European Cooperation in the field of Scientific and Technical Research

- COST -

Brussels, 22 November 2013

COST 065/13

MEMORANDUM OF UNDERSTANDING

Subject :	Memorandum of Understanding for the implementation of a European Concerted Research Action designated as COST Action IC1305: Network for Sustainable Ultrascale Computing (NESUS)

Delegations will find attached the Memorandum of Understanding for COST Action IC1305 as approved by the COST Committee of Senior Officials (CSO) at its 188th meeting on 14 November 2013.

MEMORANDUM OF UNDERSTANDING
For the implementation of a European Concerted Research Action designated as

COST Action IC1305
NETWORK FOR SUSTAINABLE ULTRASCALE COMPUTING (NESUS)

The Parties to this Memorandum of Understanding, declaring their common intention to participate in the concerted Action referred to above and described in the technical Annex to the Memorandum, have reached the following understanding:

1. The Action will be carried out in accordance with the provisions of document COST 4114/13 "COST Action Management" and document COST 4112/13 "Rules for Participation in and Implementation of COST Activities" , or in any new document amending or replacing them, the contents of which the Parties are fully aware of.

2. The main objective of the Action is to coordinate European efforts for proposing realistic solutions addressing major challenges of building sustainable ultrascale computing systems by developing collaborative activities targeting cross-layer design issues to offer a unified view of ultrascale platforms.

3. The economic dimension of the activities carried out under the Action has been estimated, on the basis of information available during the planning of the Action, at EUR 80 million in 2013 prices.

4. The Memorandum of Understanding will take effect on being accepted by at least five Parties.

5. The Memorandum of Understanding will remain in force for a period of 4 years, calculated from the date of the first meeting of the Management Committee, unless the duration of the Action is modified according to the provisions of section *2. Changes to a COST Action* in the document COST 4114/13.

A. ABSTRACT AND KEYWORDS

Ultrascale systems are envisioned as large-scale complex systems joining parallel and distributed computing systems that will be two to three orders of magnitude larger that today´s systems. The EU is already funding large scale computing systems research, but it is not coordinated across researchers, leading to duplications and inefficiencies. The goal of the NESUS Action is to establish an open European research network targeting sustainable solutions for ultrascale computing aiming at cross fertilization among High Performance Computing (HPC), large scale distributed systems, and big data management. The network will contribute to glue disparate researchers working across different areas and provide a meeting ground for researchers in these separate areas to exchange ideas, to identify synergies, and to pursue common activities in research topics such as sustainable software solutions (applications and system software stack), data management, energy efficiency, and resilience. This Action will increase the value of the most active research groups in this area at the European-level by reducing duplication of efforts and providing a more holistic view to all researchers, it will promote the leadership of Europe, and it will increase the impact of those groups on science, economy, and society.

Keywords: Ultrascale computing sustainability, programmability of ultrascale applications, energy efficiency, data management, resilience.

B. BACKGROUND

B.1 General Background

The ever-increasing data and processing requirements of applications from various domains are constantly pushing for dramatic increases in computational capabilities. Today, we have reached a point where computer systems' growth cannot be addressed anymore in an incremental way, due to the huge challenges lying ahead, in particular scalability, energy barrier, data management, programmability and reliability. Action is urgently needed to bring together all the researchers working in the design and implementation of sustainable solutions for ultra-scale computing systems.

An ultrascale computing system (UCS) is envisioned as a large-scale complex system joining parallel and distributed computing systems, maybe located at multiple sites that cooperate to provide solutions to the users. As a growth of two or three orders of magnitude of today's computing systems is expected, including systems with unprecedented amounts of hardware, lines

of source code, numbers of users, and volumes of data, sustainability is critical to ensure the environmental, social, and economic impact of those systems. The NESUS Action aims at collaboratively rethinking the current basis of development of system software for scalable computing systems in order to pave the way towards a sustainable future scale growth by improving the coordination of efforts between complementary communities. Today's scientific community envisions and supports the emergence of Exascale systems within the next decade. In parallel, many companies and researchers are engaged in efforts of scaling data centres and system software to meet the requirements of diversifying on-line cloud applications and services. Both communities are already facing big data challenges and are developing their particular solutions to similar problems. While there is an emerging cross-domain interaction (for instance the need for high-performance in clouds or the adoption of distributed programming paradigms such as Map-Reduce in scientific applications), the cooperation between HPC and distributed system communities towards building the ultrascale systems of the future is still weak.

The NESUS Action aims at fostering research coordination for sustainable ultrascale computing. By coordinating researchers across Europe, NESUS will enable to more efficiently understand the limitations of current approaches of developing applications and system software and to study realistic solutions that address major challenges of building sustainable ultrascale systems: scalability, programmability, data management, energy efficiency, and resilience. Given the magnitude of the problem and the particular strengths of each community, bringing together communities that do not collaborate sufficiently at the moment will result in cross fertilization of ideas, elimination of redundant efforts, and the creation of a critical mass to address the huge challenges lying ahead.

This COST Action is needed to bring Europe together to tackle the fundamental challenges blocking our progress towards sustainable ultrascale systems. It will provide: strong coordination between different research teams with significant expertise in related fields; innovative cross-fertilization of ideas across fields (especially HPC, cloud computing, data centres, and big data); novel multi-disciplinary collaborations; avoid duplication of efforts; a better visibility of the European research; an increasing cooperation between scientists and industry. Moreover, any advance in sustainability will redound in societal benefits, given the quasi-ubiquity of computers in modern societies.

B.2 Current state of knowledge

The interest of governments, industry, and researchers in very large scale computing systems has

significantly increased in the last years, as a steady scale growth of computing infrastructures, such as data centres and supercomputers, has been produced and this evolution is expected to continue. A large number of worldwide research projects investigate system software solutions for supporting this scalability in cloud and exascale computing. However, each community is addressing only a part of the ultrascale sustainability problem, with sometimes overlapping efforts and wasting resources. Research in cloud computing is addressing a large variety of challenges related to sustainable scalability such as economy-of-scale, agile elastic scalability, energy-efficiency, scalable storage, automatic management of resources, resilience, hybrid clouds, etc. At the same time, several Exascale research projects explore novel approaches for advancing the system software towards new scalability levels, while addressing the main challenges of energy efficiency, high-performance, programmability, and resilience. International initiatives, such as the International Exascale Software Project (IESP) or the collaboration between European PRACE and US XSEDE, are mostly devoted to large scale parallel systems and scientific data systems. The advent of the Big Data challenges has generated new initiatives closely related to ultrascale computing systems in large scale distributed systems. It is noticeable that none of the former calls addresses ultrascale and sustainability in a holistic manner, while most of current research projects focus on specific types of systems.

The term ultrascale (or ultra-large scale) goes beyond the former concepts to characterize complex computing systems that are of unprecedented scale, with operationally and managerially independent subsystems, with evolutionary development, emergent behaviour, geographic distribution, and considering failures norm rather than exception. Due to the complexity of those systems our ability to develop, maintain, and manage such systems may fall behind their growth. There are some initiatives running on ultrascale computing systems, such as CUCIS and ULS in USA and LSCITS in UK, and some research activities on extreme-scale systems funded at international level, but they are fragmented in research communities (exascale, cloud computing, big data, energy-efficient data centres) with complementary strengths. If brought together, with a focus on sustainability, these research communities could provide a significant critical mass for addressing the current challenges of ultrascale computing. It is essential to take a holistic view to this challenge, if one is to meet the demand for future ultrascale systems.

NESUS Action will be innovative in addressing the existing and emerging problems of achieving sustainable ultrascale systems at two levels. First, specialized working groups will open up collaborative efforts for building innovative approaches, techniques, and algorithms to further advance towards Sustainable UCS. Second, the Action will collaboratively pursue a holistic approach to promote the integration of both horizontal and vertical solutions to improve the

sustainability of hybrid architectures mixing large scale distributed and parallel systems by means of better programmability, data management optimization, energy efficiency, and strong resilience.

B.3 Reasons for the Action

There are three major reasons to develop this Action at European level: the magnitude of the technical challenges; the interdependencies among areas that require multidisciplinarity and collaboration; and the scope of the effort, which must be international to reach a global impact.
In this regard, the main goal of this Action is to coordinate European groups working on parallel and distributed systems and to develop collaborative activities toward sustainability and scalability among the groups involved in this research topic. Thanks to this effort for defragmentation and coordination of the European research in this area, the expected outcomes of this Action are:

- Increasing awareness on sustainability of ultrascale systems by developing joint activities and publications in the field of sustainable ultrascale systems, the elimination of overlaps in research activities and improvement of collaborations targeting the same goal.

- Building a multi-disciplinary environment by attracting multi-domain applications scientists and industrial partners.

- Understanding and promoting this topic by exchanging researchers (faculty and PhD students), consortium meetings and workshops, and focused working meetings. This Action will contribute to train a new generation of scientists, by helping the career of starting researchers and bringing students to this area of research.

- Sharing and disseminating best practices and experiences to achieve sustainable systems by providing innovative approaches, techniques, and algorithms for reaching unprecedented levels of scalability in future computing platforms.

To achieve these goals, the Action will launch specific activities such as schools, workshops, meetings, short-term research visits, etc. where scientists will meet each other and work on the scientific problems identified as major challenges in this domain.
It should be noted that 57 EU researchers from 20 countries had already expressed their interest in this network at the stage of the full proposal for this Action. However, the network is intended to be open to new participants interested in this research area, with special emphasis on the incorporation of partners coming from countries which have recently joined the EU and partners from less-

favoured regions. Industrial partners are also interested in the activities of the Action.
The Action is both beneficial from a scientific/technological and economic/societal point of view. The potential impact of this Action is increased by the fact that experts participating in the proposal for this Action were already cooperating in major worldwide initiatives or in bilateral agreements. This is an indicator that opportunities exist for future collaborations. Moreover, any advance on sustainability (energy, programmability, resilience) will result in societal benefits.

B.4 Complementarity with other research programmes

There are several European research initiatives related to high-performance computing and large distributed systems. The Seventh Framework Programme (FP7) of the EC stimulates research, development and innovation in computing systems (Obj. 3.4), Exascale (Obj. 12), cloud computing (Obj. 1.2), and data centres (Obj. 6.2), all of them considered as separated niches with different focus. The PRACE network is oriented towards providing shared infrastructure for HPC. HiPEAC and PlanetHPC are NoE for computing systems and HPC. Exascale projects and initiatives such as CRESTA, MONT BLANC, DEEP, or EESI target to build massively parallel platforms, which will provide Exaflop performance within the next decade. Data centres projects are mainly devoted to energy and environmental performance of data centres. In the cloud computing area, ENISA and the EC have published “A Roadmap For Advanced Cloud Technologies Under H2020”, in which they identify scalability and sustainability among the main topics to be addressed to insure necessary advance in cloud environments.
The NESUS Action is complementary to the above European research efforts in the sense that a coordinated focus is currently missing, as most of them address sustainability at large scale as a research niche, and not in a holistic approach to bring together related communities. The Action focuses on research coordination, while most of the calls within the FP7 finance research and infrastructures. Second, this Action targets cross-layer issues in the software design and applications to offer a unified view of ultrascale platforms, with emphasis on sustainability and usability of those systems, while most of current research projects target specific types of systems. The holistic approach of the Action will allow to complement other research programs with ideas and solutions that could be applied to their areas and to make more effective use of the resources already devoted to research of those topics in the EU.

C. OBJECTIVES AND BENEFITS

C.1 Aim

The main objective of the Action is to coordinate European efforts for proposing realistic solutions addressing major challenges of building sustainable ultrascale computing systems (such as application programmability, adaptive runtimes, big data, energy saving, and resilience, while guaranteeing performance), as well as the development of collaborative activities among the involved research groups. While most of current research projects address specific types of systems, this Action targets cross-layer issues in the system software and applications design to offer a unified view of ultrascale platforms, with emphasis on sustainability and usability of those systems. The Action objectives can be summarized as follows: (O1) To foster original high quality research and to achieve the creation of a critical mass of researchers in the field of sustainable ultrascale computing (deliverables: position paper/survey on existing practices; innovative ideas guidelines/brochures, workshop proceedings, researcher exchanges, book on the projects topic, joint publications, training schools, applications catalog, and benchmark suite); (O2) To give coherence to the European research agenda in areas with traditionally different, but complementary research communities such as HPC/Exascale, big data, and large-scale cloud computing, in aspects related to sustainability (deliverables: best-case scenarios lists, workshops and meetings, white papers on good practises, research roadmap); (O3) To build a multi-disciplinary forum for cross-fertilization of ideas aiming to reach sustainable ultra-scalability, by addressing current challenges of HPC/Exascale, big data, and large-scale cloud computing (deliverables: project wiki, discussion lists, postings at popular scientific forums, newsletters, LinkedIn and Twitter groups). Notably the different leverages must be regarded together since they might influence each other, thus detailed studies on the impact of applying sustainable ultrascale computing practices to real problems of relevance in Europe will be published considering the holistic approach proposed in this Action.

C.2 Objectives

The Action aims at becoming a main reference point for the European research in sustainable ultrascale computing systems. This in turn will facilitate the launch of collaborations and joint focused activities within the European community, non-European entities working in the same field, and industry (already 4 companies interested).
The Action will also become an attraction point for researchers interested in taking into account the different aspects of sustainability in computing systems in their work, with a holistic approach merging the interests of the parallel and distributed systems communities. Sharing and merging existing and innovative good practices will lead the Action to propose and disseminate innovative

approaches of reforming the software stack, dynamic adaptation techniques, and scalable algorithms to increase sustainability in very large scale complex systems. Those practices will be oriented towards combining distributed and parallel computing, considering cross layer issues, such as energy and resilience, for reaching unprecedented levels of scalability on future computing platforms. A line of action will also be devoted to sustainable scaling the applications in preparation for the future ultrascale systems.
Continuous learning to identify the trends and needs in the field, by tracking the existing activities in the field and knowledge sharing, will allow to create and stimulate a discussion forum aimed to establish a common view among the involved communities, which may ultimately lead to EU project proposals and to standardization procedures in the field.
To broaden the Action societal impact, it will support and raise consciousness in sustainability for different audiences: system administrators, experienced scientists and engineers, industry CEO, general public, etc.

C.3 How networking within the Action will yield the objectives?

To achieve the objectives of this Action, management, training, mobility, and dissemination activities are foreseen to foster original high quality research in the area of sustainable ultrascale by different means, including:

- A **management and coordination** part which reflects in the Action organization in Management Committee (MC), Working Groups (WG), Focus Groups(FG), and coordination among them (see part E for details).

- **General Workshops** (one per year), that will be a multi-disciplinary forum for cross-fertilization among Action members and external players (research, industry, etc.). The theme will be decided each year by the MC after consulting members and other collectives interested. After each meeting, the results will be summarized and posted on the web page.

- Three **training schools** aimed to PhD students and young researchers from the groups in the project, but also to external participants. (Years 2, 3, and 4).

- Three **Doctoral Symposiums** to provide young researchers a good opportunity to share information and knowledge, and to present their current research. (Years 2, 3, and 4).

- Short-term **research visits** (Approximately sixty) mainly open to young researchers, which will receive a cross-disciplinary training and take advantage of existing resources.

- Fostering the organization of **international workshops** in major conferences, centred on the Action topics. Key speakers will be invited to participate to provide leadership in formulating a long-term roadmap to the different aspects of sustainability in ultrascale computing.

- **Dissemination** activities to increase Action impact and to raise scientific and societal awareness on sustainability through Web site, email, newsletters, blogs, social networks, etc. A Dissemination Board (DB) will be created.

- The production of annual **reports**, activity reports, etc. on the Action progress.

C.4 Potential impact of the Action

The overall benefits of the Action can be summarized as follows (ref. Section B and C1):

1. Fostering collaboration of very active European research groups in the field of sustainable ultrascale computing, increasing the cross-fertilization of scientists coming from different communities to structure and federate a disseminated community in Europe and to bridge the gap between theoretic research and real world applications by including stakeholders and industry players. If brought together, those communities will provide significant advances for the current challenges of ultrascale computing.

2. Facilitating the adoption and usage of sustainable ultrascale computing systems, by providing innovative solutions to advance the knowledge of designing sustainable ultrascale software and systems, which will be the basic facilities for new discoveries in science and technology and will have a direct impact on economic growth, society, and environment at European level.

3. Raising consciousness for sustainable computing systems, addressing application programmability, runtimes, big data, energy saving, and resilience, while guaranteeing performance.

4. Ultrascale computing could benefit key science areas like biology, astrophysics, climate science, or energy research, and provide tools to cope with big data-processing problems, like data analytics and management for social networks or Internet of Things. A catalog of benchmarks and applications will be elaborated.

5. Developing international collaborations, exchanges, and training for a critical mass of young researchers that will represent the next generation of specialists in this frontier area.

6. Providing highly useful experiences at the European academic level to form high-quality researchers and promoting science to university students and public at large.

Benefits beyond the Action time span can be expected, as several experts participating in the proposal for this Action were already cooperating and they have strong links, which is an indicator that further opportunities exist for getting funding for future collaborations.

C.5 Target groups/end users

The outcome of this network will provide significant benefits to researchers working in sustainable ICT and large scale systems. The Action will promote inter-disciplinary collaborations with different scientific groups for exploring efficient solutions to highly demanding problems (HPC, Cloud, data-intensive, etc.). The infrastructure, methodology, software, knowledge, and expertise,

developed in the context of the network will provide significant benefits to these groups and their scientific fields (e.g. biology, astrophysics, climate science, big-data, or energy research). The network will be open to new groups, which could collaborate in all the proposed topics by applying the research and knowledge in their areas (particularly in WG6).
The Action will actively attract industry partners interested in improving the sustainability and scalability of their system software stack and parallel and distributed applications. At the proposal stage four major industry players showed already interest to participate in the Action.
The Action will serve as a reference point to provide expertise to EU officials and authorities on the research in sustainability in ultrascale computing systems and ICT in general. Last, but not least, enhancing the sustainability of computing systems will have a direct impact on society, as will help improve many applications from various socio-economical domains such as ecology, business, medicine, etc., and by increasing the consciousness toward more sustainable solutions.

D. SCIENTIFIC PROGRAMME

D.1 Scientific focus

The NESUS Action will focus on a cross-community approach of exploring system software and applications for enabling a sustainable development of future high-scale computing platforms. In details, the Action will work in the following scientific tasks:

- First, the current state-of-the-art on sustainability in large-scale systems will be studied. The Action will strive for continuous learning by looking for synergies among HPC, distributed systems, and big data communities in cross cutting aspects like programmability, scalability, resilience, energy efficiency, and data management.

- Second, the Action will explore new programming paradigms, runtimes, and middlewares to increase the productivity, scalability, and reliability of parallel and distributed programming.

- Third, as failures will be more frequent in ultrascale systems, the Action will explore approaches of continuous running in the presence of failures. The Action plans to find synergies between resilient schedulers that handle errors reactively or proactively,

monitoring and assessment of failures, and malleable applications that can adapt their resource usage at runtime.

- Fourth, future scalable systems will require sustainable data management for addressing the predicted exponential growth of digital information. The Action plans to explore synergistic approaches from traditionally separated communities to reform the handling of the whole data life cycle, in particular: restructure the Input/Output (I/O) stack, advance predictive and adaptive data management, and improve data locality.

- Fifth, as energy is a major limitation for the design of ultrascale infrastructures, the Action will address energy efficiency of ultrascale systems by investigating, promoting, and potentially standardizing novel metrics for energy monitoring and profiling, modelling, and simulation of energy consumption and CO2 emission, eco-design of ultrascale components and applications, energy-aware resource management, and hardware/software codesign.

- Finally, the Action will identify applications, high-level algorithms, and services amenable to ultrascale systems and investigate the redesign and reprogramming efforts needed for applications to efficiently exploit ultrascale platforms, while providing sustainability.

D.2 Scientific work plan methods and means

The scientific programme will be developed in six Working Groups addressing different topics identified in D1 part of this Memorandum of Understanding. Only deliverables specific to each Working Group (WG) are mentioned. Shared actions such as training schools, meetings, workshops, etc. are specified in Section C3.

WG1: State of the art and continuous learning in Ultra Scale Computing Systems.
WG1 will continuously study the state of the art in hardware platforms and software stack of very large scale systems (data centres, post petaflop systems, multi-cloud environments, etc.). This Working Group will gather the participants experiences and knowledge about the possibilities to

adapt the underlying software stack of large scale systems in order to increase their sustainability. It will be useful for WG2, 3, 4, 5 and 6.
This working group will invite influential speakers from research and industry to give keynotes related to the interests of the Action. It will be closely related to training.
Deliverables

- A reference report on software techniques to increase sustainability in ultrascale systems. Updated every year. M6, M18, M30, M42

- Roadmap for research challenges and new starting points for prospective research in sustainable ultrascale systems. M21, M42.

- Book on sustainable ultrascale computing encompassing and gathering the whole experience. M36 to M48.

WG2: Programming models and runtimes.
This WG will promote new sustainable programming and execution models in the context of rapidly changing underlying computing architecture. The participants will explore synergies among emerging programming models and run-times from HPC, distributed systems, and big data management communities. The efforts will focus on improving the programmability of future systems that will likely reach substantially higher levels of concurrency and have heterogeneous architectures. This WG will explore programming models and run-times that facilitate the task of scaling and extracting performance on continuous evolving platforms, while providing resilience and fault-tolerant mechanisms to tackle the increasing probability of failures throughout the whole software stack. Further, this WG fill focus on data centric computing, as the increasing data demands of applications from various domains require novel programming paradigms and run-times for efficiency processing data, while keeping the energy consumption within acceptable budgets.
Key objectives: Scale handling (optimal usage of resources, faults), improve programmability, adaptation to rapidly changing underlying computing architecture, data-centric programming models, resilience, energy-efficiency
Topics: programming run-times for adapting to failures and load variability, trade-offs among programmability, performance, scalability and energy-efficiency.

Deliverables:

- Tutorial on sustainable programming models and runtimes. M15, M27.

- White paper on sustainable programming models and run-times. M36.

- Research papers on synergistic approaches for programming models and run-times. M6 to M48

WG3: Resilience of applications and runtime environments.

This WG will explore approaches of continuous execution in the presence of failures. The partners will work on innovative techniques to deal with hardware and system software failures or intentional changes within the complex system environment: resilient, reactive schedulers that can survive errors at the node and/or the cluster-level, cluster-level monitoring and assessment of failures with pro-active actions to remedy failures before they actually occur, and malleable applications that can adapt their resource usage at run-time.

The benefits of the COST Action are tremendous since the innovative techniques developed by each WG partner that only cover part of the software stack can immediately be presented and discussed in a much broader community such that the necessary integration of point solutions can take place.

Key objectives: monitoring and assessment of failures in Ultra-large-scale systems; Going beyond fail-stop errors to manage hard, transient, and failures in the SW stack; fault-tolerance at the ultrascale system level, devising integrated design approaches to get continuous service in the presence of continuous streams of errors; and understanding HW & SW dependencies and monitoring changes and their impact within complex systems.

Topics: fault tolerance in partitioned global address space (e.g. PGAS, MPI, hybrid) and federated cooperative environments; proactive actions based on efficient and reliable fault-prediction mechanisms. Exposing errors to the whole system and collective decision making process; dynamic replication of data and/or behaviour; supporting resilience at the application-level: developing malleable applications; support resilience at the infrastructure level: developing resilient schedulers at the node (shared-memory) or cluster (across nodes) level; automatic HW&SW context and dependency detection; and methods for verifying (qualitative) process performance against predefined requirements in dynamic environments

Deliverables

- Position paper on mutual influences between HPC and distributed computing. M15.

- Research papers on integration of resilience techniques across multiple levels of the software stack (e.g. across application-level and runtime level, across runtime level and hardware level). M6-M48.

- Meetings combining resilience issues with cross topics like energy, performance, data, and Quality of Service (QoS). M15, M27, M42. During an Action meeting.

WG4: Sustainable data management.

This WG will study data management lifecycle on scalable architectures in a synergistic approach combining HPC and distributed computing, as future scalable systems will require sustainable data management for addressing the predicted exponential growth and complexity of digital information. The goal of this WG is to explore and re-think the relationship between the data management lifecycle and scalable architectures in order to pave the way towards reaching sustainable ultrascale on the next generation of computing platforms. Significant emphasis will be placed on the cross-fertilization between HPC and distributed computing in addressing challenges such as scaling I/O stack, expose and exploit data locality, energy efficient data management, improve the scalability of big data applications and data analytics. Additionally, based on scalable data management the Action is looking to build a multi-disciplinary environment by attracting applications scientists from different domains who face big data problems.

Key objectives: contributing to the evolution of the storage I/O stack towards higher-levels of scalability and sustainability; data sharing/integration (globalization of data); improving the programmability of data management and analysis; improving the exploitation of data workload predictability and manage uncertainty through adaptivity.

Topics: synergies between HPC and distributed computing (e.g. HPC in the data clouds, Map-Reduce in HPC); scalability of storage I/O stack, reducing bottlenecks and exploiting data organization to pave the way towards ultrascale file systems; analyse the impact/integration of novel memory/storage technologies in the large-scale systems (Storage Class Memory, Shingled Disks, 3D stacked memory); investigate the role of data locality throughout the memory and storage hierarchy; predictive and adaptive data management for performance, elasticity, resilience; energy-aware data management in big data analysis techniques and applications.

Deliverables

- Position paper on mutual influences between HPC and distributed computing in data issues. M18.

- Scientific papers and novel research about reforming the storage I/O stack, integration of novel memory/storage technologies, data analysis techniques, predictive and adaptive data management. M6 to M48.

WG5: Energy efficiency.

Energy consumption is one of the main limiting factor for the design of ultrascale infrastructures. Multi level hardware and software optimisations must be designed and explored in order to reduce energy consumption for these large scale equipment. The WG5 will address the issue of energy efficiency of ultrascale systems in front of other quality metrics. The goal of this WG is to explore the design of metrics, analysis, frameworks and tools for putting energy awareness and energy efficiency at the next stage. Significant emphasis will be placed on the idea of "energy complexity", reflecting the synergies between energy efficiency and Quality of Service, resilience and performance, by studying computation power, communication/data sharing power, data access power, algorithm energy consumption, etc.

Key objectives: Exploring sustainability in ultrascale systems: Contributing to the monitoring and energy analysis of large scale infrastructures; Proposing new holistic models of energy consumption for ultrascale systems; Designing and studying energy aware software components; Exploring energy consumption and energy efficiency metrics for multiple resources (computing, storage, networking); Helping users to express interest in energy issues for their applications; More than « energy in usage of IT » : taking care of more complete life cycle of data and computing centres (production, transport, cooling, recycling) for energy models.

Topics: Profiling and reporting energy usage of ultra scale computing by taking into account the full system (IT and infrastructure (cooling)) with multi-objective adapted metrics; Monitoring systems for heterogeneous ultrascale systems (low power servers, Graphics processing units (GPUs),…); Models of energy consumption as basis for ultra scale infrastructures software components; Analyzing energy efficiency metrics in the context of ultrascale systems; Impact of applying green leverages at large scale: shutdown and slowdown of resources; Energy efficient resource management and scheduling; Synergies and links with other WGs: tradeoffs between energy and performance, resilience, and eco-design of applications.

Deliverables:

- Position paper on energy monitoring and analysis at large scale. M18.

- Joint workshops dealing with energy efficiency in HPC. M24.

- Meetings combining energy issues with cross topics like performance, resilience, data, QoS. M24 to M48.

- Research papers on energy efficiency and sustainability. M6-M48.

WG6: Applications.

This WG will identify algorithms, applications, and services amenable to ultrascale systems and will study the impact of application requirements on the ultrascale system design. Approaches from various but related fields are expected. The different approaches that will be studied include earth sciences, astrophysics, physics and chemistry (such us molecular dynamics), material sciences, biology and life sciences (such as genomics and HPC sequencing), health science, high energy physics (such as QCD), fluid dynamics, scalable robust multiscale and multi-physics methods and diverse applications for analysing large and heterogeneous data sets related to social, financial, and industrial contexts.

These applications have a need for ultrascale computing due to scientific goals to simulate larger problems within a reasonable time period. However, it is generally agreed that applications have to be rewritten substantially in order to reach ultrascale computing dimensions. Accompanying issues in the redesign of application codes are the inclusion of cross-cutting dimensions identified to be crucial for sustainable ultrascale computing and which will be explored in the other working groups of this Action.

Key Objectives: selecting a set of key applications with need for ultrascale computing concerning computational power and data storage; evaluation of the needs of the selected applications concerning scalability, programmability, portability, resilience; identification of computational patterns for expressing the applications at a higher level of abstraction for leveraging programming for ultrascale systems; categorization of applications for ultrascale systems concerning important characteristics, such as data requirements and distribution, computational structure and data access patterns.

Topics: analysis of the data requirement and data access behaviour of selected applications; identification of a set of key characteristics to determine a priori whether applications are amenable

for ultrascale computing; interaction with other WGs to explore selected aspects addressed in the WGs.
Deliverables:

- Report on the requirements for ultrascale systems from the applications' perspective. M6.

- Catalogue of applications with their specific requirements for ultrascale computing. M18.

- Benchmark suite for a standardized evaluation of ultrascale systems. M36.

- Joint research papers on sustainable ultrascale applications. M6 to M48.

The COST Action framework is the suitable scheme for gathering the energy of the different partners. Partners will individually (or in subgroups) explore various leverages to improve sustainability, proposing advances in research taking into account some software adaptation possibilities. The different scientific approaches will nourish each other from the experiences of the partners. The cross-fertilization of the approaches from distinct but related fields is expected. The different approaches that will be studied by the different partners will be complimentary with some necessary overlap. This is also reflected by organization of the WGs, vertical for WG2, WG4, and WG6, but horizontal for cross-layer issues as WG1, WG3, and WG5 to foster coordination and integration between participants.

E. ORGANISATION

E.1 Coordination and organisation

A **Management Committee** (MC) will be appointed to supervise the Action and to ensure the achievement of Action goals. It will be composed of at least a member per country with a maximum of two.

- A MC Chair will be elected at the first meeting. He or she will steer the project day by day, interfacing requests with the MC, as well as document the annual and final reports. Chairs of each particular activity will be also appointed (administration, dissemination, coordination, etc.).

- The MC will meet every six months to plan and design the activities of the Action (scientific program, reporting, etc.), to supervise activities implementation, and to allocate resources and manage budget.

Concerning **research coordination**, the Action will work as follows:

- The scientific work of the Action will be structured in **Working Groups**. For each WG, a Working Group Leader (**WGL**) will be appointed by the MC. Each WG will choose a WG Steering Board (**WGSB**) to coordinate the WG activities and to organize the scientific meetings. WG membership will be decided on the basis of research focus and interest, but each member will register to at least one WG. Young researchers will be encouraged to create small Focus Groups (FG) to explore specific research topics. Horizontal contacts across WG will be emphasized to exchange research directions, information and expertise of young and senior researchers.

- General Workshop of the Action once a year. Themes will be decided each year by the MC after consulting the Action members. External key speakers will be invited regularly. A report will be made for each workshop and published.

- STSM to foster bilateral collaborations among young researchers of research groups. A coordinator will be appointed by the MC to work with the Action Chair to choose and approve the visits.

- Training activities for young researchers will be implemented through three Training Schools, organized every year starting from Year 2, conditioned to budget availability. Topics of each training school will be decided each year by the MC after consulting the Action members. They will be complemented with three Doctoral Symposiums to allow

young researchers to share their PhD research. A report will be made for each activity and published.

- A Dissemination Board (**DB**) will prepare and update a dissemination plan and coordinate dissemination actions with the partners. A Dissemination Chair (**DC**) will be elected at the first meeting. The DB will guide the setting up of an Action website.

- A Collaboration Board (**CB**) will promote internal and external collaboration with other research projects and stakeholders, leaded by the Collaboration Chair (**CC**).

E.2 Working Groups

The research program of the NESUS Action presented in section D is decomposed in 6 Working Groups (WGs):

- WG1: State of the art and continuous learning in Sustainable Ultra Scale Computing Systems.

- WG2: Programming models and runtimes.

- WG3: Resilience of applications and runtime environments.

- WG4: Sustainable data management.

- WG5: Energy efficiency.

- WG6: Applications.

E.3 Liaison and interaction with other research programmes

The following key fora are the most likely to benefit and/or provide interactions for mutual interest:

- The Action will establish links with the FP7 European NoE Hipeac and Support Actions PlanetHPC and HP-SEE since they gather a lot of researchers involved in the systems and HPC community. NESUS experts already participate in Hipeac and HP-SEE, thus, close cooperation is ensured.

- Liaison will be fostered with FP7 research projects related to NESUS: next generation of storage (Eurocloud, RDA), programmability (RELEASE, REPARA), exascale (CRESTA, DEEP, and MONTBLANC), cloud computing (MODAClouds, mOSAIC, HOST - HPC), data centres energy (CoolEmAll and All4Green), and resilience (TIMBUS, WF4EVER, EUDAT). Experts from those projects will be invited to NESUS meetings and seminars.

- Important for the Action will be the cooperation with European Exascale Software Initiative (EESI2), the International Exascale Software Project (IESP), and the Energy Efficient High Performance Computing (EE HPC) working group. They will allow wide diffusion at international level. NESUS members already cooperate in those initiatives.

- The Action will interact with EU decision makers (such as e-Infrastructure Reflection, e-Infrastructure Policy Forum, PRACE, and EGI-InSPIRE), stakeholders and industry players associations (as ETP for HPC, EuroCloud, GreenGrid, or Big Data Public Private Forum), and standardization organizations (as ISO and OGF). Experts from the Action already participate in several of these organizations, thus cooperation is ensured. Links with these project will give a voice to this Action to a broad spectrum of users, managers and standardization bodies to increase the impact of the Action.

- Several NESUS partners participate in the SHIWA community (http://www.shiwa-workflow.eu/) that supports interoperability among different scientific workflow systems (ASKALON, MOTEUR, Triana, P-GRADE) and Distributed Computing infrastructure. The cooperation with this community will ensure that the NESUS

research, especially in WG2, WG3, and WG4, will reach a broad international scientific community, thus improving interoperability at the European level.

The List of Experts having participated in the proposal for this Action included individuals and institutions involved in most emerging initiatives in the very large scale systems in Europe, with links existing with similar initiatives worldwide (particularly US and Japan). This fact, together with CB activities, will allow for a cross-fertilization among different research groups and approaches. At the same time, it represents an indication that the Action has the critical mass to act as a central point of expertise in the research community.

E.4 Gender balance and involvement of early-stage researchers

The Action will promote the involvement of female experts in the different activities, especially at the different levels of management, and to adhere to gender equality for new recruits. Female scientists were already involved in the elaboration of the Action proposal. Gender balance will also be promoted in STSMs and Training School places.
The Action will promote the participation of early-stage researchers in different activities, fostering leadership in FG. STSM of Early Stage Researchers (ESR) will be a central pillar of the Action to develop contacts. Training schools will be devoted to ESR to build up expertise and learn early about different techniques and methodologies. ESR will participate in the organization of Training schools.
In all cases, encouragement of applications from all sectors of the community (including minorities) will be clearly stated. No discrimination based on sex, race, religion or citizenship will be considered. The systems developed also provide means to improve equal opportunities for disabled persons.

F. TIMETABLE

This NESUS COST Action is **planned for four years**.
The timetable for the Action activities is shown below:

- A Management Committee meeting is planned every 6 months.

- Each Working Group will meet at least twice a year, every year, at months Year+[3 to 4], and Year+[9 to 10].

- The DB will deliver the dissemination plan in month 3. It will be updated at least every year, meaning at Months 12, 24, and 36. Issuing a deliverable for each new plan. WG deliverables are specified in Section D.2.

- The planning of short-term visits will start after the first Working Group meetings, once participants will have identified common research interest. Up to 15 STSM will be appointed per year, conditioned to the budget.

- Three training schools will take place (Year-2, Year-3, and Year-4).

- Three Doctoral Symposiums will take place (Year-2, Year-3, and Year-4), joint with schools.

- A final general meeting will be organized at the end of the Action, in Month 48.

- At the end of the Action, Months 47 and Month 48, the MC will produce the final report of the Action. In the report, the MC will evaluate and monitor appropriately the achievements of the objectives of the Action in respect with section C.

The general timeline for the WG activities is the following:

- WG1: starts M0. It will issue a first deliverable after 6 months, that which will be updated every year, at Months 6, 18, 30, 42. Contribute to the annual reports of the Action.

- WG2, WG3, WG4, WG5, WG6: starts M+3. Ends M+46. Contribute to the annual report of the Action.

- All WG: M47 and M+48, contribute to the final report of the Action.

The main milestones for the first year of the Action are shown in the following table.

Month	Milestones	Actions
M+0	First Management Committee meeting and Kick-off meeting.	Elect WG leaders, decide on the upcoming activities, dissemination board, coordination board, and nominate chairs for all foreseen activities (training school, STSM, …).
M+2 to M+3	First Working Group meetings.	Identify scientific synergies, present work plan, decide topic of next meeting. Dissemination plan delivered.
M+3	Start of the first year short-term visit campaign.	Announce up to 15 STSM. Look for interested partners.
M+6	Working meeting and Management Committee meeting.	Presentation of early results, identified synergies, website, presentations and discussions on the theme of the meeting.
M+9 to M+10	Second WG meetings.	Present progress and topic of the next meeting.

At **M+13** the Management Committee will meet again and decide the organization and milestones for the second year. If the organization of the first year is found successful it will be reuse for the next years, but including new activities such as training schools and doctoral symposiums. Otherwise a more effective organization will be proposed.

G. ECONOMIC DIMENSION

The following COST countries have actively participated in the preparation of the Action or otherwise indicated their interest: AT, BE, BG, CY, CZ, DE, EL, ES, FR, HU, IE, IT, LT, NO, PL, PT, RO, SE, SI, UK. On the basis of national estimates, the economic dimension of the activities to be carried out under the Action has been estimated at 80 Million € for the total duration of the Action. This estimate is valid under the assumption that all the countries mentioned above but no other countries will participate in the Action. Any departure from this will change the total cost accordingly.

H. DISSEMINATION PLAN

H.1 Who?

As target audience for project dissemination the Action will address a multitude of groups, exposing scientific results and raising the awareness of sustainability in very large scale parallel and distributed ICT systems.

Academia and research. The Cloud, Grid, supercomputing, and data centres communities are directly concerned with the efficient and sustainable usage of vast computing resources. Since many Action outcomes will be based on parallel/distributed algorithms and runtimes developed by the research team members of this Action, the Action addresses the whole scientific computing community. Thus, many academic groups that have high computational needs could benefit from the results of the Action (e.g., astronomy, multi-physics, biology, etc.).

Industry is a main target for the Action, especially IT staff and managers of data and computing centres and Internet Service Providers, who could enhance efficiency and quality of service due to the results of NESUS Action, reducing a non-negligible parts of their budgets through sustainable solutions. It also addresses software vendors that will find in the Action innovative algorithms and programming methods to enhance scalability, resilience and energy efficiency. There are four industrial partners interested in the Action, but the Action will pursue the potential involvement of new start-ups from the HPC, cloud, and big-data areas. Relations with hardware vendors, who can incorporate our research into products, will be also fostered to promote that our research reaches the marketplace.

Standardization Bodies. The research teams involved in the Action use standard tools (such as MPI, OpenCL, OpenMP, OpenStack Compute, …) and open popular frameworks (such as MapReduce, OpenNebula, …). Moreover, members of the Action already participates in standardization committees, (ISO SC22 C++), Cloud distributed services (ISO SC38), Energy efficient Data Centres (ISO SC39), and OGF.

European Union. The Action also addresses the EU as a whole, since the focus of the work is to enhance very large scale computing environments making them more sustainable. This goal fits to all supercomputing installation, data centres, and very large scientific installations. As a result of the Action activities, a white paper will be elaborated including recommendations for EU policy makers.

General public will also benefit from the Action, as easier to use and more efficient and reliable large scale services could be accessible from home, opening them new business opportunities.

H.2 What?

The results of this Action will be disseminated publicly through the following means:
-Publications:

- **Research papers**. Peer-reviewed articles in high-quality journals and conferences. The Action will encourage co-publications between partners.

- **Technical reports**. Based on the research results are created the involved partners will create internal technical reports, available in the Intranet of the Web site. They will consist of case study reports, state of the art reports, manuals, integration results, activity results, etc.

- Publication of at least one **book** on the topic of sustainable ultrascale computing in a high-quality thematic publisher.

- Annual Action and Working Group **reports** will be regularly published on the public Web site.

- DB will regularly produce a **newsletter** relating the advances and activities of the Action (minute reports, Action status, events, news, membership, STSM, …). Available online at the Web site and emailed to partners and interested bodies.

-Event organization and participation:

- **Annual Action Workshop** will be a forum to show research results of the partners. Open to invited speakers, other research groups, and industry, workshop proceedings will be published with ISBN and will be posted on the Action Web site according to the COST Office policy.

- **Research Workshops** organization and participation in major conferences where the community is already actively participating (Eurosys, EuroPar, EuroMPI, HeteroPar, IPDPS, ICPS, CLUSTER, CCGrid, WWW, EnaHPC, SC, Hot Storage, etc.).

- **Seminars** for industry players where current knowledge and research results are presented to the public, people from software industry, hardware vendors etc. These seminars will be organized together with other bodies that are working in related research areas.

-Web diffusion and social networks:

- **Public web site** with open information about the Action and a moderated blog to foster discussions about the Action research topics. It will display information about the Action itself, partners, general information, and news about current events, activities, and results.

- **Electronic communication** network by using email lists, forum, Twitter, and LinkedIn groups.

- **Intranet** for internal communication and support of WG meetings and drafts. It will include management information, a project Wiki, forms and documents, reports, publications of the partners, etc.).

- **Application catalog** with sustainable ultrascale applications and benchmarks provided by the members of the Action.

-Non-technical publications.

- **Press** official statements, newspapers, leaflets, etc. The partners will send articles about their work and research results to newspapers, magazines etc. for dissemination of the topic to the general public.

- Information for industrial partners and stakeholders through the Web site.

H.3 How?

To guarantee an appropriate dissemination of the results of the Action, the DB will prepare a

dissemination plan to ensure that all dissemination methods are properly addressed and implemented, during and after the Action run time, and that all the target audiences are addressed. The plan will include a detailed time schedule for each dissemination tool and activity within the first year of the Action, and it will be continuously updated and revised during the Action (see Section F). The DB will ensure that the necessary steps are taken to implement this schedule, for instance by choosing locations for the workshops, or scheduling meetings for discussing books, Web contents, etc.

As part of the dissemination effort, the partners will create **documents and reports** of results of the WGs, pursuing horizontal integration. As part of the annual workshop, a plenary report will be issued to disseminate the results inside and outside the NESUS consortium. The annual meeting of the Action and summer school will be the perfect occasion for strengthening the dissemination by inviting key actors of the domain (industry, non-EU country, IEEE or ACM, other COST Actions, etc.).

For disseminating the scientific results, partners will organize **workshops at major conferences** and will promote joint publication of journal papers, white papers, books, etc. Several experts involved in this Action are involved in the editorial boards of well-known journals in the field as well as organizers of scientific events. Moreover, the Action will send representatives to the major events of the domain of high performance and distributed computing for participation to forums, booths, panel discussions, etc.

A **newsletter** will be produced every six months including Action activities, results, and information for interested bodies. It will be distributed through email, social networks, and the Web page, also to make dissemination for the general public, that will be promoted by encouraging Action experts to publish **press** official statements, newspapers, leaflets, articles in magazines, etc. To disseminate Action results to industry players and EU policy makers, the Action will organize **seminars**, which results will be published in the Action Webpage and included in the newsletter. All action results will be available on a dedicated website. The public part of this **website** will be composed of: description of the Action (goals, means and results); list of members; description on the procedure to join the Action; relevant published documents (technical reports, presentations, seminar and schools results, newsletter, standardization proposal, etc.); event calls (conferences, meetings, organized workshops, etc.); news on the life of the Action; and social communication (LinkedIn, Twitter, blogs, and forum). The private part will include a wiki, where members of the Action will be able to collaborate on the edition of a knowledge base, blog, forum, and internal documents published. Moreover two mailing lists will be set: one for internal communication and one for non-members interested in the results and the activity of the Action. The DB will monitor

the popularity and usage level of the Web site, taking measures to increase impact, such as inviting regularly relevant people to place comments or articles in the website, and creating links to other websites or events related to the topic.

Dissemination activities that should take place **after the Action run time** will be collected during the first half of the Action and written into a dissemination plan that will be constantly updated during the Action time. This dissemination plan will be then included into annual report of the penultimate Action year for DC approval. The plan must make sure that partners agree to carry out the dissemination within a defined time frame, and that the necessary financial means are available.