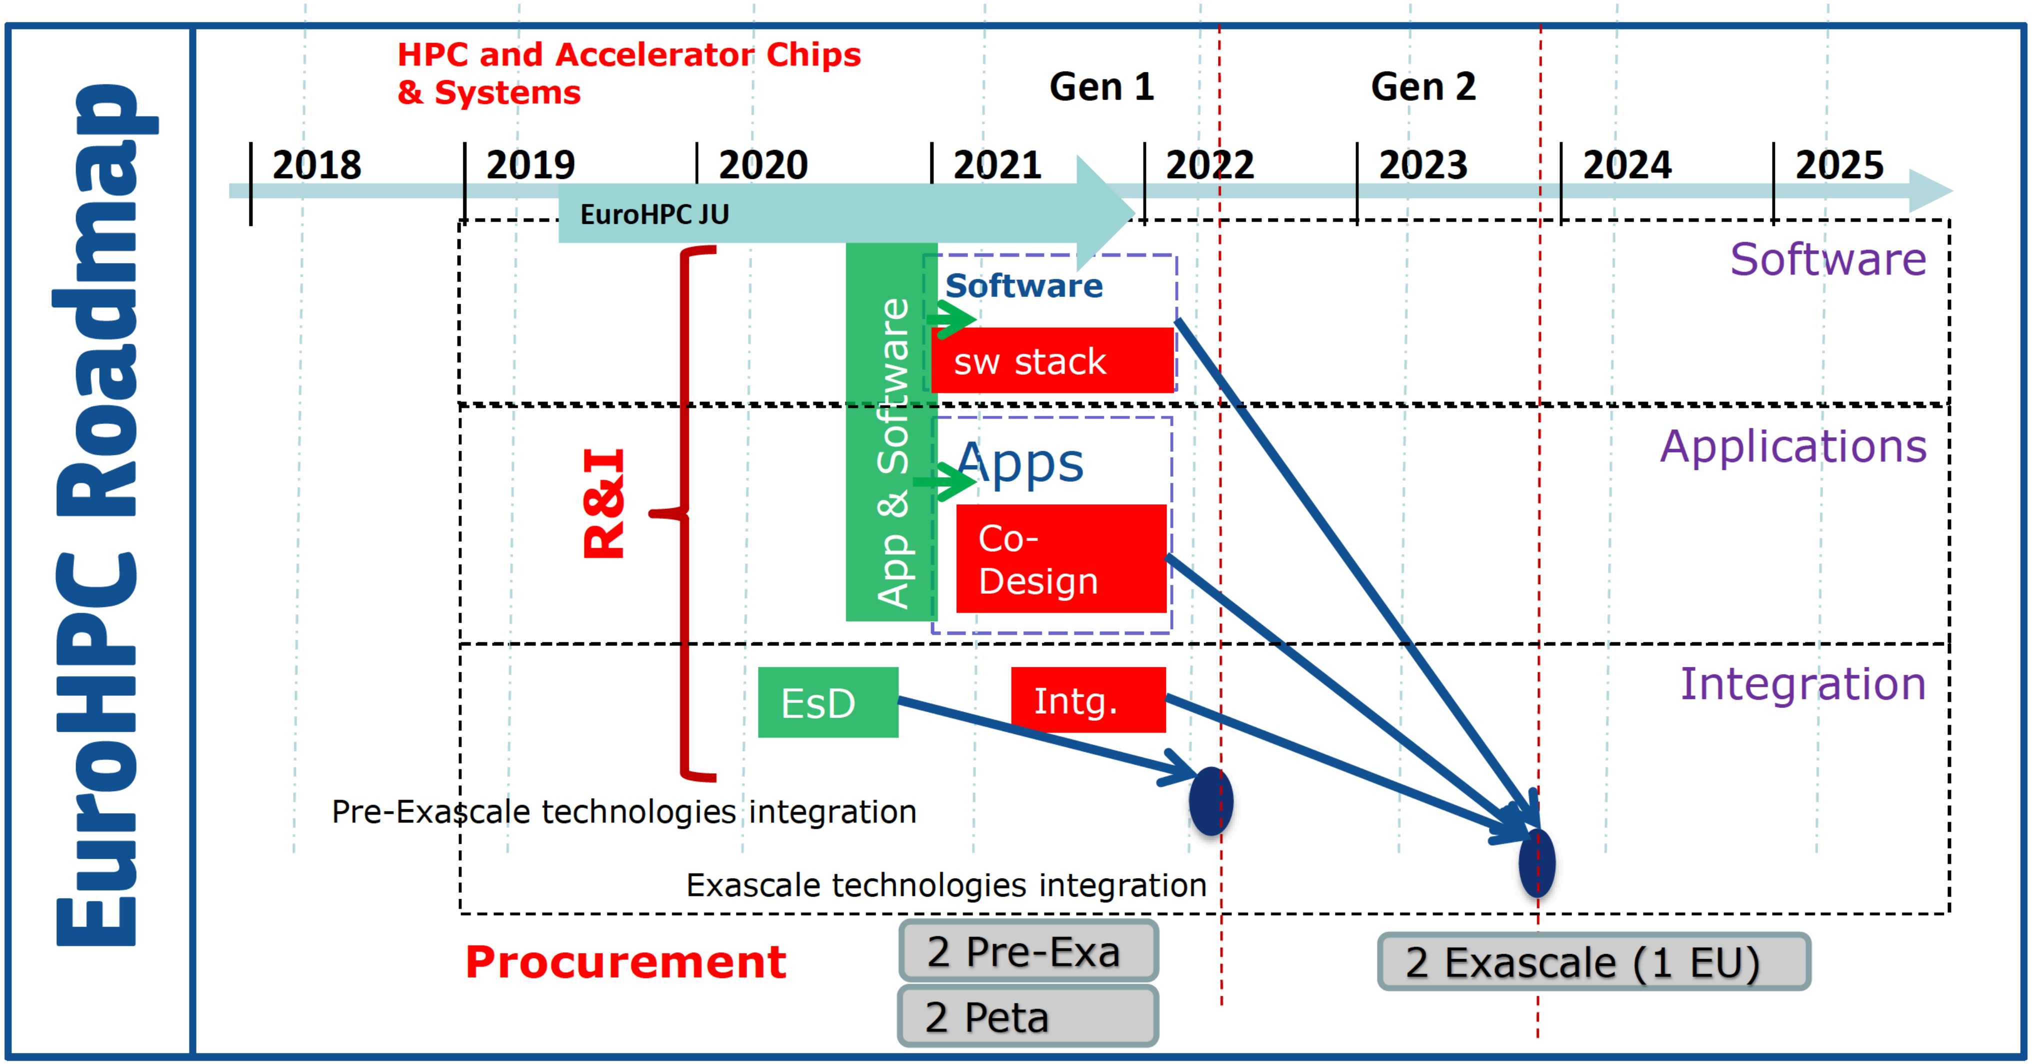
EuroHPC Roadmap
HPC and Accelerator Chips
& Systems
Gen 1
Gen 2
2018
2019
2020
2021
2022
2023
2024
2025
EuroHPC JU
Software
Applications
Integration
R&I
App & Software
Software
sw stack
Apps
Co-
Design
EsD
Intg.
Pre-Exascale technologies integration
Exascale technologies integration
Procurement
2 Pre-Exa
2 Peta
2 Exascale (1 EU)

EuroHPC Activities 2019-2020

HPC Ecosystem: Infrastructure & Operations	HPC Ecosystem: R&I, Applications & Skills	JU Admin/Running costs		
~270	min 180	10	486 m€	EC
~290	~186	10	486 m€	Participating States
560	392	20	972 m€	Total
0	~420 (in kind)	2	422 m€	Private Members

- **Infrastructure + Operations**

  Procurement of 2 pre-exascale machines and several (tbd) mid-range machines

- **Applications & Skills + R&I**

  R&I, exascale technologies and systems (incl. low-power processor); applications

- **JU Admin/running costs**
  - **JU Operation: 2019 to 2026**

*** Indicative figures!**

EPI

European Processor Initiative

Exascale Technical challenges

- Important hardware diversity/opportunity, extreme/multilevel parallelism, strong memory hierarchy,…
- Programming models , runtimes
- I/O , fault tolerance, online processing , visualisation
- Workflow , ensemble run
- Linear Algebra
- Algorithmic issue , numerical schemes
- Methodology : performance monitoring, continuous integration, performance engineering,…

Long term vision is a critical issue

- Applications are very long term efforts
- Need a holistic approach from the system to the application
- The software stack is can mitigate the strong technological diversity
- Find proper balance between complex community codes and *personal* tools