

Deliverable D8.4

Report on second technical workshop

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

Due date of deliverable: 31 September 2014

Actual submission date: 31 September 2014

Publication date: 31 September 2014

Start date of project: 1 September 2012

Duration: 30 months

Name of lead contractor for this deliverable: PRACE-CINECA Giovanni Erbacci

Name of reviewers for this deliverable:

Abstract: This document describes the organisation and the outcomes of the second technical workshop held in Bologna-(Italy) on the 3rd and 4th of June, 2014. About 40 people attended this workshop, which was aimed to present each working group results and roadmaps.

Project co-funded by the European Commission within the Seventh Framework Programme (FP7/2007-2013)		
Dissemination Level		
PU	Public	X
PP	Restricted to other programme participants (including the Commission Services)	
RE	Restricted to a group specified by the consortium (including the Commission Services)	
CO	Confidential, only for members of the consortium (including the Commission Services)	

Table of Contents

1. EXECUTIVE SUMMARY	3
2. INTRODUCTION.....	4
3. SETTING AND AGENDA	5
4. REPORT ON DAY 1	7
4.1 INTRODUCTION: GOAL & EXPECTED OUTPUT OF THE MEETING	8
4.2 PRESENTATIONS OF THE WORK PACKAGE 3: WORKS OF WGs, FIRST R&D RECOMMENDATIONS	8
4.3 PRESENTATIONS OF THE WORK PACKAGE 4: WORKS OF WGs, FIRST R&D RECOMMENDATIONS ...	11
4.4 PRESENTATIONS OF THE WORK PACKAGE 5: WORKS OF WGs, FIRST R&D RECOMMENDATIONS ...	14
4.5 EESI2 IN THE CONTEXT OF EC DG-CONNECT, EXASCALE AND HPC EC STRATEGY	16
4.6 PRESENTATION OF IDC HPC STUDY	17
4.7 SYNTHESIS OF THE PARALLEL SESSIONS ON DISRUPTIVE TECHNOLOGIES.....	17
5. REPORT ON DAY 2	20
5.1 PRESENTATIONS OF THE WORK PACKAGE 2: WORKS, FIRST R&D RECOMMENDATIONS	20
5.2 PRESENTATIONS OF THE WORK PACKAGE 6: WORKS, FIRST R&D RECOMMENDATIONS	21
5.3 PRESENTATIONS OF THE WORK PACKAGE 8	21
5.4 SYNTHESIS OF THE PARALLEL SESSIONS ON R&D PROGRAMS RECOMMENDATIONS	22
6. CONCLUSION OF THE MEETING	24
7. ANNEX – WORKSHOP PARTICIPANTS	25

Glossary

Abbreviation / acronym	Description
COTS	Commodity-of-the-shelf
DoE	Department of Energy (US)
EC	European Commission
EESI	European Exascale Software Initiative (Europe)
ESC	Exascale Software Centre (US)
HPC	High Performance Computing
HPCI	High Performance Computing Infrastructure (Japan)
IDC	International Data Corporation
IESP	International Exascale Software Project
IPR	Intellectual Property Rights
JSC	Juelich Supercomputing Centre
NCF	National Computing Facilities Foundation (the Netherlands)
NNSA	National Nuclear Security Administration (US)
NSF	National Science Foundation (US)
WG	Working Group (in EESI)
X-stack	Exascale software stack

1. Executive summary

This deliverable reports on the second EESI2 internal workshop which took place in Bologna (Italy), on June 3 and 4, 2014. This two-day meeting has brought together about 40 experts in the areas of software development, performance analysis, applications knowledge, funding models and governance aspects in High Performance Computing, coming from academia and industry in several European countries.

The Internal workshop was aimed for each working group (WG) to present the updating of the results and roadmaps addressed during the first year of activity and, generally speaking, to address cross cutting issues and disruptive technologies between the different fields. The aim of the meeting was also to elaborate the second set of concrete recommendations for the future of HPC in Europe.

The first day, on June 3, was mainly dedicated to the presentations of the work packages 3, 4 and 5 which are at the heart of the project. It was the occasion for the work packages leaders to present a synthetic vision of the S&T perimeter of their work: hurdles / cross cutting issues, questions to the others. The first day was also the occasion for the DG-CONNECT to present the EC strategy for Exascale and HPC in Horizon 2020, and for IDC to present the study on HPC Return on Investment in Europe. Moreover, one global session on disruptive technologies was organised with all the experts, to brainstorm on this hot topic and identify common methodologies. The global session was followed by two specific parallel sessions which better focused on specific disruptive technologies specific arguments.

The second day, on June 4, started with two parallel sessions on “Numerics and Applications” and “Enabling Technologies” with the objective to improve and define new R&D program recommendations. Then the workshop continued with the presentations of the work packages 2 and 6. Two more parallel sessions were held on the organisation of the BDEC event (January 2015) and the coordination of the European Applied Mathematics community, to address complex themes toward Exascale. Finally, after a synthesis of the two sessions, the agenda for the next steps was discussed. It concerned the expectations for the deliverables and the next milestones.

The main results of the internal workshop are:

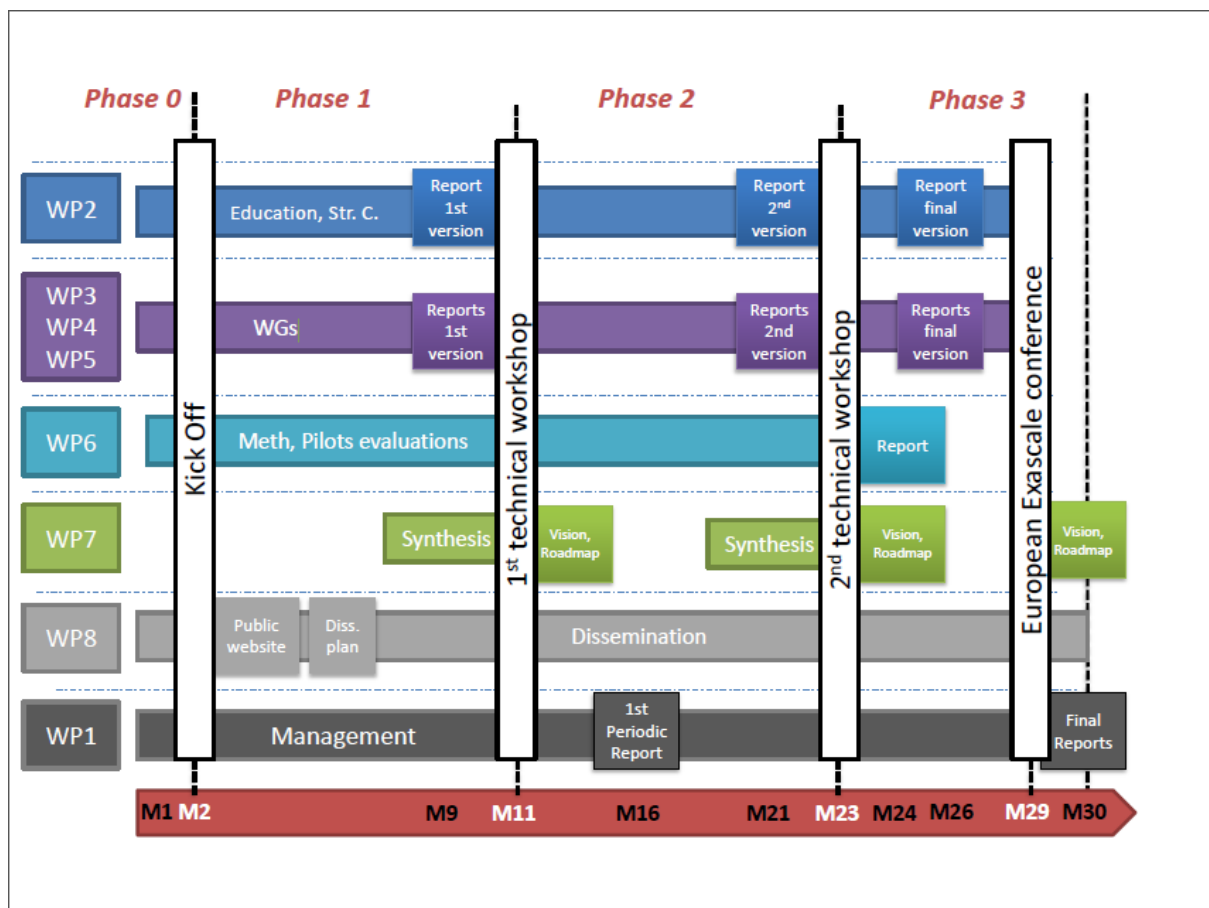
- A very good attendance of all work packages chairs and vice chairs as well as experts from the working groups;
- A presentation of the results of the working groups showing a good overall progress;
- Exchanges between working groups and their respective work package leader and the project leader about missing points in the report to address during the summer for the deliverables expected;
- A set of parallel sessions which will feed the deliverables and guide final thoughts;
- Fruitful discussions with P. Tsarchopoulos from DG-CONNECT.

2. Introduction

It is widely recognized that High Performance Computing (HPC) will be increasingly important to address global scientific, societal and economic challenges. Although the projected evolution of hardware is a technological challenge in itself, more and more concern is expressed on the ability of scientific software to efficiently use the future hardware architectures.

EESI is the European Exascale Software Initiative and is meant as a collaborative project which is expected to coordinate the European efforts in exascale challenges. EESI2 is the step following EESI. The objective of EESI2 is to build on the work done within EESI1 and to extend this role of external and independent representative of the European Exascale community.

The agenda of EESI 2 is divided in 3 main phases, the phase 1 from month 2 to month 11, the phase 2 from month 11 to month 23 and the phase 3 from month 23 to month 30.



3. Setting and Agenda

The Internal workshop was organized in Bologna (Italy) on June 3 and 4, 2014 at the location “*I Portici Hotel*” (see Figure 1 below), in the centre of the city but not far from the Cineca Supercomputing Centre, in charge for organising the workshop.

For the registration to the workshop a web page was prepared at:

http://www.eventbrite.it/e/eesi-2nd-technical-meeting-registration-11328863937?utm_campaign=new_eventv2&utm_medium=email&utm_source=eb_email&utm_term=eventurl_text

The number of participants to the meeting was around 40. The names and affiliations of the participants have been listed In Appendix A,.



Figure 1: General overview of the workshop location

Program of the Workshop

The program of the workshop is reported below. The overall schedule has been respected, only with some modifications on the order of the speakers during the two days meeting.



EESI 2- 2nd Technical Workshop Bologna, 3-4 June 2014

Day 1, June 3 2014	
08:40	Registration & Welcome
09:00	Introduction: goal & expected results of the meeting, EESI2 visibility in EU Philippe Ricoux
09:30	WP3 WGs activity: New results and R&D recommendations Applications / Jean Claude André Jean Claude André - Italo Epicoco – Godehard Sutmann
11:00	Coffee break
11:30	WP4 WGs activity: New results and R&D recommendations Enabling Technologies / Rosa Badia Iain Duff – Felix Wolf – Michael Ott
13:00	Lunch break
14:00	Horizon 2020: Exascale and HPC EC strategy Panagiotis Tsarchopoulos EC Project Officer
14:30	WP5 WGs activity: New results and R&D recommendations Cross Cutting issues / Giovanni Erbacci François Bodin – Anne Laure Popelin – Simon McIntosh Smith – Franck Cappello
16:00	Coffee break
16:30	Presentation of the IDC Study Joseph Earl, IDC
17:00	WP3, WP4 and WP5 WGs Disruptive Technologies: <i>Brainstorming and common methodology</i> <i>One global session with all experts and task leaders followed, if need, by 2 or 3 parallel sessions</i> Uli Ruede – Rosa Badia – Carlo Cavazzoni
18:30	Guided tour in the centre of Bologna
20:30	Social Dinner at the Hotel I Portici

Day 2, June 4, 2014	
08:30	R&D Programs Recommendations : Previous and Next Philippe Ricoux
08:50	Synthesis of the Disruptive technologies session
09:20	Two parallel sessions on improved and new R&D Program Recommendations: - <i>Numerics and Applications</i> - <i>Enabling technologies</i> All WP Leaders and Experts
11:00	Coffee break
11:30	Presentations WP 2: Education - Center of excellence - International collaboration Uli Ruede /Jean Yves Berthou
12:00	Presentations WP 6: Codes Maturity Bernd Mohr
12.30	Synthesis of the two parallel sessions on R&D Programs Recommendations
13:00	Lunch break
14:00	Two parallel sessions - BDEC Preparation (January 2014) - EESI2 of Scientific Workshops (European Applied Maths community coordination to share complexity / parallelisation of time) All WP & WG Leaders
15:00	Synthesis of the parallel sessions
15:30	Agenda for the next steps: • Synthesis report, expectations, templates • EC view and expectation Philippe Ricoux / Thierry Bidot
16:00	Conclusion and closure of the meeting

4. Report on Day 1

4.1 Introduction: goal & expected results, EESI 2 Visibility in EU

After the registration the workshop has been introduced by Philippe Ricoux (Total), the project leader of EESI 2. The second technical meeting of the project, this event was a unique opportunity for the different experts and WP leaders to meet and exchange on the objectives of the project.

During the introduction, some main issues to be addressed towards Exascale were recalled:

- Unified Simulation Framework
- Multi-physics simulation
- Mesh-generation tool, automatic and adaptive meshing
- Standardized efficient parallel IO and data management
- New numerical methods
- 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
- Development of standards programming models
- Human resources, training

This meeting was a unique opportunity to give strong guidelines for paving the path of EESI2 activities.

Finally, the objectives and expectations of this meeting, in terms of organisation were detailed. Deadlines were set-up for the deliverables expected by the European Commission.

After the introduction, the presentations of the working groups started.

4.2 Presentations of the Work Package 3: works of WGs, first R&D recommendations

1. WG 3.1 – Industrial and engineering applications

This group includes already 10 experts from academia as well as industry (5). It is presently focused in Aeronautics, Energy and automotive and includes representatives from ISVs (NAFEMS, Dassault Systemes).

Most EESI1 Industrial roadmaps are still valid. In term of industrial needs Exascale must provide both capability with already visible needs not only for Exascale but also for Zetascale, and capacity including multi-physics, multi-scales, multi-disciplines.

Few unified simulation framework are available and there is a lot to do on coupling altogether elements such as CAD, mesh generation, data setting tools, computational scheme editing aids, visualization, data management, etc.

Training is also seen as a critical issue.

A collective brainstorming with others working groups will be done to identify disruptive technologies. A first idea is the use of composite instead of steel in the automotive industry.

2. WG 3.2 – Weather, climatology and earth sciences

The WG3.2 on “Weather, Climate and Earth Sciences” (WCES) reported about the following items:

- Membership and meetings: membership has been slightly completed. Two electronic meetings have been organized (April 17, May 16). The WG benefits from the work done by ENES (European Network for Earth-System Modelling) HPC Task Force, which already organized two workshops (Lecce, December 2011, and Toulouse, January 2013) and a number of conference calls (5 in 2012, 3 to date in 2013).
- Report on Disruptive technologies: the report has been sent in due time to WG 3.5. It includes suggestions under 3 categories (Algorithms, Software and Hardware) and underlines that trends to massively parallel systems with relatively poor data bandwidth would be very negatively disruptive!
- Report on Data: this topic is still under discussion within WCES. A questionnaire has been prepared and sent to the experts. Replies are due by June 10.
- Future work on Modeling: this topic is on the agenda for future WCES work. It will benefit from previous analysis done during the two earlier IS-ENES workshops (see above).
- Discussion of Centre of Excellence for Climate: the issue is still in a brainstorming phase.

3. WG 3.3 – Fundamental Sciences

The working group covers a broad range of scientific domains, comprising nuclear physics, laser-plasma physics, nuclear fusion, quantum chemistry, soft matter physics, materials science and astrophysics/cosmology, thereby integrating a wide range of length- and time-scales which are partly connected via multi-scale approaches.

During the presentation for the first technical workshop, an update from EESI-1 was given which contained new developments in software, performance and scalability issues, national and international initiatives and actions including community organizations.

Indicative for performance considerations is the Gordon Bell prize, which is awarded once a year for best performances during the Supercomputing Conference.

In 2011 and 2012 codes from materials science and cosmology, running on the K-computer and Tsubame in Japan as well as at LLNL in US attracted high attention.

In 2012, cosmology could be identified as a driver for high performance computing, as it was awarded in two categories for a large simulation as well as demonstration of large scale data-analysis, which shows that not only production but also analysis of data is now a central topic in supercomputing.

From the European side also a paper from cosmology was highly ranked demonstrating for the first time a simulation of the complete observable universe. An important development for Europe is seen with the acceptance of the Graphene FET flagship programme, which is topically covered by the WG. Although there is no direct link to Exascale initiatives, it is foreseen that it might also act as a European driver in extreme computing and first demonstrations of highest scalability on the European Tier-0 architecture Curie are reported.

International initiatives including participation of EU, US and Japan are seen in the G8 projects, where one of the six projects is addressing Nuclear Fusion, demonstrating porting, optimization and highest scalability for a number of European codes. Further initiatives are reported from community organizations, e.g. CECAM, where a strong interest exists in fostering and strengthening software development for complex simulation- and analysis-codes in material science, soft matter and biophysics, including scale-bridging methods. National initiatives in Germany are promoting high

performance computing in special calls, trying to form a German HPC community or linking HPC with the industrial sector.

From discussions within the working group, a strong demand is seen for higher funding of software development and a recommendation is given to devote about 20% of budget, invested into high performance architectures to software development. A strong motivation for this suggestion is based on the observation that wall clock time reduction for solving a given problem might be substantially larger than a further machine upgrade to higher nominal peak performance, if an optimised software implementation is available.

Finally it is recommended that software for system and applications should be provided within the Open Innovation concept, which provides a more efficient way to exchange and integrate software modules and libraries, taking into account license issues, reducing time-to-solution of optimized software development and minimizing the duplication of implementations.

4. WG 3.4 – Life sciences & health

The WG identified 4 main areas in Life Sciences that will require Exascale computing: Genomics, Systems Biology, Molecular Simulation and Organ simulation.

Cartography update:

- Genomics: We presented the challenges reported in EESI-1 (1000 Genomes, International Cancer Genomics Consortium) and new ones: the Encode project (30 papers with 6 natures), that opens an unmet level of annotation to human genome and the Metagenomics project (Metahit) that aims to analyse genomes of human gut organisms;
- Systems biology: One of the most relevant advances in the fields is the publication of a whole-cell model (cell 2012);
- Molecular Dynamics: Long simulation times are one of the main challenges of Molecular Simulation. New advances in development of data frameworks and environments for automatic preparation and launch of simulations on HPC;
- Organ simulation: The Human Brain Project is granted with a FET-Flagship and the Virtual Physiological Human consortium is very active connecting various FP7 projects.

Recommendations:

- Creation of a co-development centre to set up a platform for Exascale applications for Life Sciences

A face-to-face meeting of Life Sciences panel in Barcelona, September 2013 has been announced.

5. WG 3.5 – Disruptive Technologies

The working group is concerned with disruptive approaches in applications, which raises several questions and problems that have to be addressed first.

To form a working group it has to be assured that it represents in a way the topic and that it might be relevant to Exascale computing. A main problem is that there is (by definition of the word “disruptive”), no dedicated community and consequently no community organization which can be addressed. Indeed, disruptive approaches, developments or trends must be identified among a number of different communities. The topics exhibit an inherently interdisciplinary and multi-disciplinary character which includes e.g. a variety of length- and time-scales and diversity of different methods, e.g. stochastic versus deterministic, lattice versus mesh-less methods, optimization, dynamic versus sampling methods including multi scale methodologies.

The key element of a disruptive approach might nevertheless be found on a common ground. A disruptive approach puts questions on traditional methods and schemes and proposes an alternative

approach or even completely new topics. By the “definition” of disruptiveness, a new approach is not yet generally expected to have full success nor is it fully accepted in the community. Nevertheless it should contain ingredients which might promise a break-through in case it turns out to be successful. In that respect, disruptive approaches can still be considered as risky.

To receive a basis for the working group, a survey was started among different communities, which tried to identify both disruptive methods and disruptive approaches in applications. The latter may be based on “traditional” methods, but try to explore new ways in modeling or combination of methods in e.g. a hierarchical way. Given the limited number of responses covering a diverse set of communities, a first impression about possible directions was obtained. General comments and answers from the survey addressed topics like web-based computer environments, application of ab initio approaches (i.e. not invoking a model) for a variety of applications, adaptability in accuracy and model selection or parallelization in time, to name a few.

Along the result from the survey as well as from discussions within EESI and other communities, the selection of experts will be pursued to provide an overview of activities and to come up with recommendations for possible actions.

4.3 Presentations of the Work Package 4: works of WGs, first R&D recommendations

1. WG 4.1 – Numerical analysis

This group includes 13 experts most of them were already involved in EESI1. The group met in Edinburgh on April 4th.

The field of investigation have been broken down into the following sections:

- Dense linear algebra
- Tensors
- Graph and hypergraph partitioning, mapping, repartitioning, remapping
- Sparse direct methods
- Iterative methods for solving sparse linear systems of equations
- Eigenvalues problems, model order reduction, matrix equations and around
- Optimization, Control, Automatic Differentiation
- Fast Multipole Methods
- Structured and unstructured grid calculations

As memory accesses are increasingly the bottleneck in computations algorithms need to maximise the number of useful calculations per memory access. This is frequently achieved by blocking/tiling and communication hiding. Load balancing issues mean that synchronisation points are expensive and asynchronous versions of existing algorithms need to be investigated. Often such algorithms have been suggested in the past but have been unfavoured due to stability issues that are difficult to understand and control. It is time to give these algorithms a fresh look. Dynamic scheduling based on DAG representations of algorithms is better suited to exploit all possible concurrency in computations. To make full use of its potential and to avoid unnecessary synchronisation points the traditional hierarchy of numerical library calls (matrix-vector multiplication called by an iterative solver called after a graph partitioning routine - which by design lead to a fork-join execution model) should be broken down. Other common issues that will be increasingly important are the trade-off between accuracy and reproducibility, fault tolerance and uncertainty quantification. Naturally progress towards dealing with these challenges is more advanced in some areas than in others.

2. WG 4.2 – Scientific software engineering, software eco-system and programmability

In an adaptation of ISO/IEC/IEEE 24765:2010, we define “[scientific] software engineering as the application of a systematic, disciplined, quantifiable approach to the development, operation, and maintenance of software; that is, the application of engineering to [scientific] software”.

This working group focus on methods, processes, tools, and support structures required to create robust, correct, efficient, and maintainable code under economic constraints.

The target software includes mostly highly scalable simulation codes but also other data-intensive applications such as graph analysis and is developed in both academic and industrial settings. Main challenges arise from:

- incrementally specified requirements, which lead to an organic growth of the software;
- the long lifetime of codes, which makes developers reluctant to adopt new and potentially unstable technologies;
- the difficulty to verify/validate the correctness of results that cannot be precisely reproduced in experiments;
- the lack of high-level parallel programming environments;
- multi-physics problems, which frequently require the coupling of methods across multiple length and time scales
- the desire to maintain portability across a range of modern and emerging parallel hardware platforms.

Under such constraints, performance and maintainability are often conflicting goals. While IDEs are increasingly integrated into HPC environments, the adoption of other state-of-the-art software-engineering approaches such as object-oriented design, development frameworks, and domain-specific languages have so far enjoyed only limited success in the HPC arena.

Moreover, the domain scientists who shape development practices rarely receive formal software engineering training – not to mention that software engineering curricula tailored to the specific needs of HPC barely exist.

In conclusion, we need to expand the currently mostly algorithm- and programming-centric view of HPC software development and achieve a better understanding of the (re-)design and quality management processes with the goal of providing appropriate methods and tools to support them. These must also be integrated into our training programs. Special attention must be paid to the needs of industry and community codes with a large user base, where changes may incur a very high cost.

3. WG 4.3 – Disruptive Technologies

The main objectives of working group 4.3 on “Disruptive Technologies” are to identify algorithmic possibilities for disruptive technologies and areas where breakthroughs can be made. There are aspects foreseen in the area of numeric (Monte-Carlo techniques, chaotic relaxation, and others), in areas such as stochastic programming, uncertainty quantification, new programming models, operating systems...

A disruptive technology can be defined as “a new technology that unexpectedly displaces an established technology e. g, the digital camera, the telephone, CMOS technology, RISC instruction set, smart phones...”¹

¹ Harvard Prof. Clayton M. Christensen

While identifying disruptions can be very difficult since it may only be recognised when already had happened, the group plans to follow the approach of identifying disruptions, in the sense of a change in philosophy, practice, culture... and from these disruptions identify the technologies that can cope with these disruptions.

Also, new technologies can be themselves the originators of disruptions. Since the working group aims to collect input from other working groups, its activities have been quite limited so far. Three experts have been identified so far and after the meeting topics that have been identified related to this WP make reference to efficient IO and data analysis, high-productivity programming models, a portable API for metrology for system level tools, VVUQ tools, SW engineering methodologies and tools for Exascale, and resilient OS.

4. WG 4.4 – Hardware and operating software Vendors

The main objectives of working group 4.4 on “Hardware and Software Vendors” are:

- to establish and maintain a global network of contacts with vendors in the HPC industry
- to leverage this network to investigate state of the art and trends related to an Exascale roadmap in the HPC hardware and software industry
- to propose initiatives for Europe based on these inputs

To address these objectives, a network with 13 HPC experts from the hard- software industry has been established.

During a face-to-face meeting, the experts identified a set of hard- and software challenges that need to be addressed and classified them according to their importance. According to the experts, the most important hardware challenges are (ordered by relevance):

- energy efficiency
- reliability and resilience
- data-processing closer to data
- memory/storage (capacity, packaging, bandwidth)
- multi-level interconnection networks
- the efficient use of additional transistors

On the software side, the most pressing challenges are:

- programmability and programming environments
- standardization of APIs and libraries
- inter-/intra-node scalability to 1M tasks
- data locality and avoiding data movement
- checkpoint/restart mechanisms and fault tolerance
- the creation of a set of characteristic mini-apps/benchmarks

The experts agreed that the software challenges are more urging than the hardware challenges and hence gave the following recommendations for public R&D funding:

- applications software in strategic areas
- technologies for data intensive computing and workflow awareness
- wide adoption of parallel runtime software
- strategic application libraries and data formats, scaling of important European ISV software packages
- adaptive system software optimization

- programming environments

4.4 Presentations of the Work Package 5: works of WGs, first R&D recommendations

1. WG 5.1 – Data management and exploration

Data management and I/O performance will strongly influence for the design of applications. IOs are not free anymore!

WG5.1 addresses « Data management and exploration » in Exascale applications as the organization of the scientific discovery workflow. This vision is complementary to the EESI1 one that had a strong technological focus. In particular, this WG studied the potential synergies between big data and Exascale as well as the consequences on applications design and development. Discussions among the experts point out that actions to promote connections between the communities HPC and Database are needed.

It was also highlighted that there is an issue in finding the right tradeoffs between storage and (re)-computation, in/out-situ analysis, incremental small computation vs large computation...

It was also considered that data storage management must be flexible enough to adapt to the data exploitation evolution during the scientific discovery process.

This translates in the need to build an ecosystem where computing, storage, network resources uses/deployments (and corresponding business model) are carefully planned and stable overtime to allow an efficient local (e.g. scientist view) and global (e.g. computing center operator) utilization. Support engineering teams, able to provide insights to scientists from the design phase to the implementation phase of the applications, will be a key component of this ecosystem.

One main recommendation issue during the meeting is to set up actions to address “End-to-end techniques for efficient I/O and data analysis » to describe the full life-cycle of data for a set of applications in order to produce designs/workflows that are consistent all the way from the production to the analysis of the data while considering locality, structures, metadata, right accesses, sharing etc.

One suggested manner to study this topic would be to specify scenarios for technology deployment and the available options for organizing the data storage and processing flow.

2. WG 5.2 – Uncertainties

HPC and uncertainty quantification have a two-sided relationship. On the one hand, the ever increasing size of the computational data leads to increasing sources of uncertainties, due to the accumulation of numerical errors. On the other hand, HPC gives access to computational power that can be used to tackle explicitly the evaluation of uncertainties, be it by embedded methods or by design of experiments.

A workshop dedicated to the tools and methodologies for uncertainty analysis was held in Paris on April 22-23. It was the opportunity for the experts to meet and exchange on the topic. Talks discussed various methods (spectral methods, Gaussian processes, reduced basis models, sensitivity analysis) and presented the way two software tools (URANIE and Openturns) exploit HPC computers.

First recommendations from the work group are the following:

- Educational aspects - diffusion of tools and practices
- Axes for progresses in computational efficiency of numerical methods
- ✓ Reduced basis models

- ✓ Adaptive designs of experiments for metamodels (NISP, kriging, NN, polynomials)
- Software and architectures
- ✓ - Taking into account DOE (Design of Experiments) -based methods in middleware
- ✓ - DOE tools checkpoint/restart procedures
- ✓ - Multiple levels of parallelism

3. WG 5.4 – Resilience

The working group 5.4 on resilience is composed of the following members: Franck Cappello (leader), Luc Giraud (INRIA), Torsten Hoefler (ETH), Simon McIntosh-Smith (Bristol), Christine Morin (INRIA), Bogdan Nicolae (IBM), Pascale Rosse-laurent (BULL), Osman Unsal (BSC).

The group considered several documents as a starting point for the gap analysis: IESP road map, 2011 EESI1 report, 2012 Report of the ICIS workshop: [Snir 2013] and a recent report from DoE: [Geist 2013]. The working group focuses on two different kinds of problems: 1) Process crashes (fail stop errors) and 2) Data corruptions (silent soft errors). Note that data corruptions could ultimately lead to process crashes. From the EESI1 recommendations, the working group noted that there is still no available fault model. Worst, the community has different views on failure rate, silent soft error increase (or not) with aggressive power saving techno. The working group considers important to investigate the RAS system (Reliability, Availability Serviceability system) reliability, the Runtime: MPI, task based resilience, High Performance Check pointing, High performance context saving/restore, Multilevel check pointing, Advanced fault tolerant protocols, Failure prediction, Algorithm resilience.

After a conference call organized on April 26 2013, the working group produced the following recommendations that were presented at the EESI2 Tremblay meeting in May 2013:

- Take seriously the fault model issue (there is no clear leader on this topic and Europe could play a leading role)
- Push Checkpoint restart as far as possible (USA and Japan are leading: ANL, UIUC, LLNL, UTK, Titech, U. Tsukuba, etc.)
- Fault prediction (USA is leading: UIUC, Argonne, IIT, Oak Ridge)
- Tasks based programming models (Europe is leading: OMPSS)
- Fault notification/management back plane (No clear leader: CFTS is no longer funded)
- Resilient OS and Runtime (USA is leading with MPICH and OpenMPI. No leader for OS yet)
- Resilient Algorithms (USA is a clear leader: UTK, Colorado)

So the main recommendation of the working group is to take a complementary position to USA and Japan:

- Provide a well-defined fault model
- Develop resilience for tasks based programming models
- Contribute to the development of a Fault notification/management back plane
- Contribute to the definition of a Resilient OS (make sure that OS confine faults)
- Collaborate on Resilient Algorithms
- To avoid overlap with USA and Japan and focus on complementarities the working group suggest establishing a forum/working group

4. WG 5.5 – Disruptive technologies

The presentation of the WG 5.5 was focused on the results of the discussion we had with the experts during a face to face workshop which took place in Milan (Italy) on the 15th of April 2013.

In particular disruptive technologies able to induce sharp changes in the way machine are designed, deployed, powered, cooled, and in their usage model, have been presented. The potential disruptions have been divided into different domains regarding:

- semiconductor technology;
- packaging;
- data transfer;
- memory;
- network;
- cooling and infrastructure;
- I/O subsystem

Whereas, the specific disruptive technologies identified by the WG are:

- Near Threshold Voltage Chip (NTV);
- NVRAM as a substitute for DRAM and/or Disks;
- 3D Stacking DRAM and Chip to short distances;
- Data Center Energy reuse;
- Electrons vs Photons as data carrier;
- End-to-end data exchange;
- Microfluidics.

After the list of disruptive technologies, the gap analysis has been presented. Here it is a list of the main results discussed during the talk:

- limited silicon photonics technologies should become available before 2020;
- NVRAM should reach consumer market by 2020 and become available (at moderate costs) for an Exascale system;
- 3D chip stacking will be there but limited to Memory;
- NTV chip will be ready by 2020;
- Microfluidics will be still limited to few packing geometries;
- liquid direct cooling will be the only option to obtain the required system density;
- Posix file system will become critical for performance.

The final recommendations to Exascale 2020 presented regards: investigate new I/O strategies for a tiered not only posix I/O subsystem, invest in new paradigm to favour locality and task parallelism (data-flow inspired), develop new API to couple applications to workload manager and energy monitoring system, even in an "introspective" way, and develop new end-to-end data transfer middle-ware and paradigms.

4.5 EESI2 in the context of EC DG-CONNECT, Exascale and HPC EC Strategy

P. Tsarchopoulos from DG-CONNECT gave a presentation on Horizon 2020 recent updates and expected schedule with a focus on topics potentially related to Exascale.

ETP4HPC, Prace and Centres of excellence are expected to work in close collaboration.

FET and Exascale actions will be most probably managed within the Research infrastructure area. This should facilitate cross links.

H2020 expected calendar is presently:

- Summer 2013 draft Work Programs
- Sept 2013 Program committee
- End 2013 first calls
- Spring 2014 close of first calls

4.6 Presentation of IDC HPC Study

This presentation was done by Earl Joseph from IDC (ejoseph@idc.com) and covered the following topics:

- IDC HPC Activities (www.hpcuserforum.com)
- HPC Market Update and Trends
- National competition for supercomputing leadership
- Creating economic models for HPC and ROI and for HPC and innovation

The last item is a presentation of a new IDC study that describes how increases in HPC investments can significantly improve economic success and increase scientific innovation. The concept of the study and preliminary results were presented.

The study includes creating two unique models:

- A macroeconomic model which depicts the way HPC investments result in economic advancements in the form of ROI, growth and jobs
- An "Innovation Index" that provides a means of measuring and comparing innovation levels, based on the level of applying HPC computing resources towards scientific and technical advancement

The presentation was followed by a deep discussion on the new IDC study.

4.7 Synthesis of the parallel sessions on Disruptive Technologies

Ecosystem (Carlo Cavazzoni)

This session has been organized in the form of a round table, where all EESI2 and invited experts discussed freely on the main topics emerged in their analysis and presentations.

The discussion starts from the definition of ecosystem and disruption itself, at the end of this first round of opinion the conclusion was that a disruption in the ecosystem is something that can cause an application to be totally rewritten or used in a complete different way. Experts also noted that a disruption can trigger another disruption and something a disruption does not introduce new paradigm but it can make an existing paradigm easy to use and implement by everybody. Like OpenMP that makes shared memory programming enough easy to be used by the majority.

Then the experts start discussing specific subjects than can cause disruption in the application world, like the developments of "Domain Specific Languages", and the use of python as "glue" for different application or components. The introduction of introspection functionalities and API can have also been discussed.

Then the discussion moved to the problems related to disruptive changes caused by energy limits, power capacity and heterogeneity that will most probably be the main architectures constraints for an Exascale system.

We shortly discussed about the impact of adopting NVRAM and a byte addressable I/O subsystem with the development of new APIs.

Moving to I/O, everybody agree on the fact that this is one of the components of main concern about ecosystem. We discussed disruption in the application to mitigate the problems related to the I/O like the possibility to rewrite application to perform analysis on the fly rather than writing to the disks/file system. It will be probably less costly to compute something again rather than waiting the time datasets are written or read from the disks.

Remaining on the I/O experts discuss about disruption coming from new way of managing I/O at large scale, like hadoop and map-reduce or the possibility to allocate and define (within the application) persistent objects.

Finally experts discussed the possible disruption coming from new APIs that will change the way how execution threads are scheduled at node and system levels.

Numeric (Godehard Sutmann)

The discussion included members from application domains, software engineering, computer scientists and hardware experts. It started with a clarification of the notion of disruptiveness and tried to come up with topics of common interest for further recommendations.

Disruptiveness was considered on a relative basis, meaning that disruptions may not necessarily affect all elements in a chain or hierarchy e.g. there might be disruptive technologies entering into the construction or production of a mobile phone without affecting the end user.

An example for a disruption which affects hardware and software is the advent of graphics cards as compute devices. Initially there was no disruption for the hardware, but the programming paradigm and the software support to address the GPU from a running program had to be changed. Later on double precision arithmetic and the need for direct communication between graphics cards disrupted also the hardware and the notion of General Purpose Graphics Processing Units (GPGPU) was established.

The meta language XML, substituting SGML, could be considered as disruptive which, although being not very different, offered a much larger spectrum of applications. For programming paradigms, domain specific languages (DSL's) are another example. At least in massively parallel simulation programs, DSL's are not established at present and therefore, if further developed, have the potential to disrupt common programming models.

Disruptions can also be expected on the algorithm side with respect to a paradigm shift in accounting metrics. If not only CPU-time is measured, but also energy consumption, memory and disk usage, which might be the case if memory and energy costs might exceed costs for cycles in future, this would have a consequence that energy aware algorithms and implementations with low memory usage together with optimization protocols and runtime models were to be considered. Another disruption is to be foreseen in the generation and handling of data. Due to the increasing gap between CPU performance and I/O-bandwidth capacities, the programs which generate the data will have to take over higher responsibilities to process data before writing them to disk. Therefore in-situ methodologies have to be developed which might consist of analysis-on-the-fly methods, (hierarchical) data filtering or automatic hot-spot detection to substantially reduce the amount of data.

As a consequence asynchronous execution models and autonomous load-balancing procedures have to be developed.

Applications (Iain Duff)

In the discussions of disruptive technologies in the context of Numerics and Applications, after identifying potential key disruptive technologies, the emphasis was on the requirements from applications that could establish disruptive breakthroughs in the application.

Weather prediction would benefit hugely if good algorithms and software were available for parallelizing the computation in the time domain as this is currently the bottleneck to exploiting parallelism in this area. In particular, this will be necessary if the potential of Exascale is to be realized.

The advent of Exascale could radically change the approach to solving many application challenges. For example, there are 10^{13} cells in the body so it is conceivable to use these as units in a computation where the computational power is at 10^{18} . Indeed breakthroughs are possible in much of the Meso scale domain and particle based methods might prove worthwhile to investigate.

The extension of data assimilation techniques used in climate modelling, to fields such as aeronautics might produce useful benefits.

In automotive applications, a great many relative small optimization computations need to be performed and techniques in both continuous and combinatorial optimization are required. The core requirement here and for many other applications is the availability of a good highly-parallel solver.

Exascale computing also opens the door to big advances in the solution of inverse problems both in the oil and gas industries but also in drug design. Computations with varying time scales (say in combining combustion with unsteady fluid flow) and in uncertainty quantification and associated stochastic approaches are also ripe for considerable advances of a truly disruptive nature.

5. Report on Day 2

5.1 Presentations of the Work Package 2: works, first R&D recommendations

WP2 aims to establish and maintain a global network of expertise and to act as a European voice for the Exascale Software Community.

It will investigate the state-of-the-art, trends, and future needs in HPC training and education. The goals are organized into four tasks, according to which WP2 will evaluate the following issues and will make recommendations.

- Investigate and describe state-of-the-art, trends, and future needs in HPC training and education. Work is in the process to evaluate the state of education, distinguishing between the training of scientists in the use extreme scale supercomputers, and education in a wider in the computational sciences sense. The latter survey aims at university education from the undergraduate to the PhD level where disciplinary boundaries prohibit the exploitation of computational techniques to their full potential. This affects Exa-scale computing across all fields of science and engineering.
- Establish and maintain a global network of expertise and funding bodies in the area of Exascale computing: a questionnaire addressed to all organizations financing research in the Xtreme Computing and Big data field has been designed and sent to a small group of experts for validation:
 - enquetes.agencerecherche.fr/index.php/survey/index/sid/781935/lang/en
Feedbacks from these experts are expected mid-June. The questionnaire will be sent to NFOs, European experts and projects beginning of July.
The contact point for this cartography is Sophie Despinoy at ANR.
- WP2 will also investigate and describe the establishment and landscape of co-design centers in the area of HPC and specifically Exascale computing. EESI will monitor the functioning of international existing centers. A work plan has been suggested which will take effect on time with M 10.
- Act as a proactive European voice and representative into the International Exascale Software Community:
 - ✓ A workshop has been organized on November 12, 2013 during SC'13 (Salt Lake City, USA) between the G8 NFOs, projects funded by the Exascale G8 2010 call and the research community. This workshop has not been funded by EESI2.
 - ✓ A new series of workshops US/Asia/Europe have been launched, the Big Data and Extreme-scale Computing (BDEC) workshops. The first edition took place in Charleston/USA - April 30 to May 01. The main focus was on applications (Natural Sciences) and Big Data.
 - ✓ The presentations are available at: www.exascale.org/bdec/agenda/charleston and white papers produced at: www.exascale.org/bdec/documents/charleston.
 - ✓ The second edition will take place in Japan, Kobe in February 2014, the third one in Europe (location to be defined) end of 2014.

5.2 Presentations of the Work Package 6: works, first R&D recommendations

The objective of work package 6 ("Operational software maturity-level methodology") are:

- Develop and document a methodology for estimating the level of maturity of Exascale software components
- Identify 3 software stack components from existing and near future European Exascale projects and apply the defined methodology
- Examine existing "equivalent" centres and propose structure adapted to Exascale software

Work in the first nine months concentrated on item 1. Items 2 + 3 will be done in the future once item 1 is finished. The work was mainly done by the four chairs and task leaders of the working group (Francois Bodin, Andrew Jones, Lee Margetts, and Bernd Mohr) and coordinated via monthly telephone conferences.

We started by researching existing approaches like Software Maturity Level, Capability levels, and various Technology Readiness Level approaches used by US DoD, US Air Force, NASA, European Space Agency and large companies. Methods which are closest our intentions are the QualiPSO Maturity Model (OSS) (see http://www.qualipso.org/trustworthy_process) published by the EU FP6 project QualiPSO and work done in the UK Software Sustainability Institute (see <http://www.software.ac.uk/>).

We also started to discuss and define a first draft of HPC software maturity proposal. We decided to restrict the proposal to four classes of HPC software components (application codes, frameworks and libraries, development tools, and programming model implementations). To assess the maturity of these HPC software components, we started to define a set of criteria (metrics) and methods to measure them (Documentation, Support, Availability, Coverage, Portability, Scalability, Performance, and Quality). For details please see the slides presented at the workshop.

In the near future, we want to finalize our proposal with the help of an expert meeting (external and members from other EESI2 work packages) and then document our findings in deliverable D6.1 "Report on operational software maturity level methodology".

5.3 Presentations of the Work Package 8

Main objectives of WP8 are:

- Organise two technical workshops (T8.1)
- Organise a European Exascale Conference (T8.2)
- Design and execute a communication strategy (T8.3)
- EESI in Europe and world-wide (T8.4)

It includes 5 deliverables:

- D8.1: Public website: M3
- D8.2: Dissemination Plan
- D8.3 and D8.4 for T8.1 (each technical workshop)
- D8.5: Exascale Conference for T8.2

The first technical workshop has been organized at Domaine du Tremblay and is presented in this report. The second workshop is planned on M25 at Cineca premises in Italy.

The European Exascale Conference is planned at M29 (January 2015) in Amsterdam. Final date will be decided in September 2013. 200 attendees are expected.

The dissemination plan (D8.2) is finished. Main contents are:

- Definition of target audiences
- EESI2 website
- Materials:
 - ✓ Posters, flyers, brochures
 - ✓ Press releases
 - ✓ Logo
 - ✓ Templates
 - ✓ Other promotional materials
- Events:
 - ✓ European Exascale Conference
 - ✓ Technical workshops
 - ✓ International community workshops
 - ✓ Supercomputing conferences
- Collaborations and monitoring

The EESI2 website based on an update of the EESI1 web site is available. A completely new site using up-to-date internet capabilities should be available in July 2013.

- <http://www.eesi-project.eu/pages/menu/homepage.php>

Task EESI in Europe/world-wide is done closely with T2.4 (Towards a public collaboration). One workshop has been done in USA early May 2013. Two others workshops are planned in Japan in March 2014 and in Europe in January 2015. A decision to synchronize or not this last meeting with the European Exascale Conference will be decided.

5.4 Synthesis of the parallel sessions on R&D Programs recommendations

The session Ecosystem and Applications provided the following recommendations:

- End-to-end techniques for efficient I/O and data analysis. Techniques to reduce I/O volume or improve I/O performance through better organization of the I/O and in-situ data analysis.
- High productivity programming models (DSL, rapid prototyping) which support heterogeneous architectures and exploit aspects such as dynamic data structures, locality, load balancing and communications.
- Improving code VVUQ (Verification, Validation and Uncertainty Quantification) tools to take into account multiple levels of parallelism
- Software engineering methodologies and tools tailored to the needs of Exascale with a focus on design and quality management (correctness, testing, performance, maturity)
- System software that adapt to resource variations (modeling) including Portable API for sys. Software and application for resource metrology and management
- New consistent (across software layers) resilience approaches for system and applications based on a clearly defined fault model
- Simulator : it seems to be very useful to everybody with some concern about cost and availability in a reasonable time

The session Numeric and Applications provided the following recommendations:

- Big data
 - ✓ Development of scalable methodologies for
 - handling big data from end to end
 - performing Data Analytics and massive processing
 - ✓ Issues: data driven algorithms with dealing with asynchronous communications for dynamic streams structures, structure representation and indexing beyond current model...
 - ✓ First step: create a European network of experts in scientific computing and data management
- Towards next generation couplers and associated methodologies
 - ✓ Multidisciplinary approach
 - ✓ International context of improving couplers
 - ✓ Code coupling algorithmic
 - ✓ Scalability of code coupling (memory, communication,...)
 - ✓ development of standard API for enabling interoperability
 - ✓ Next steps, creation of a network of excellence for identifying multidisciplinary needs and existing expertise, identify bottlenecks, extract concrete cases
- UQ & Optimisation methods and tools
 - ✓ Development of new methods and algorithm, standard APIs
 - ✓ UQ could be declined for optimization and can take into account data assimilation and inversion (for example in oil exploration or NDT or climate)
 - ✓ Next steps, creation of a network of excellence for identifying multidisciplinary needs and existing expertise, identify bottlenecks, extract concrete cases
 - ✓ 2 phases approach
 - identification of needs and expertise
 - Identification of a development strategy (algorithms, tools or full EU unified platform?)
- Exascale software simulators and mini apps
 - ✓ Development of open source mini apps "lighter versions of complex HPC applications embedding code essential features while demonstrating state-of-the-art implementations relying on modern programming paradigms"
 - ✓ International context with US and Asia
 - ✓ Link of ETP and PRACE prototypes made available
 - ✓ Link with training issues
 - ✓ Next steps : Call for proposal for
 - ✓ identifying and develop mini apps
 - ✓ Feasibility study and development of software simulators
- Development of ultra-scalable algorithms with quantifiable performance for realistic apps
 - ✓ Per nature interdisciplinary
 - ✓ Challenges: massive parallelism, heterogeneity, resilience, better fit with new meshing approaches...
 - ✓ Previously hidden into calls by applications (was the case in FP4, disappeared later)
 - ✓ Need to have strong European numerical libraries
 - ✓ Need strong link between fundamental and applied R&D
 - ✓ Some examples: linear algebra, time integration solvers, optimization, eigenvalues, particle methods, meshless methods...

6. Conclusion of the meeting

During the conclusion, the scientific key issues to be tackled were recalled.

The importance to deliver the reports in due time was underlined and an explanation for the content of each deliverable was given by the Project leader.

Guidelines for the recommendations and disruption chapters were presented. There is a need to elaborate concrete, informative and useful recommendations, which implies an important work.

7. Annex – Workshop Participants

	Surname	Name	Organization	Work Group
1	ANDRE	Jean-Claude	Jca Consultance & Analyse	WG 3.2
2	Ashworth	Mike	STFC	WG 4.2
3	Badia	Rosa M	Barcelona Supercomputing Center	WG4.3, WP4 leader
4	Bartolini	Andrea	ETHZ	WG 5.5:
5	Benini	Luca	Università di Bologna	T5.5
6	Berthou	Jean-Yves	ANR	2, 7
7	Bidot	Thierry	NEOVIA	WP 1
8	Bodin	François	University of Rennes 1	WP5.1, 5.4: WP6
9	Bogdan	Nicolae	IBM Research	Resilience
10	Brown	Nick	EPCC	co-design task)
11	Cappello	Franck	ANL	WG 5.4
12	Cavazzoni	Carlo	CINECA	WP5
13	DEMICHEL	Patrick	HP	WG 5.5
14	Duff	Iain	STFC - Rutherford Appleton Lab	WG 4.1 and WG 4.3
15	Earl	Joseph II	IDC	WP2
16	Epicoco	Italo	University of Salento & CMCC	WG3.2
17	Erbacci	Giovanni	Cineca	WP5 leader
18	Filippone	Salvatore	Università Roma Tor Vergata	WG 4.1
19	Goni	Ramon	BSC	WG 3.4
20	Grigori	Laura	INRIA	WG 4.3.
21	Ltaief	Hatem	KAUST	WG 4.3
22	Ludwig	Thomas	German Climate computing Center	WG 5.3
23	Margetts	Lee	University of Manchester	WG 4 and WP 6
24	McIntosh-Smith	Simon	University of Bristol	WG 5.3
25	Mohr	Bernd	Jülich Supercomputing Centre	WP6
26	Mueller	Matthias	RWTH Aachen University	4.2
27	Mueller	Thomas	FZJ	WG 3.3
28	Mula	Olga	UPMC Paris	WG 4.3.
29	Ott	Michael	Leibniz Supercomputing Centre	4.4
30	Popelin	Anne-Laure	EDF R&D	WG 5.2
31	RICOUX	Philippe	TOTAL SA	EEESI Coordinator
32	Ruede	Ulrich	FAU Erlangen	WP 2
33	Sutmann	Godehard	FZJ	WG 3.3
34	Tsarchopoulos	Panagiotis	European Commission	EC Project Officer
35	Unsal	Osman Sabri	Barcelona Supercomputing Center	WG 5
36	Vasseur	Xavier	CERFACS, Toulouse, France	WG 4.1
37	Vázquez	Mariano	Barcelona Supercomputing Center	WG 4.1
38	Winkel	Mathias	USI Lugano Insti. of Computat. Science	WG 3
39	Wolf	Felix	Germain Research School for Simul. Sci.	WG 4.2

