D7.3

EESI2 Final Report

2015 Update Vision & Recommendations

D7.3
2015 Update
Vision & Recommendations

CONTRACT NO EESI2 312478
INSTRUMENT CSA (Support and Collaborative Action)
THEMATIC INFRASTRUCTURE

Due date of deliverable:
Actual submission date: 30th September 2015
Publication date:
Start date of project: 1 September 2012 Duration: 34 months
Name of lead contractor for this deliverable: TOTAL SA, Philippe RICOUX

Revision: V1.0

Summary

The document presents the final EESI2 update, from European Exascale Software Initiative Experts, of the roadmap/vision and the recommendations for critical R&D challenges to be funded in order to develop efficient applications at Exascale computing.

This final update is very close to the one emitted end 2014 and could be resumed by

- Exascale architecture is a breakthrough, then Exascale applications should be developed with breakthroughs,
- Extreme Computing and Extreme Data are linked definitively. All algorithms for Exascale should take in account massively parallel programming but also data movement and data processing.
- Yes, efficient Exascale applications could be at the end of the tunnel, but this need to urgently define specific R&D program to tackle different key issues. There are ESSI2 recommendations.

These vision and recommendations are focused on the issues and challenges at Extreme Computing and Extreme Data, which are much more complex than classical HPC and which cannot be tackled only by pursuing the development of known HPC technologies and tools.

The roadmap towards the implementation of efficient Exascale applications and the consecutive recommendations are gathered in three large pillars:

- Tools & Programming Models
- Ultra Scalable Algorithms
- Data Centric Approaches

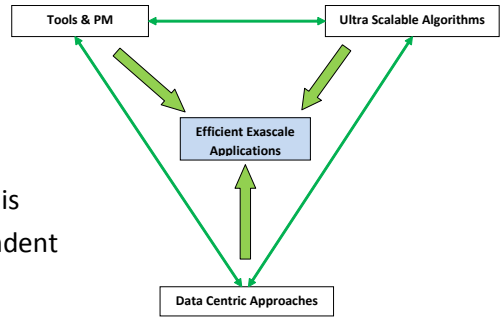

Note that the Data Centric vision is very new in Europe but is essential for approaching the ultra complex and interdependent challenges of Extreme Computing and Extreme Data.

As already advised by EESI2 (and by US DOE), Exascale requires a new and different approach compared to classical HPC. There is an urgent need for specific and disruptive R&D programs targeting Exascale software.

All EESI2 recommendations are aligned with this approach: they are coherent and aimed at assuring the efficiency of tools and applications at Exascale.

European Exascale Software Initiative Final Recommendations

Vision
Scientific and Strategic Recommendations

The EESI2 Initiative is clearly oriented toward the development and implementation of efficient Exascale applications, algorithms and software for enabling the emergence of a new generation of data intensive and extreme computing applications. The driver for this is the disruptive nature of Exascale computing with its potential for massive return on investment by addressing huge economic, societal and scientific challenges. New thinking is required to develop new programming models, new algorithms, new tools, new data processing methods ... not only bigger than the present ones, which will remain useless, but far beyond the required innovation. Exascale creates fundamental new opportunities, but it also brings fundamental new challenges.

The EESI vision is in coherence with International R&D programs funded on Exascale in particular in the US and Japan. Europe clearly has strengths (applications, scalable algorithms, couplers...) but also is clearly late on some Exascale key issues (languages, programming tools ...). EESI is very active in the International Initiative Big Data & Extreme Computing (BDEC) towards Exabytes and Exaflops from EU, US and Japan.

The principles underlying the recommendations are:

- Exascale is not only a "bigger HPC". There is an urgent need for new specific and disruptive ultra-scalable improvements in order to realize its full potential

- 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 ...).
- It is urgent that the EU funds large projects 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.

The EESI Vision of Exascale

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 th**e *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.

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 …)*

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.*

The EESI Scientific Recommendations

In this context, and clearly in accordance with the EESI roadmap towards Exascale, the experts of the EESI initiative have developed a set of coherent recommendations. Each recommendation proposes an initiative to tackle one of the key issues. As a whole the initiatives comprise a unique and coherent objective: to supply all necessary actions towards efficient Exascale software and applications.

The EESI recommendations can be grouped into three pillars: Tools and Programming Models including resilience, Ultra Scalable Algorithms, and Data Centric Approaches. Each EESI recommendation is detailed in the next section of this document.

Exascale comprises both Exaflops and Exabytes. Power efficiency, computing and overall performance are heavily dependent on data locality and data movement. Therefore, Data Centric Approaches to Extreme Computing should be in the core of the Exascale strategic roadmap.

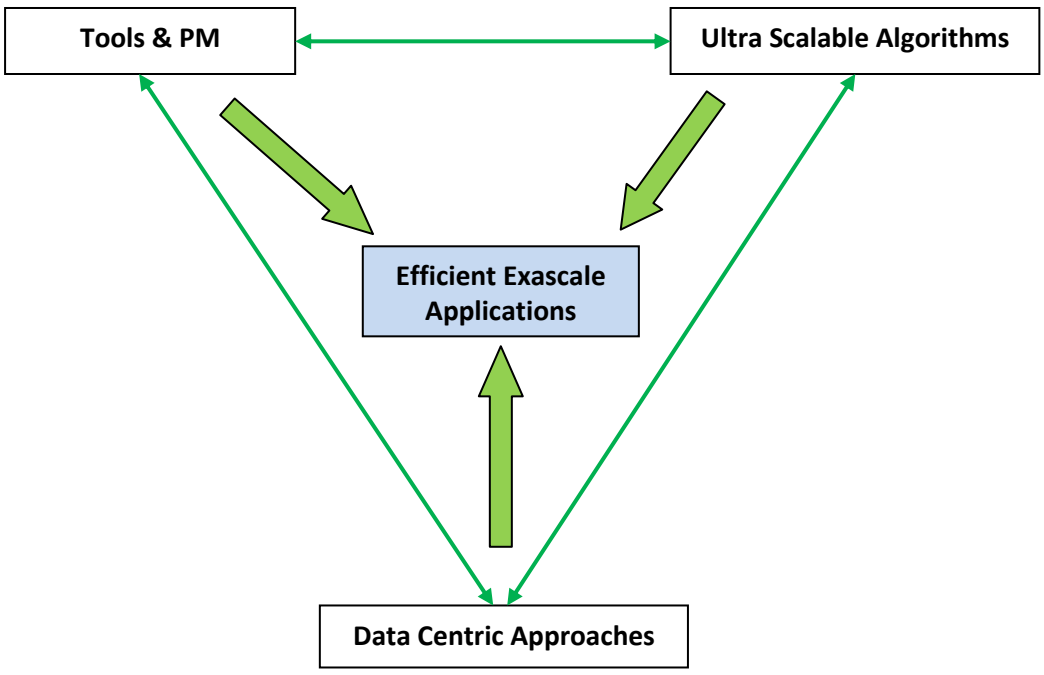

In the **Tools & Programming Models pillar**, recommendations concern programming models and methods, heterogeneity management, software engineering and cross-cutting issues like resilience, validation and uncertainty quantification with a strong focus on the specificity of Exascale in these domains.

The following recommendations are proposed for funding by the European Commission:

- 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
- Verification Validation and Uncertainties Quantifications tools evolution for a for better exploitation of Exascale capacities

In the **Ultra Scalable Algorithms pillar** recommendations concern specific and disruptive algorithms for Exascale computing, taking a step-change beyond "traditional" HPC. It will lead to the design and implementation of extremely efficient scalable solvers for a wide range of applications.

The following recommendations are proposed for funding by the European Commission:

- Algorithms for Communication and Data-Movement Avoidance
- Parallel-in-Time: a fundamental step forward in Exascale Simulations (disruptive approach)

The **Data Centric pillar** links Extreme Computing and Extreme Data. For the transition to Exascale, current data life cycle management techniques must be fully rethought, as described in the first joined document "Software for Data Centric Approaches to Extreme Computing" which is more a vision than a concrete recommendation.
This pillar gathers together key strategic issues for Exascale applications which are not enough addressed until now in Europe.

Ensuing from the EESI holistic vision of "Software for Data Centric Approaches to Extreme Computing", the following recommendations, all new at European level, should be supported and funded by European Commission:

- Towards flexible and efficient Exascale software couplers (direct or not, exchange of big data)
- In Situ Extreme Data Processing and better science through I/O avoidance in High-Performance Computing systems
- Declarative processing frameworks for big data analytics, extreme data fusion e.g. identification of turbulent flow features from massively parallel Exaflops and Exabytes simulations

Not all of these recommendations are at the same level of generalization but they are complementary and linked to each other by their global common objective: enabling the emergence of a new generation of intensive data and extreme computing applications. Some of them are fully disruptive; all need to go beyond known HPC technologies and methods.

All these recommendations should be supported and funded.

Some of these recommendations could be addressed in part by being strategic themes for new Centres of Excellence (CoEs).

Description of the EESI 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
- Verification Validation and Uncertainties Quantifications tools evolution for a for better exploitation of Exascale capacities

Recommendations of the Ultra Scalable Algorithms pillar:

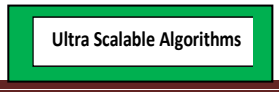

- Algorithms for Communication and Data-Movement Avoidance
- Parallel-in-Time: a fundamental step forward in Exascale Simulations (disruptive approach)

Recommendations of Data Centric Approaches pillar:

Vision "Software for Data Centric Approaches to Extreme Computing"

- Towards flexible and efficient Exascale software couplers (direct or not, exchange of big data)
- In Situ Extreme Data Processing and better science through I/O avoidance in High-Performance Computing systems
- Declarative processing frameworks for big data analytics, extreme data fusion e.g. identification of turbulent flow features from massively parallel Exaflops and Exabytes simulations

Impact of the EESI Recommendations

Here below, as a consequence of EESI recommendations, an extract of European Call for 2016-2017:

- ***Co-design*** *of HPC systems and applications (big projects)*
- ***Transition to exascale computing*** *(smaller focused projects)*
  - **High productivity programming environments for exascale**
  - **Exascale system software and management**
  - **Exascale I/O and storage in the presence of multiple tiers of data storage:**
  - **Supercomputing for Extreme Data and emerging HPC use modes**
  - **Mathematics and algorithms for extreme scale HPC systems and applications working with extreme data**
- *Exascale* ***Ecosystem Development*** *(CSAs)*

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)