FP7 Support Action - European Exascale Software Initiative

DG Information Society and the unit e-Infrastructures

Addressing the Challenge of Exascale

Towards Exascale roadmap implementation

European Exascale Software Initiative

EESI

EESI2 – Vision & Recommendations

Philippe RICOUX (TOTAL) , EESI2 Coordinator

WELCOME

to

The European Exascale Software Initiative, EESI2

Final Conference

Dublin, 28 & 29 May 2015

EESI2 Initiative : Sept 2011 – June 2015

Why Exascale?

EXAFLOP (capability and capacity)

Oil & Gas Roadmap

HPC - Physics and Numerics

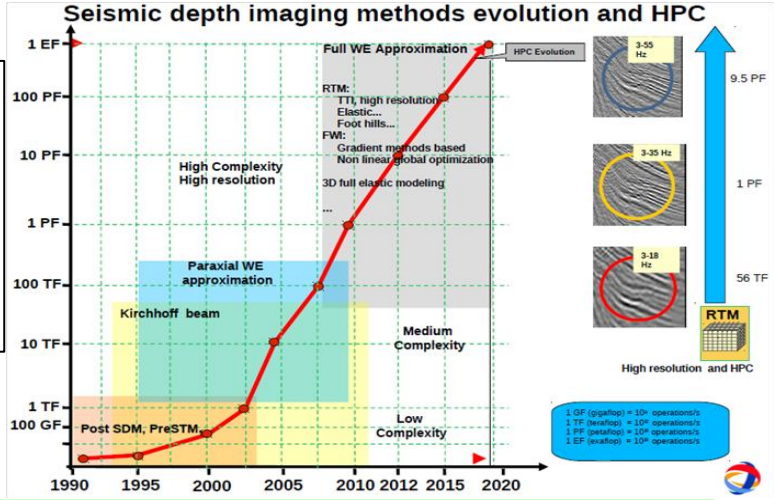

Drug design: realistic cell membrane models, including **drug permeation and binding**

Oil&Gas: huge 3D seismic wave inversion, reservoir modelling, robust optimization

Materials: material properties identification

Aeronautics-Greening the aircraft

Astronomy: **Square Kilometre Array**,

Earth Sciences: Natural hazard mitigation

Fundamental Sciences, Life Sciences,

Engineering sciences (Turbulence, Combustion, acoustics, Mechanical, chemical engineering, …)

… and/or **ExaBytes**

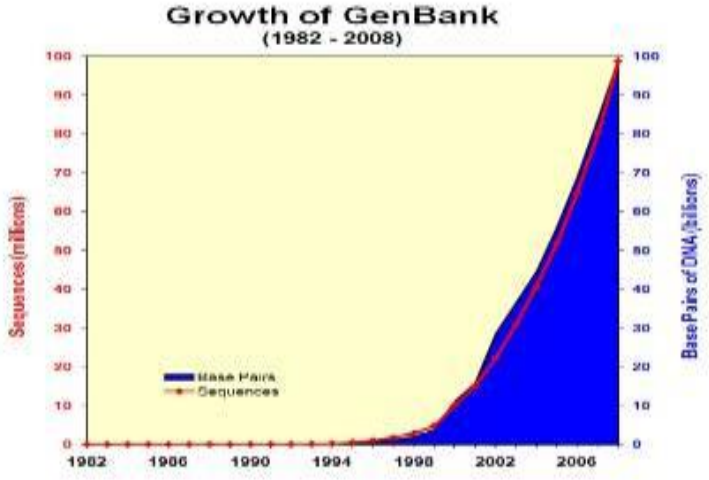

Climate: satellite sensors create floods of data (x 1000), leading to Exa-scale archives, ex. projected frequency of intense tropical cyclones in some region of the globe

Quantum Chemistry : discover a material and properties: ab initio databases of materials and molecular properties connected to existing databases of experimental properties

Industrial applications: management of data generated from micro/macro sensors and automated measurement devices

What is at stake?

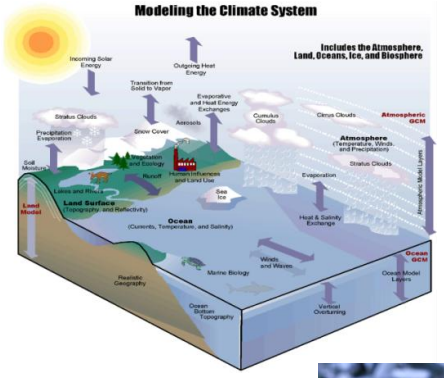

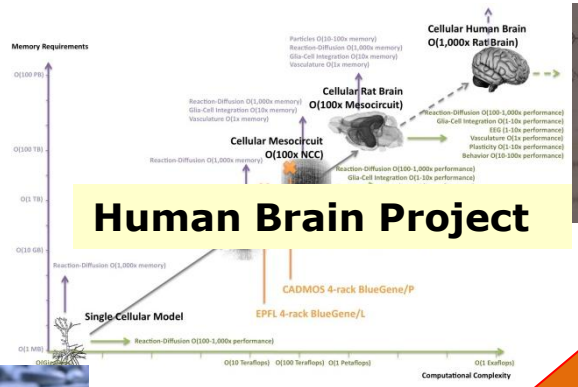

Human Brain Project

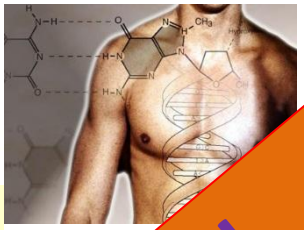

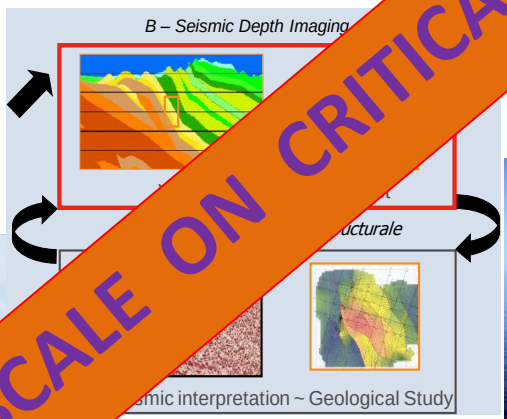

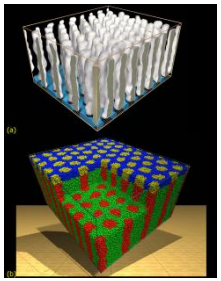

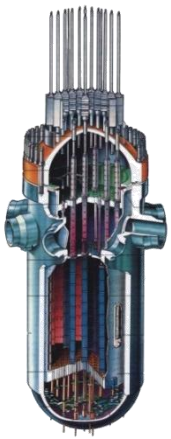

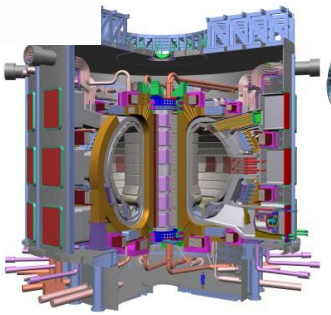

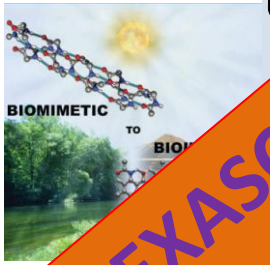

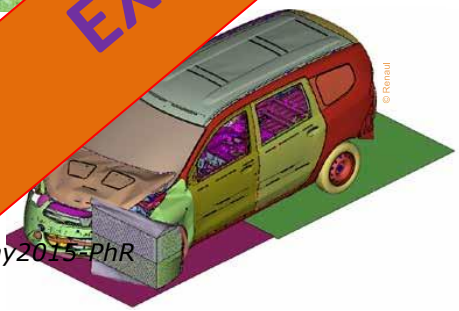

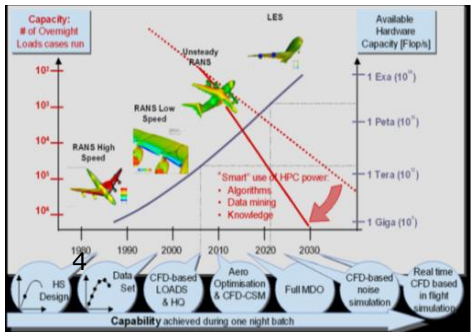

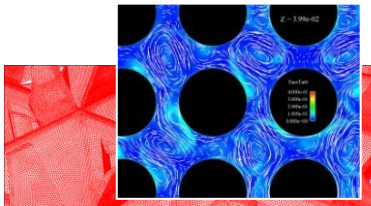

EXASCALE ON CRITICAL PATH

EESI Motivations and Organization

Origins of EESI: **IESP** (Jack Dongarra)

www.exascale.org

Build an international plan for developing the next generation open source software for scientific high-performance computing

*www.**eesi-project.eu***

EESI objectives

Build and consolidate a **vision and roadmap** at the European level, **including applications**, both from ***academia and industry*** to address the challenge of ***performing scientific computing*** on the new generation of super-computers, **hundreds of Petaflop/PBytes** in 2017 and **Exaflop/ExaBytes** in 2020/2022

Propose & Initiate International collaborations

In order to tackle the key issues

Toward EFFICIENT EXASCALE APPLICATIONS

European Strengths

- **PRACE, Infrastructure**
- **Applied Maths, Algorithms**
- **Applications (Astrophysics, fusion, bioinformatics, combustion, Life sciences, Molecular simulations, …),**
- **Industry (Energy, Aeronautics, …)**

European weaknesses

- **Simulation as a single program is not as strongly supported as observations and experiments (ITER, LHC), …**
- **Lack of coordination in development of HPC software, scientific libraries, Lack of data management (most of databases are non-European), Not enough participation in the definition of new standards for programming (MPI, OpenMP, C++, Fortran, …)**
- **Applications in various areas where European has a leadership position, will need to be redesigned, i.e Climate models, Life Science (molecular dynamics, Sequence Analysis, …), engineering structural mechanics**
- ***These applications represent X1000 man.year of development***

EESI1 vision and recommendations

EESI recommendations – general statement

- **Fund** strategic projects where Europe is strong and able to federate significant critical mass
- Insure the coordination of European efforts with the rest of the world, in particular:
  - on strategic areas where Europe lack of critical mass and,
  - in World wide initiatives in those that Europe has more value
- Reinforce Multi-disciplinary HPC Centers providing support in terms of code development, porting and optimization as well as algorithm development.
- co-design centers/centers of excellence should conciliates scientific multi-disciplinarity, international dimension, critical mass of researchers working at the same place, the balance of vertical (specialty) and horizontal (transverse) scientific domain.

Towards Exascale: Main issues to be addressed (EESI1)

- **At the level of simulation environment:**
  - **Unified Simulation Framework** and associated services: CAD, mesh generation, data setting tools, computational scheme editing aids, visualization, etc.
  - **Multi-physics simulation**: establishment of standard coupling interfaces and software tools, mixing legacy and new generation codes
  - common (jointly developed) **mesh-generation tool, automatic and adaptive** meshing, highly parallel
  - Standardized **efficient parallel IO and data management** (sorting memory for fast access, allocating new memory as needed in smaller chunks, treat parts of memory that are rarely/never needed based on heuristic algorithms, …)

50% Computer Power for Data movement Synchronization and Communication reducing algorithms

- **At the level of codes/applications:**
  - **New numerical methods**, algorithms, solvers/libraries, improved efficiency
  - coupling between stochastic and deterministic methods : Numerical scheme involving Stochastic HPC computing **for uncertainty** and risk quantification
  - meshless methods and **particle simulation**
  - **Scalable program, strong and weak scalability**, load balancing, fault-tolerance techniques, multi-level parallelism (issues identified with multi-core with reduced memory bandwidth per core , Collective communications, Efficient parallel IO)
  - ***Development of standards*** **programming models** (MPI, OpenMP, C++, Fortran, …) handling multi-level parallelism and heterogeneous architecture
- **Human resources, training (what level?)**

Exascale, a technological breakthrough

- Compare to Petascale computers: memory per core 1/10, CPU heterogeneity, total node interconnect BW & node memory 1/10, concurrency *10

  => concurrency/load balancing, data locality/Memory management, resilience/fault tolerance, energy efficiency
- Software layer and applications need to exploit these new hardware trends that cannot be handled by existing software stack
- Community codes unprepared for sea change in architecture while : designing and developing a new generation of Scientific Applications **takes 5 to 10 years,** lifetime of Scientific Applications are **several decades**

System attributes	2010	"2015"		"2018"		Difference Today & 2018
System peak	2 Pflop/s	200 Pflop/s		1 Eflop/sec		O(1000)
Power	6 MW	15 MW		~20 MW		
System memory	0.3 PB	5 PB		32-64 PB		O(100)
Node performance	125 GF	0.5 TF	7 TF	1 TF	10 TF	O(10) – O(100)
Node memory BW	25 GB/s	0.1 TB/sec	1 TB/sec	0.4 TB/sec	4 TB/sec	O(100)
Node concurrency	12	O(100)	O(1,000)	O(1,000)	O(10,000)	O(100) – O(1000)
Total Concurrency	225,000	$O(10^8)$		$O(10^9)$		O(10,000)
Total Node Interconnect BW	1.5 GB/s	20 GB/sec		200 GB/sec		O(100)
MTTI	days	O(1day)		O(1 day)		- O(10)

EXAFLOP also means a Petaflop in a box … **and 20 KW /PF!!!!!!**

Huge impact for those, academic, industrial, large and small structures, including SMEs, that will be **able to take advantage** of "Exascale" technology, not just for few heroes applications

From EESI to EESI2

EESI roadmaps, vision and recommendations ***need to be monitored, updated, on a dynamical way***

Key issues to be addressed are pointed out in EESI1 … Now EESI2 must recommend R&D CONCRETE Proposals how to tackle them

- **Extend, refine, and update** Exascale cartography (**directly in the dedicated WG for better analysis of each WG)** and **roadmaps** from HPC community, on software, tools, methods, R&D and industrial applications**, …**

new *With a Gap Analysis.*

new Including WG on ***disruptive technologies***

new **Address "Cross Cutting issues":** Data management and exploration, **Uncertainties - UQ&VQ**, Power & Performance, Resilience, Disruptive technologies

new **Investigation on** funding scheme and opportunities, education, co-design centres, international coordination

new Operational **Software maturity** level methodology, evaluation

EESI2 general picture

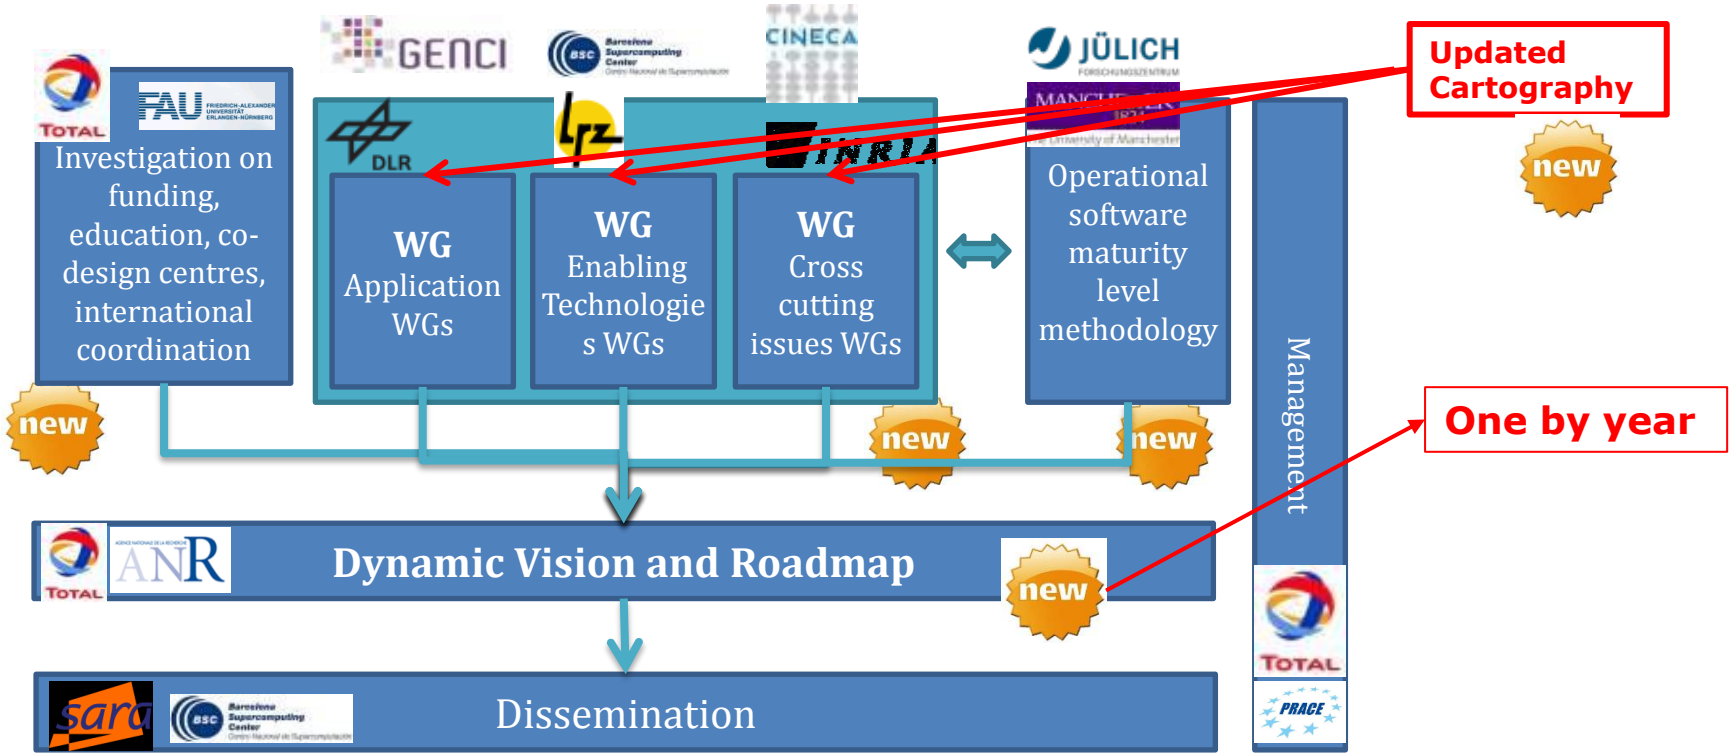

Contractual partners: TOTAL (coordinator), PRACE AISBL (acting for third parties LZR, GENCI, BSC, CINECA, EPCC, SARA…)

Contributing partners, involved in the management of EESI2 tasks but not associated to PRACE AISBL: INTEL, DLR, EDF, ANR, CERFACS, …

Supporting partners : more than 50 letters of Support

EESI2 proposal submitted in November to INFRA-2012-3.3: *Coordination actions, conferences and studies supporting policy development, including international cooperation, for e-Infrastructures*.

Requested funding: 1.5 M€ → 1.36 M€ accepted by EC

Duration: 30 months, Start 1st September 2012 - kick off 18th September 2012 (extended to 34 months)

EESI2 Working Groups

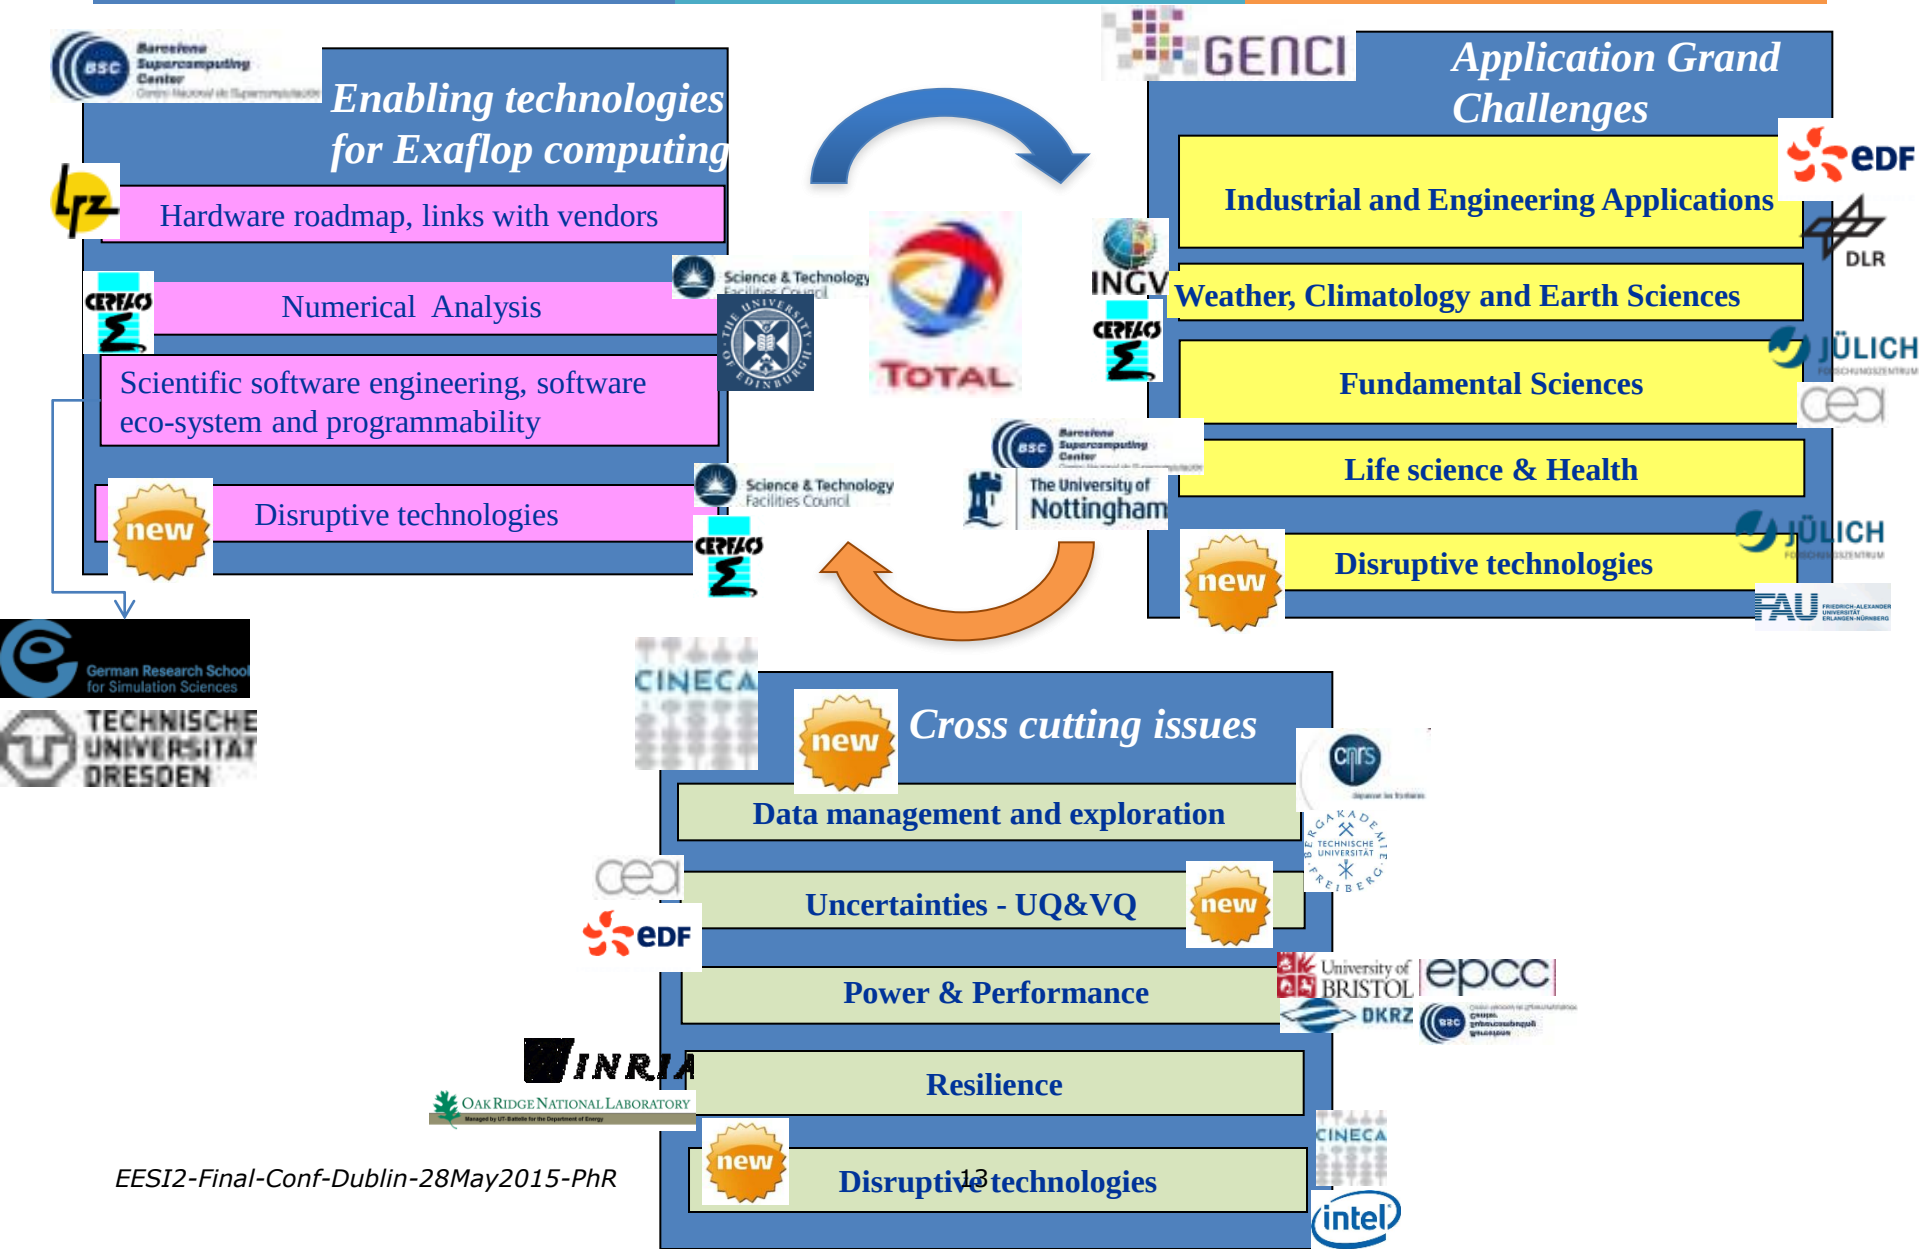

EESI: 150 experts, 100 entities

14 countries covering Europe including Russia ,
International Links, Participation of US, Japan, Israel, …
+ Independent Experts …

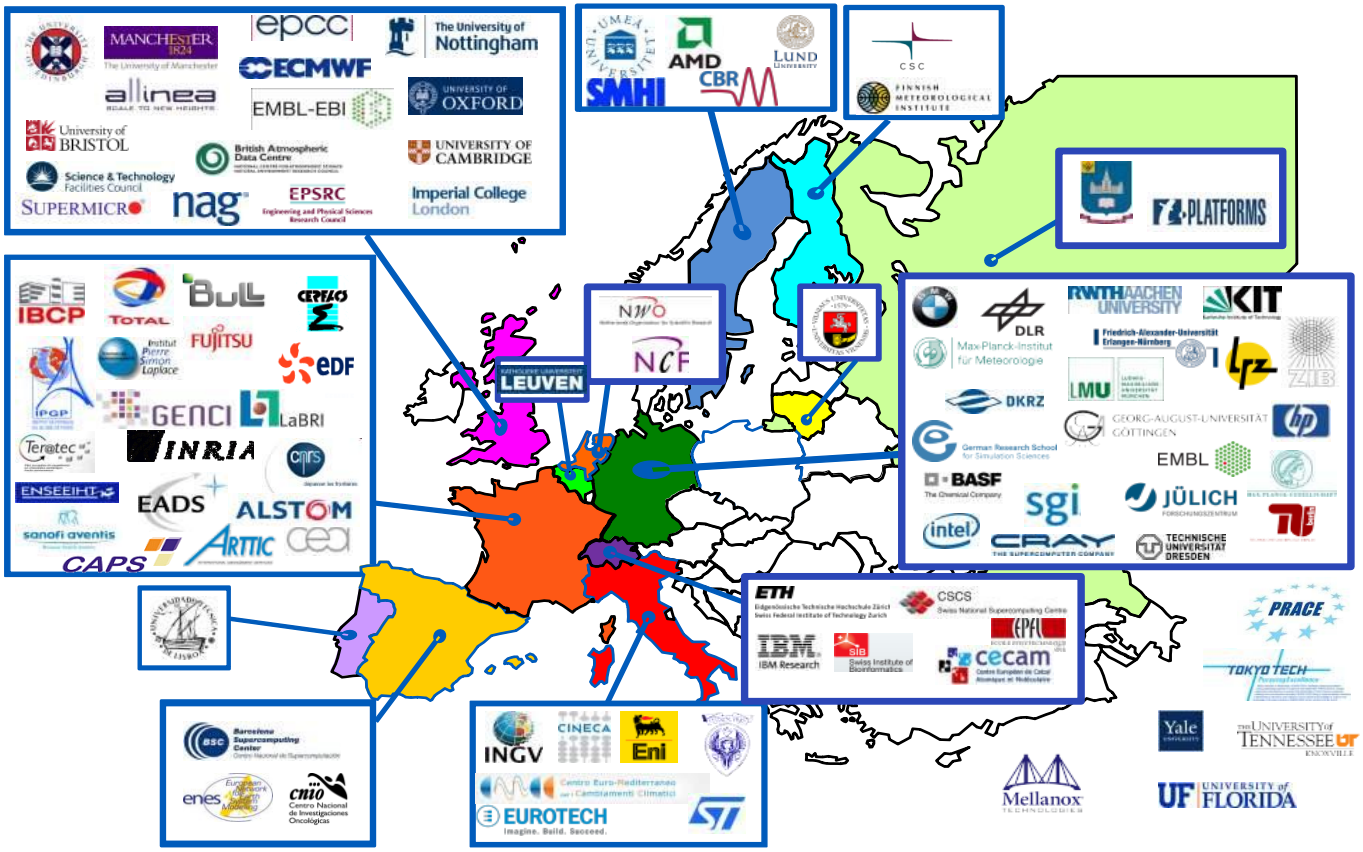

EESI2 FOCUS of Experts and WGs

- **Scientific key issues to be tackled**
  - **Fundamental, Industrial, Engineering, eco software, …**
  - **Cross Cutting**
    - **Big Data in extreme computing**
    - **Uncertainties (Quantification, propagation, ...)**
  - ***Not only challenges list but technical propositions (funding specific experts)!***

- **Detailed periodic roadmaps** on the key challenges, including:
  - **Gap analysis,**
  - ***Breakthroughs*,**
  - **Identification of priority actions,**
  - ***Recommendations* to R&D project**

- **Define educational programs for Exascale**

- **Contributing to build an international Exascale software Organization**

WGs, and Technical Meetings

Le Tremblay 2013

Bologna 2014

EESI2 VISION Summary

- **The EESI2 Vision**

  Learned from the existing "tens Petaflops" Computers , the feedback of several applications and tests running on full configurations of these systems. These tests have shown the *extreme difficulty to get some acceptable results* in term of performance on these computers. In particular the following points appear to be critical:

  - **Resilience**
  - **Error propagation**
  - **Reproductibility**
  - **Data transfert, communication**
  - **Task synchronization**

- **As a consequence, Exascale applicative software appear to be a very difficult challenge and most worldwide experts consider that this challenge will not be solved with existing algorithms**

- **In any case, the potential Exascale computing deployment is conditioning by the efficiency of scalable applications. Industries need a ROI as well in capacity as in capability.**

EESI2 VISION Summary

- **What appears presently, shared by US, Japanese and European experts, is that:**
  - ***Exascale technology will trickle down to every scale** (architecture system as well physics and time)*
  - ***Exascale cannot be justified only if we are just planning to do the usual thing but bigger***
  - ***Exascale machines will be useless without algorithms that use their specific features***
  - ***Extreme computing and Extreme Data should be tackled simultaneously***.
    - At Exascale, Extreme Computing and Extreme Data (or Big Data) are intrinsically linked since supercomputers become mandatory to analyze efficiently huge flows of data generated by large scale instruments or by massive complex simulations.
  - **Exascale applications** will be efficient only through developments by multidisciplinary teams, **optimizing the interactions between *architecture*** (nodes, cores, memories, interconnect, power, resilience ...), ***algorithms*** (programming, ultra scalable numerical methods, asynchrony, fault tolerance ...), **and *applications*** (discretization of problems, engineering tools, data processing ...).
  - ***Exascale imposes to do something different and differently***
  - **Exascale needs breakthroughs in several domains (Algorithms, Algebra, Uncertainties, Couplers, Meshing …)**

EESI2 VISION Summary

- **The following points are on the critical path to Exascale Computing::**
  - ***The use of hierarchical algorithms which reduce communications and tasks synchronizations***
  - ***The use of multi-physics methods which do not need or minimize data transfers and include multi scaling and parallel space-time methods***
  - ***The reshaping of operating systems and management tools such as MPI and OpenMP and mesh generation tools to the new developed algorithms***
  - ***Coupling Extreme Data and Extreme Computing, the use of in situ data processing,***
- ***It is urgent that the EU funds large project***s (CoE ?) focusing on complex specific Exascale challenges in particular in the domains where the EU has established strengths. Exascale 2020 (or 2022) is approaching rapidly and ***the Exascale issues are really challenging.***

EESI2 recommendations

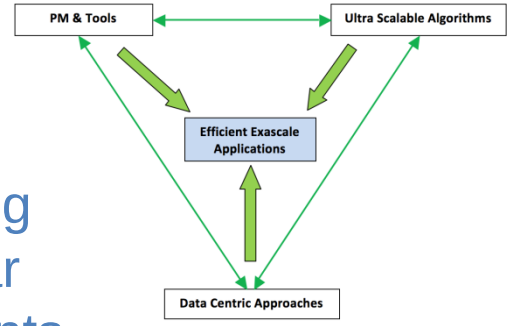

New thinking is required to develop new programming models, new algorithms, new tools, new data processing methods ... not only bigger than the present ones but far beyond the required innovation, **disruptive** improvements are needed

Extreme computing and **Extreme Data** should be tackled simultaneously

Assuring the **efficiency and productivity** of tools and applications at Exascale

EESI2 recommendations

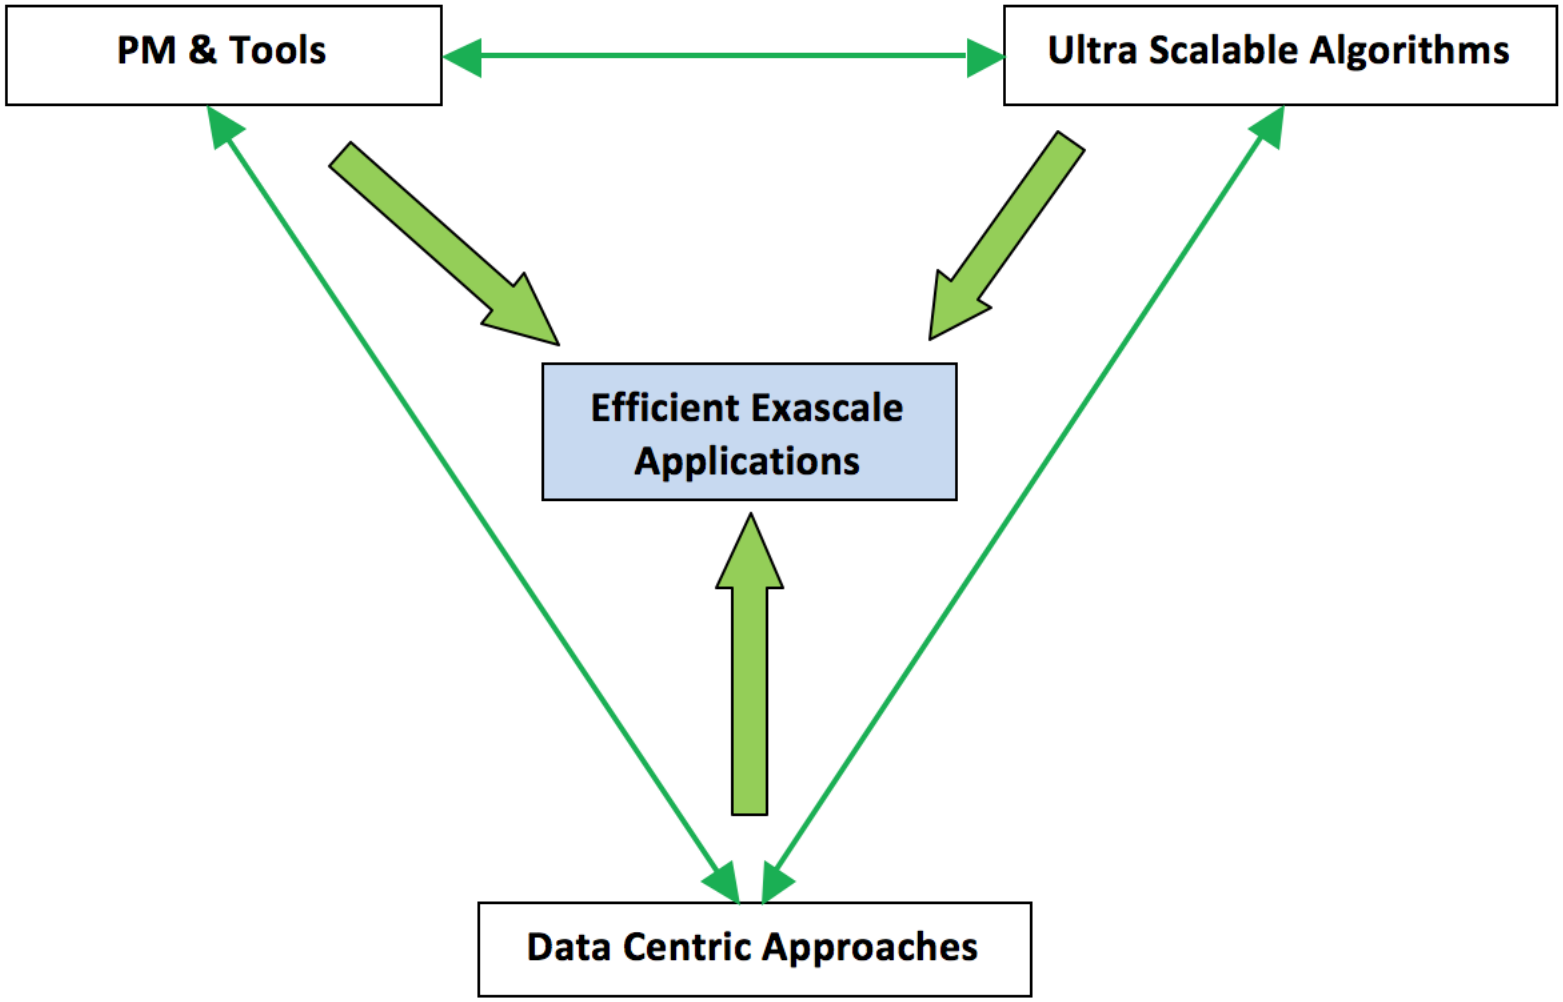

EESI Scientific Recommendations

- **Recommendations of the Tools & Programming Models pillar:**
  - High productivity programming models for Extreme Computing
  - Holistic approach for extreme heterogeneity management of Exascale supercomputers
  - Software Engineering Methods for High-Performance Computing
  - Holistic approach to resilience

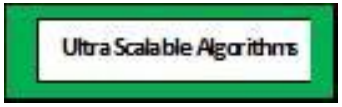

- **Recommendations of the Ultra Scalable Algorithms pillar:**
  - Algorithms for Communication and Data-Movement Avoidance
  - Parallel-in-Time: a fundamental step forward in Exascale Simulations (disruptive approach)

- **Recommendations of Data Centric Approaches pillar:**
  Vision "Software for Data Centric Approaches to Extreme Computing"
  - Towards flexible & efficient Exascale software couplers (direct or not, exchange of massive data)
  - In Situ Extreme Data Processing and better science through I/O avoidance in High-Performance Computing systems
  - Uncertainties Quantifications tools evolution for a for better exploitation of Exascale capacities
  - Declarative processing frameworks for big data analytics, extreme data fusion e.g. identification of turbulent flow features from massively parallel Exaflops and Exabytes simulations

EESI Info & Deliverables on Web Site

EESI vision and recommendations are also available for large public in readable documents and fliers on a new more "commercial" Web site (suggested by the project reviewers)

This new Site will allow to really and widely disseminated EESI works, vision, key deliverables …

but also information, events, news, research, programs, …

links with PRACE, ETP4HPC, …

links with International Community, …

All towards Exascale Implementation

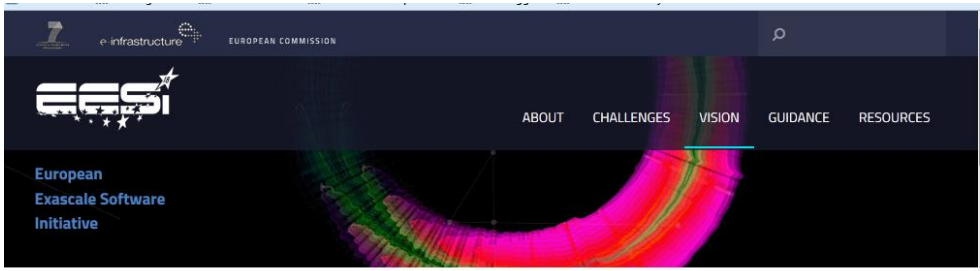

TOWARD EXASCALE COMPUTING ROADMAP FOUNDATION KEY R&D PROGRAMS

EESI2 vision: 3 pillars for Exascale

EESI2 has developed the periodic update of vision, roadmap and recommendations for an effective implementation of Exascale in Europe. This vision clearly recognizes the imperative need for breakthrough on hierarchical algorithms.

In 2013, EESI2 had published the first vision update. Eight key R&D programs for an exascale roadmap have been identified:

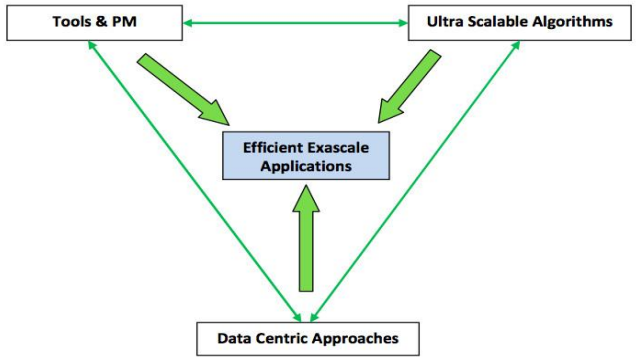

EESI Strategic & Funding Recommendations

Towards Exascale EU funding Budget must concern of course Infrastructure but also and largely the Software and Applications as pushed by EESI

It is a fact: Europe is in late vs USA (CoE, Examaths, In situ Data Processing, Resilience, …)

Exascale is Disruptive, So Europe must be disruptive in its strategy

Key R&D programs must be *specified to tackle Exascale issues* as described in the 3 EESI pillars, not mixed with classic HPC, and should be *funded at once* without waiting for 2017. Europe already lost 3 years!

Research critical mass must be assigned to each key program and each CoE during 10 years

Europe must develop a THINK TANK for Exascale (EESI … or equivalent)

For real competitiveness of European Industries and Research , Critical mass could lead to:

- **Specific Research and Development projects (Recommendations):** 20 projects; 3M€/year each during 10 years → **600 M€ over 10 years**
- **Center of Excellence program**: 5-7 CoE , 10M€/year each → **500 – 700 M€ over 10 Years**
- **Exascale Infrastructure platforms: 500 – 700 M€ for 2 to 3 platforms, over10 years**
- **Technological transfer** through a **European Exascale Software Center**: 20 M€/year → **200 M€ over 10 years**

Cumulative Global estimated budget : ≈ 2 Billions € over the next 10 years (Confirming EESI1)

Much Higher than EC plans!

Agenda of EESI Final Conference

Day 1:

- EESI works, vision and R&D programs recommendations towards efficient Exascale Applications
- Presentation of EESI key Recommendations pillar by pillar
- Presentation of Recommendation of Exascale Maturity Center

- Round Table on the 3 Pillars Programs

Day 2

- EESI vision on Education, Co-design center and International Collaboration
- Industrial accounts on expectations towards Exascale
- Vendors/Sponsors accounts

- Beyond EESI
  - EC Strategy
  - Next Initiative EXDCI

- Round Table on What next after EESI towards Exascale?

Agenda of EESI Final Conference

THANK YOU to ALL EXPERTS AND ACTIVE PARTICIPANTS TO EESI WORKS

All did A GREAT JOB in order to propose programs to TACKLE so COMPLEX CHALLENGES

THANK YOU for the "disruptive" contributions

EESI was and will be the real EU HPC THINK TANK

ENJOY THE PRESENTATIONS during this EESI Final Conference

HAVE A NICE and FRUITFUL TIME

THANK YOU