Deliverable D2.3

Report on International Activities

CONTRACT NO EESI 261513
INSTRUMENT CSA (Support and Collaborative Action)
THEMATIC INFRASTRUCTURE
START DATE OF PROJECT: 1 June 2010 DURATION: 18 months

Name of lead contractor for this deliverable: NCF
Authors: Peter Michielse (NCF), Patrick Aerts (NCF)
Name of reviewers for this deliverable:

Abstract: This deliverable reports on progress in exascale projects worldwide. Its main source of information are the IESP meetings. This report will be extended after each IESP meeting.

Release N° 1

Due date of deliverable:	30.11.2011
Publication date:	<To be completed >

Release N° 2

Due date of deliverable:	
Updated sections:	<To be completed > ….
Publication date:	<To be completed >

Release N° 3

Due date of deliverable:	
Updated sections:	<To be completed> …
Publication date:	<To be completed >

Project co-funded by the European Commission within the Seventh Framework Programme (FP7/2007-2013)		
Dissemination Level : PU		
PU	Public	X

Table of Contents

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)
GENCI	Grand Equipement National de Calcul Intensif (France)
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 Center
LRZ	Leibniz Rechnung Zentrum
NCF	National Computing Facilities Foundation (the Netherlands)
NNSA	National Nuclear Security Administration (US)
NSF	National Science Foundation (US)
SKA	Square Kilometre Array
WG	Working Group (in EESI)
X-stack	Exascale software stack

1. Executive summary

*** No need for this now, as it is a living document. ***

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. In order to address the expectedly increasing gap between projected exascale-class hardware and efficient software usage, a worldwide effort has been set up at the beginning of 2009, which is known as the International Exascale Software Project (IESP), [1].

With seed funding from key government partners in the United States, European Union and Japan, as well as supplemental contributions from some industry stakeholders, IESP has been formed around the following mission:

The guiding purpose of the IESP is to empower ultra-high resolution and data-intensive science and engineering research through the year 2020 by developing a plan for (1) a common, high-quality computational environment for petascale/exascale systems and (2) catalyzing, coordinating, and sustaining the effort of the international open source software community to create that environment as quickly as possible.

As has been extensively described in [1], there exist good reasons to conclude that such a plan is urgently needed. Firstly and foremost, the scope of the technical challenges for software infrastructure that the novel architectures and extreme scale of emerging systems bring with them is daunting. These challenges, which are already apparent on the leadership-class systems of the US National Science Foundation (NSF) and Department of Energy (DOE), as well as on systems in Europe and Asia, are more than sufficient to require the complete redesign and replacement of the operating systems, programming models, libraries, and tools on which high-end computing necessarily depends. Secondly, the complex web of interdependencies and side effects that exist among such software components means that making sweeping changes to this infrastructure will require a high degree of coordination and collaboration. Failure to identify critical holes or potential conflicts in the software environment, to spot opportunities for beneficial integration, or to adequately specify component requirements will tend to retard or disrupt everyone's progress, wasting time that can not be afforded to be lost. Since creating a software environment adapted for extreme-scale systems (e.g., NSF's Blue Waters) will require the collective effort of a broad community, this community must have good mechanisms for internal coordination. Third, it seems clear that the scope of the effort must be truly international. In terms of its rationale, scientists in nearly every field now depend on the software infrastructure of high-end computing to open up new areas of inquiry (e.g., the very small, very large, very hazardous, very complex), to dramatically increase their research productivity, and to amplify the social and economic impact of their work. It serves global scientific communities who need to work together on problems of global significance and leverage distributed resources in transnational configurations. In terms of feasibility, the dimensions of the task—totally redesigning and recreating, in the period of just a few years, the massive software foundation of computational science in order to meet the new realities of extreme-scale computing—are simply too large for any one country, or small consortium of countries, to undertake on its own.

During 2009 IESP has organized a series of three international workshops, one each in the United States, Europe, and Asia in order to work out a plan. This has led to a technology roadmap to address a critical challenge that now confronts modern science and is produced by a convergence of three factors:

1. the compelling science case to be made, in both fields of deep intellectual interest and fields of vital importance to humanity, for increasing usable computing power by orders of magnitude as quickly as possible;

2. the clear and widely recognized *inadequacy* of the current high end software infrastructure, in all its component areas, for supporting this essential escalation;
3. the near complete lack of planning and coordination in the global scientific software community in overcoming the formidable obstacles that stand in the way of replacing it.

This document is organized as follows. Section 3 will briefly document the four IESP workshops held before the start of the EESI project (three in 2009 and one in 2010). Section 4 will report on the first IESP workshop held during the actual EESI project (in October 2010 in Maui, Hawaii, USA), including a summary of the major findings and actions as a result of the workshops.

This is a living document which will be updated after each IESP workshop, during the course of the EESI project.

3. Overview of IESP workshops until April 2010

3.1 Introduction

As has been documented in [1], IESP develops a plan for producing a new software infrastructure capable of supporting exascale applications. In order to do so, the following sequence of objectives needs to be addressed during the workshops:

1. *Make a thorough assessment of needs, issues and strategies*: A successful plan in this arena requires a thorough assessment of the technology drivers for future peta/exascale systems and of the short-term, medium-term, and long-term needs of applications that are expected to use them. The IESP workshops brought together a strong and broad-based contingent of experts in all areas of HPC software infrastructure, as well as representatives from application communities and vendors, to provide these assessments. Also, leverage has been taken of the substantial number of reports and other material on future science applications and HPC technology trends that different parts of the community have created in the past three years.

2. *Develop a coordinated software roadmap*: The results of the workshop attendees' analysis have been incorporated into a draft of a coordinated roadmap intended to help guide the open source scientific software infrastructure effort with better coordination and fewer missing components.

3. *Provide a framework for organizing the software research community*: Development of an organizational framework to enable the international software research community to work together to navigate the roadmap and reach the appointed destination—a common, high quality computational environment that can support extreme-scale science on extreme-scale systems. The framework will include elements such as initial working groups, outlines of a system of governance, alternative models for shared software development with common code repositories, and feasible schemes for selecting valuable software research and encouraging its translation into usable, production-quality software for application developers. This organization must also foster and help coordinate R&D efforts to address the emerging needs of users and application communities.

4. *Engage and coordinate with the vendor community in cross-cutting efforts*: To leverage resources and create a more capable software infrastructure for supporting exascale science, the IESP is committed to engaging and coordinating with vendors across all of its other objectives. Industry stakeholders have already made contributions to the workshops (i.e.. objectives 1 and 2 above) and IESP expects similar, if not greater participation, in the effort to create a model for cooperation as well as coordinated R&D programs for new exascale software technologies.

5. *Encourage and facilitate collaboration in education and training*: The magnitude of the changes in programming models and software infrastructure and tools brought about by the transition to peta/exascale architectures will produce tremendous challenges in the area of education and training. As it develops its model of community cooperation, the IESP plan must, therefore, also provide for cooperation in the production of education and training materials to be used in curricula, at workshops and on-line. This roadmap document, which focuses objectives 1 and 2 above, represents the main result of the first phase of the planning process. Although some work on tasks 3–5 has already begun, we plan to solicit, and expect to receive in the near future, further input on the roadmap from a much broader set of stakeholders in the computational science community. This version of the roadmap begins that process by including more extensive input from the science application community, international funding agencies, and vendor partners. The additional ideas and information we gather as the roadmap is disseminated are likely to produce changes that need to be

incorporated into future iterations of the document as plans for objectives 3–5 develop and cooperative research and development efforts begin to take shape.

The following sections will give a brief overview of the IESP workshops, held before the start of the EESI project on June 1, 2010. All IESP meetings so far have been by invitation only.

3.2 Workshop in Santa Fe, NM, USA

After an initial meeting during SC08 in Austin, the first workshop has been held on April 7, 8, 2009, in Santa Fe, New Mexico, USA. This first workshop brought together an initial set of representatives from various countries, all active in the area of HPC and software development for demanding scientific applications. Appendix A lists the attendees of all workshops so far.

During this initial meeting, time has been reserved for an introduction to many topics: HPC situations in the USA, Europe and Japan, science drivers, hardware feasibility and challenges, software bottlenecks and funding and organisational issues. It was agreed to develop a version of the IESP roadmap after the third workshop, if possible before SC09 in Portland. This roadmap should consider HPC architectures and hardware, software roadmap and computational science research strategies, but also approaches to funding, on how to collaborate with vendors which also have their own commercial interests. Initial ideas have been raised, brainstorming has been done, which were planned to be developed in detail during the next two workshops in Paris and Tsukuba.

3.3 Workshop in Paris, France

The main goal of the Paris meeting (June 28, 29, 2009) was to initiate development of a coordinated software roadmap and get the participants fully engaged with the process of developing a draft version. Since the roadmap is intended to be a living document, used to guide and coordinate the work of the international open source software community over time, it is intended to serve as the vehicle through which progress on the other objectives in the plan are expressed. For example, the most fundamental set of inputs into the roadmap are the assessments of the short, medium and long term needs, issues, and strategies emerging from an analysis of emerging peta/exascale architectures and the scientific research agendas IESP aspires to serve. Consequently the majority of the effort at the second workshop was concentrated on those analyses and assessments. On the other hand, the other objectives of the plan can be thought of as a derived product of the roadmap, once it is created. These objectives include the following:

- Provide a framework for organizing the software research community;
- Encourage and facilitate collaboration in education and training;
- Engage and coordinate vendor community in crosscutting efforts.

After plenary sessions to establish common context, the participants at the workshop divided into four breakout groups –software, applications, vendors (industrial collaborators), and funders. The division was intended to facilitate the analysis and assessment of the underlying needs, issues and strategies that must feed into the roadmap. The software group, which was the largest, found it natural to divide again in order to attend to intra-processor and inter-processor issues separately; the results of their discussions were then consolidated at the end of the workshop.

In general, the second workshop provided the framework that the project needed to pursue in going forward. More specifically, the key components of roadmap have been identified and their content has begun to be filled in. An initial group of authors who will lead the effort to take the roadmap components forward has been identified, though it is clear many volunteers will be needed to help finish this work in or for the project to stay on schedule.

3.4 Workshop in Tsukuba, Japan

The goal of the workshop in Tsukuba (held from October 18-20, 2009) was to prepare a draft version of the roadmap, for which the framework was set in the Paris meeting. This basically meant that during break-out groups in the workshop, many topics that were identified in the framework roadmap have been discussed in break-out groups, eventually leading to pieces of text to include in the draft roadmap. After the Tsukuba workshop, lots of effort has been put in editing the contributions, presenting a consistent document and avoiding too much overlap.

In fact, the IESP software roadmap is a planning instrument designed to enable the international HPC community to improve, coordinate and leverage their collective investments and development efforts. After the Tsukuba workshop, the software components for exascale computing (X-stack) have been described, together with the principle of co-design, in which hardware and software development will go along with application needs and perspectives.

A key aspect here is a timeline, starting now up to the end of the decade. For each X-stack component, targets and milestones have been defined, which must lead to a full-blown, tested X-stack in the 2018-2019 timeframe.

Another result of the Tsukuba workshop has been that is was recognized that organizational and planning issues need to be discussed as well. This has been an important subject of the next workshop in Oxford, UK.

3.5 Workshop in Oxford, UK

The workshop in Oxford, UK, has been organized from April 12-14, 2010. Important organizational aspects have been discussed, among which:

- Cooperation between IESP and HPC vendor communities;
- IESP organization and governance;

In order to address these aspects, more focused break-out groups have been set-up for the funding agencies on organizational and governance aspects and also for the vendors and facilitators on their considerations.

These discussions have eventually led to version 1.0 of the roadmap, May 30, 2010.[1]

[1] International Exascale Software Project: Roadmap 1.0, www.exascale.org.

4. Workshop in Maui, HI, USA

4.1 Introduction

The fifth IESP workshop has been organised in Maui, HI, USA, on October 18-19, 2010. At the start of the workshop, the goals included the following:

- Refine software roadmap for software and algorithms on extreme-scale systems based on feedback since initial release;
- Refine the prioritized list of software components for Exascale computing as outlined in the Roadmap;
- Refine the assessment of the short-term, medium-term and long-term software and algorithm needs of applications for peta/exascale systems;
- Explore how laboratories, universities, and vendors can work together on coordinated HPC software;
- Explore governance structure and management models for IESP;
- Explore IP/Open Source issues.

The morning session of the first day of the workshop was devoted to updates from the international regions: US, Japan, China and Europe. In the afternoon of the first day, focus has been on the vendor activities and on the co-design proposals the Department of Energy (DoE) in the US is currently evaluating. Sections 4.2 to 4.4 will report on these activities. The second day of the workshop has been organised in discussion sessions in break-out groups, initially focused on hardware, software and applications. Section 4.5 will report on this.

4.2 Update on international efforts

4.2.1 US

Barbara Helland from the Office of Science of the DoE presented the current efforts of the DoE on exascale computing. With respect to the funding of Exascale initiatives in the US, the DoE is active in calls for proposals in various exascale-related areas:

- Applied Mathematics: Uncertainty Quantification (6 proposals funded at $3M/yr);
- Computer Science: Advanced Architectures (6 funded at $5M/yr);
- Computer Science: X-Stack (11 funded at $8.5M/yr);
- Computational Partnerships: Co-Design (21 Proposals requesting ~ $160M/year, decisions expected early 2011).

It has turned out that there is quite some interest in these calls, as the cumulated value of the proposals is much larger that the amount of funding available. Further calls and funding are under development, both from the X-stack perspective and the co-design perspectives. Planning efforts are underway with respect to an exascale software center and five co-design centers. The initial proposals are there, refinement will need to take place, reviewing is expected in spring 2011.

Beckman continued with presenting the planning effort for the aforementioned Exascale Software Center (ESC), with the following initial scope:

- Deliver high quality system software for exascale platforms ~2015, ~2018;
- Identify software gaps, research & develop solutions, test and support deployment;
- Increase the productivity and capability and reduce the risk of exascale deployments;

- Cost:
  - Applied R&D: ~10-20 distributed teams of 3 to 7 people each;
  - Large, primarily centralized Quality Assurance, integration, and verification center;
- Schedule Overview
  - 2010 – Q1 2011: Planning and technical reviews;
  - April 2011: Launch;
  - 2014, 2017: SW ready for integration for 2015, 2018 systems respectively.

Apart from important organizational aspects (like quality assurance and interfacing of software components), the principle of co-design of hardware, taking applications into account during the design process, is regarded as an important way forward. As described above, the DoE is reviewing proposals that apply for co-design centres. The co-design process is where system architects, application software designers, applied mathematicians, and computer scientists work closely together to produce a computational science discovery environment that fully leverages these significant advances in computational capability.

Many important aspects have been identified already, one of the most important is the relation (to be formalized) between the software developers and the vendors, with respect to development, support, risk and acceptance.

It is also recognized that expertise on certain areas of the software stack is concentrated in certain institutes in the US, which makes distributed project staffing a sensible option. Each such “component team” must have a certain mass of development staff, but also with quality assurance and testing manpower available.

4.2.2 Japan

Satoshi Matsuoka and Yutaka Ishikawa described the HPC challenges in Japan, and basically focused on two major aspects:

- The Next-Generation Supercomputer Project

  After the evaluation of the Riken project, the Next-Generation Supercomputer became known as the K (KEI) computer, targeted at 10 Pflop/s to become available during the course of 2012. With more than 640k cores in 80k nodes, more than 1 PB of main memory, stored in 800+ cabinets, the system is massive. Full installation is scheduled for 2011, fine-tuning and further improvements for 2012. Processor cores are based on Fujitisu’s SPARC-8FX technology. The system will be installed in Kobe, around 450 km west of Tokyo.

  Apart from the K-system, other Pflop/s facilities are expected to become available, starting at end 2010, early 2011 (Tokyo Institute of Technology, University of Tsukuba, University of Tokyo and Kyoto University, together T3K alliance). The Tsukuba system will focus on (green) GPU-technology

- The strategic program to promote HPC activities in Japan

  During the third Science and Technology Basic Plan (2006-2010), strategic applications for Japan have been identified for execution on the high-end systems:

  - Life sciences and drug development;
  - Materials science and energy;
  - Global climate change;
  - Engineering and manufacturing;
  - Origin of matter and universe.

  In order to support these research areas, the High Performance Computing Infrastructure (HPCI) has been created. The K-facility at Kobe will act as a centerpiece for HPCI. Many organisations

have already joined HPCI, ranging from research institutes to technology providers, but also to industrial organisations which can apply for compute cycles on the systems. The goals for HPCI are:

- Move from fundamental research levels to actual applications both in academic and industrial sectors;
- Create new mechanisms to let experimentalists use simulations as daily research tools;
- Enable collaboration between computational scientists and computer scientists for effective use of many-node systems.

With respect to further plans, it is expected that there will be around 30 M$ available for system software oriented projects, which will have a runtime of 5 to 7 years. These plans towards exascale and exaflops are under discussion in the fourth Science and Technology Basic Plan (2011-2015).

Finally, Japan and France are funding collaborative calls, including software and algorithmic aspects of HPC. Japan also takes part in the G8 calls, for which many projects have been proposed and which are in the second stage of selection.

4.2.3 China

Xue-bin Chi described the development of HPC in China. After a brief history of the computing resources at the SCCAS (Supercomputing Center of the Chinese Academy of Sciences), details on algorithmic development and applications were given, together with a strong increase in the amount of users of the SCCAS computing systems. Dr. Chi mentioned climate applications, radar models, cosmology, engineering applications in the car manufacturing industry. Not only emphasis is given on standard compute clusters, but also on implementation of applications on GPU-clusters. Several GPU-clusters, each with several hundreds of Tflop/s peak performance are operational at SCCAS.

Since 2001, China operates the China National grid, which connects universities and scientific institutes throughout the country. Currently, the gird is mainly composed of Lenovo and Dawning systems, reaching peak performances in the Pflop/s range. As of November 2010, China currently has two systems which have a LINPACK performance of more than one Petaflop/s. The first one is located in Tianjin, the second one at the Dawning facilities in Shenzen. Both systems are equipped with dual-processor 6-core Intel nodes, extended with Nvidia Fermi GPU cards.

For the period 2011-2015, China expects to operate many petascale systems, including a top system with 50-100 Pflop/s peak performance. In the period 2016-2020, a system in the range of 1 to 10 Exaflop/s is expected. These systems may use China's own development on its low-power processor, based on MIPS-compatible CPUs. The processor, known as Loongson, Godson or Dragon, will be developed at the Institute of Computing Technology, Chinese Academy of Sciences.

4.2.4 Europe

The situation in Europe was covered in three presentations: Berthou, Aerts and Johnson.

Berthou kicked off on European initiatives, the EC vision and funding. He started with an overview of the European Exascale Software Initiative (EESI), which coordinates the European participation in IESP. An important aspect is the concept of Working Groups (WG's), four of which deal with scientific application areas and four of which deal with hardware, system software, numerical algorithms and software engineering. The expected output of EESI is a roadmap and a set of recommendations to the funding agencies, shared by the European HPC community, on software (tools methods and applications) to be developed for exascale supercomputers. For the EC, such output must provide an analysis of European strengths and weaknesses, the willingness of European stakeholders to build exascale systems, in relation to the recently issued IDC report.[2]

[2] "Development of a Supercomputing Strategy in Europe", IDC, October 2010.

Berthou mentioned the EC initiatives on exascale computing in the FP7 program (described in detail by Aerts), and the collaboration with Russia for a total of 6 M€.

With respect to exascale research (hardware and software), there are some initiatives in European countries between universities/research labs/computing centers on one side and industry on the other side. A brief overview:

- Exascale Innovation Center, Germany, JSC (Juelich) and IBM;
- EX@TEC, France, CEA, GENCI, UVSQ and Intel;
- Flanders ExaScience Lab, Belgium, IMEC, five Flemish universities and Intel;
- Exascale Stream Computing Collaboratory, Ireland, IRCSET, four universities and IBM;
- Exascale technology Centre, UK, EPCC and Cray

Berthou concluded with a brief overview of actual projects, installations and plans in the individual European countries.

Aerts continued with an overview of the current EC activities in the FP7 framework program. During his presentation, he raised the question on reciprocity: what should the EC expect from the international cooperations in the exascale domain ? Aerts illustrated this question with the observation that for instance processor design is not a European activity, but that details must be shared with Europe as well to design efficient software.

- Call for Computing Systems

  This 45 M€ call is targeting projects to start mid 2011 and to end in 2014. Most of the call (40 M€) is dedicated to research projects, for which the research vision has been laid out in the European Network of Excellence on High Performance and Embedded Architecture and Compilation. The research topics include:

  - Parallel and concurrent computing;
  - Virtualisation;
  - Customisation;
  - Architecture and technology.

  The expected impacts are:

  - Improved programmability of future systems;
  - Efficient and ubiquitous use of virtualisation for heterogeneous multi-core systems;
  - Accelerated system development and production;
  - Reinforced European excellence in multi-core computing architectures, system software and tools;
  - Strengthened European leadership in cross-cutting technologies.

  The call closes on January 18, 2011.

- Call for exascale computing, software and simulation

  This 25 M€ call is dedicated specifically to exascale computing, which marks the commitment of the EC to support research at the leading edge of HPC. Most of the (24 M€) is dedicated to research projects, which are expected to be collaborations between one or more supercomputing centers, technology and system suppliers and industrial and academic centers. The intention of the call is:

  - To develop a small number of advanced computing platforms (100 petaflop/s in 2014 with potential for exascale by 2020), platforms relying on vendors' proprietary hardware or on COTS hardware;

- To develop optimised application codes driven by the computational needs of science and engineering and of today's grand challenges (e.g. climate change, energy, etc.);
- To address major challenges of extreme parallelism with millions of cores (programming models, compilers, performance analysis, algorithms, power consumption ...).

The expected impact is:

- Put Europe in the frontline of international efforts for the development of HPC system software and tools;
- Strengthen European industry supplying and operating HPC systems: preparing European industry and research organisations to achieve world-leadership in this area;
- European excellence in exascale level simulation codes for the benefit of society, industrial competitiveness and policy making; emergence of EU top-class simulation centres for exascale systems;
- Reinforce cooperation in international endeavours on exascale software and systems.

It is further expected that proposals:

- develop software as open source;
- split the effort roughly 40/60 in applications and simulation vs. systems development;
- demonstrate synergies with efforts under the Capacities programme on the deployment of leadership-class HPC systems;
- include international cooperation components that are essential and complementary to European expertise.

It is clear that international collaboration is needed for implementation of the Exascale Software Roadmap. Not only scientific issues are challenging, also organisational issues like the coordination of efforts, the governance model and the collaboration with vendors and Exascale Software Centers in the US and Asia, etc., are subject for investigation.

This call also closes on January 18, 2011.

Johnson presented the activities in the PRACE project and in the PRACE legal entity, which is responsible for granting access to the PRACE Tier-0 systems. Currently, there is one PRACE Tier-0 system (1 Pflop/s IBM BlueGene/P at JSC in Juelich, Germany). An interim Board of Directors, representing the hosting partners (Germany, France, Spain and Italy), with one representative of the non-hosting partners, has organised the first access calls, including the peer review. The legal entity (AISBL in Belgium) is governed by a council, consisting of all partners (currently 20 countries). A Scientific Steering Committee (SSC) advises the council, an Access Committee (AC) will be recruited from the SSC and will lead the process of assignment of computing time on the Tier-0 system(s). The following schedule of Tier-0 systems is foreseen:

- IBM BlueGene/P at Gauss Center for Supercomputing (GCS), Juelich, Germany;
- Bull system, Intel-based, GENCI, France (initial installation by end 2010, full installation by end 2011);
- 3rd PRACE system at HLRS, Stuttgart, Germany, 2011;
- 4th PRACE system at LRZ, Munich, Germany, 2012;
- 5th and 6th PRACE systems in Italy and Spain, 2013.

The EC intends to finance four PRACE projects: one preparatory phase (has run from January 2008 until June 2010), and three implementation phases (first has started on July 1, 2010). The EC is expected to contribute 70 M€ in total. Partners will match the EC contributions up to 100% maximum. The PRACE projects focus on several aspects:

- Set-up and operational model of the Research Infrastructure;
- Operational aspects, best practices;
- Application usage, optimisation, parallelisation and scalability;

- Benchmarking;
- Future hardware and software developments, including prototyping and novel programming languages;
- Dissemination, education and training;

An important aspect for PRACE is to watch technology and component developments. Since this may interfere with procurement cycles of Tier-0 systems, STRATOS (STRAtegic TechnOlogieS) has been established, which will be able to communicate with vendors on these aspects without being hindered by procurement rules.

4.2.5 Example high-end research infrastructure: SKA

Tim Cornwell presented the Square Kilometre Array (SKA) as an interesting project for co-design of future hardware and software environments. SKA is a 2020 era radio telescope, based on the principle of the Lofar telescope in the Netherlands. SKA will be sited in Australia or South Africa, and will contain 3600 antennas. The Lofar project already uses huge amounts of compute cycles to interpret the received signals, which will need to be order of magnitude more for SKA. It is for this reason that co-design of hardware and software, given the requirements and technical specifications of SKA, could be beneficial.

SKA is expected to need around 100 Tflop/s compute capacity and 100 PB storage capaity a year by 2012, increasing to 1 Pflop/s by 2014. requirements are expected to grow into the exascale area before 2020.

4.3 Vendor perspectives on exascale and co-design

4.3.1 Cray

Peg Williams (Cray) mentioned the key challenges for exascale, together with the focus areas for Cray:

- Power (system infrastructure, network, heterogeneous processors);
- System software (OS, file system scalability, jitter reduction);
- Programming systems (libraries, locality, compilers, tools, languages);
- Resiliency (system, application).

Cray has currently two exascale research initiatives in Europe: at EPCC in Edinburgh, UK, and at CSCS in Lugano, Switzerland. Apart from software life cycle management and the availability of test platforms and test strategies, an important topic for collaboration between vendors and the open source community is the topic of vendor differentiation and IP protection. Definition of Application Programming Interfaces (APIs) could be a valuable vehicle to enable vendors to differentiate without giving full details on their plans.

4.3.2 Fujitsu

Ross Nobes (Fujitsu Laboratories in Europe) presented Fujitsu's activities in capability computing. Fujitsu works with RIKEN on the K-computer in Kobe, Japan. Fujitist looks at:

- SPARC64 processor development;
- 6D mesh/torus interconnects;
- Direct water cooling packaging;
- Programming model: users should not worry on multiple cores on a single chip, Fujitsu's vectorisation experience in compilers will solve that;
- Hybrid task/thread model.

Open initiatives:

- Exascale Application and Data Initiative – XcalableMP: directive-based language extension to avoid using MPI;
- Open Petascale Libraries Project – global collaboration to develop advanced numerical software for supercomputing.

4.3.3 IBM

Robert Wisniewski (IBM) explained IBM's activities on exascale. Internally, there are a lot of projects going on for a few years already, covering both hardware and software developments. Full software ecosystem should be addressed: programming models, performance tools, OS. All aspects need to be re-designed with scalability in mind. A typical example are atomic operations, for which the time to synchronize all threads increases with the number of threads. General messages:

- Exascale software delivery will have challenges, but is tractable;
- Co-design and collaborative development are key components;
- Existing programming models will need to migrate;
- Open Source and community provided software will play a significant role – find models for:
    - Vendor product schedules and delivery;
    - Platform-enabling software must remain proprietary;
    - Define Open Source contributions

4.3.4 Intel

David Lombard (Intel) explains the investments Intel has done in various collaborations on exascale software (Juelich, Leuven, Paris). Various aspects will be covered in these labs. As Intel is a hardware vendor, focus is on hardware enhancements:

- Extreme voltage scaling to reduce core power;
- More parallelism 10x – 100x to achieve speed;
- Re-architecting DRAM to reduce memory power;
- New interconnect to lower power and distance;
- Non-volatile memory to rduce disk power and accesses;
- Resilient design to manage unreliable transistors.

Co-design is viewed as essential, just as collaboration with the Open Source Community, although new approaches will likely be required.

4.3.5 Microsoft

Dan Reed (Microsoft) did not focus on technical challenges, but much more on challenges in business model. Computing is already everywhere, and will increasingly be important in everyday life. Think of all kinds of devices, data and information, but also more and more natural user interfaces. The key question is whether exascale software development for HPC-type of applications will be such that it can be leveraged by the much broader and commercially much more attractive commodity, everyday software environments.

4.3.6 HP

Rob Schreiber (Hewlett-Packard) confirmed that HP is committed to pursue exascale systems for HPC. He mentioned the fact that HP labs is investigating several new technologies, which will be crucial to get to exascale systems: photonics, memory technology (non-volatile, memristor). In the

software area, HP predicts a huge growth of the needs in the commercial arena. Various exascale aspects will need to be solved there as well.

4.3.7 Bull

Jean-Francois Lavignon (Bull) stated that the road to exascale is more of a journey than of a real roadmap, given all challenges to be solved underway. These challenges are not only technological, but also from a business point of view, for instance how to reuse or apply HPC developments for exascale in the more commodity areas. Co-design as a hardware, application, algorithm and software collaboration may help (an example is Bull's Extreme Computing lab, with CEA and others), but no valid business model yet seen for doing this based on a single (set of) application(s). Bull's software strategy is based on Open Source plus added value provided by Bull experts.

4.4 Co-design planning efforts

The Office of Advanced Scientific Computing Research (ASCR) of the Office of Science (SC), U.S. Department of Energy (DoE) has announced a call for proposals from integrated teams of scientific researchers, applied mathematicians, computer scientists and computer architects with the goal of setting up so-called Exascale Co-Design Centers. Co-design refers to a computer system design process where scientific problem requirements influence architecture design and technology and constraints inform formulation and design of algorithms and software. To ensure that future architectures are well-suited for DoE target applications and that major DoE scientific problems can take advantage of the emerging computer architectures, major ongoing research and development centers of computational science need to be formally engaged in the hardware, software, numeric methods, algorithms, and applications co-design process that will be responsible for making key tradeoffs in the design of exascale systems.

At the time of the Maui meeting, the DoE was in the process of evaluating and ranking the received proposals. As said in section 4.2.1, 21 proposals for a total of 160 M$ had been received.

During the Maui meeting, five co-design centers presented themselves. Their presentations have not been made publicly available, as a result of the fact that the DoE is still evaluating the proposals. The presenters have been:

- Robert Harrison, Oak Ridge National Laboratory (ORNL), on the proposal for "Chemistry Exascale Co-design Center" (CECC);
- Andrew Siegel, Argonne National Laboratory (ANL), on the proposal for "Center for Exascale Simulation of Advanced Reactors" (CESAR);
- Alice Koniges, NERSC/Lawrence Berkeley Laboratory, on the co-design proposal for fusion;
- Sriram Swaminarayan, Los Alamos National Laboratory (LANL), on the co-design proposal for materials science;
- Carlo Graziani, The University of Chicago, on the co-desihn proposal for high energy density physics.

4.5 Break-outs on hardware/software/applications

The second day of the workshop focused on the three pillars to get to efficient exascale computing: hardware, software and applications. The discussions have been set up around the following questions:

- What do the software people need from the application and hardware groups ?
- What do the applications people need from the hardware and software groups ?
- What do the hardware people need from the software and application groups ?

In practice, the software and applications group merged, so there were two parallel sessions: hardware and software/applications.

Summary of hardware break-out group

Key question: What do the hardware people need from the software/applications group ?

The key constraints for hardware development and production for exascale computing are power budget, procurement budget and delivery schedules. These constraints get back when considering design trade-offs for the nodes, interconnect, I/O and accelerator hardware. Needed from the software/applications group are answers/insights:

- With respect to node design:
    - Heterogeneous nodes: ratio heavy-weight nodes and light-weight nodes;
    - Threads: importance of single-thread performance and the amount of threads effectively used per coherence domain;
    - Data movement: separate data mover development;
    - Memory hierarchy (cache sizes and hierarchy);
    - Memory design (non-volatile): what about capacity versus latency;
    - Power management of cores, memory, …;
    - Hardware features for resiliency;
- With respect to interconnect:
    - Relative importance of bandwidth, latency and messaging rate;
    - Evaluation of network technologies – how;
    - Importance of inter-node communication with respect to performance;
- With respect to I/O:
    - What I/O rates do we need;
    - New approaches to data analysis and resiliency.

General questions include:

- What types of acceleration can be put in hardware ?
- What hierarchical and hybrid programming models need to be supported ?
- Will there be a standard and stable API to access new features ? Who is responsible ?
- Van HW development use simulation/emulation tools to rapidly characterize application behavior ?

Summary of applications break-out group

Key question: What do the software/applications people need from the hardware group ?

This summary is structured through the question “What Co-Design Vehicles want”. The break-out group not only discussed what answers were needed form the hardware group, but also posed questions to themselves. There have been 21 candidate CDV’s identified, which probably do not represent yet all computing areas, but are a good starting point.

Current dominant findings on applications:

- Most applications are bulk synchronous, MPI-based, some kind of decomposition (domain, particle, …), but not all of them (e.g. electronic structure codes);
- Already running hybrid MPI/OpenMP, OpenMP to several tens of cores;
- Weak-scalability can be done efficiently on current largest systems (300k cores);
- Majority of CDV’s memory bandwidth limited;
- Flop/s per bye of storage is linear or log-linear in weak scaling;
- Intensive I/O in most cases only during start-up and finish, although check-pointing may become more important;

- Global collectives on synchronisation required (time steps, implicit linear algebra);

Algorithms and library dependencies:

- PDEs: equilibrium (implicit) and evolution (explicit);
- FFT;
- FMM;
- Particles pushing;
- Adaptive mesh refinement;
- Sparse and dense linear algebra;
- ODE integrators;
- MPI, GlobalArrays, GasNet;
- ParMetis;
- BLAS, ScaLAPACK;
- UMFPACK, SuperLU, MUMPS;
- PETSc, hypre, Trilinos, SUNDIASL;
- Chombo, SAMRA;
- FFTW;
- GraphLib;
- VisIt, VTK;
- Triangle;
- PALM;
- SILO, ADIOS, HDF5;
- BOOST.

Applications wish list for hardware/software:

- Programming models that optionally expose machine characteristics such as memory hierarchy, cores, etc.
- Memory related
    - Exposure of memory hierarchy;
    - Hardware gather-scatter operations;
    - Programmable caches (not just DMA engines: it would be nice to specify a memory access pattern or to mark regions of memory as write rarely/read often;
    - Access to DMA;
- Compiler related:
    - Single source across different implementations of exascale;
    - Auto vectorisation;
    - Masked SIMD operations;
- Power related:
    - Ability to power on/off pieces of hardware;
    - Ability to hint on power patterns to hardware;
- Tools:
    - Tools for early adopters vs. Tools for general use;
    - Tools for hot spot analysis and bottleneck identification;
    - Hardware counters available to users.
- Implementations of exascale:
    - Just GigaHertz*MegaNode*KiloCore ?
    - Or also other implementations: heterogeneous, GPU, accelerators.

Questions to the software/application group themselves:

- Size of the coherence domain;
- Exploring new programming models on the following aspects:
    - Manycore and heterogeneous nodes;
    - Data locality enhancements;
    - More loose synchronisation models;
    - Data movement becomes relatively more expensive compared to computing;
    - Mixed-precision algorithms;
- Fault tolerance:
    - Categorise type of faults, and identify there severity;
- Reproducibility:
    - Is bit-level reproducibility required for applications results ?
    - What would be the hardware/power advantages when bit-level reproducibility would be relaxed ?
- Big Data:
    - What would the availability of hundreds of TB to tens of PB main memory mean for application design ?

Other aspects:

- Investigate scalability on three levels: weak (one instance of application), strong (one instance of application), workflow scalability (across instances);
- Load balancing becomes extremely critical when using exascale sizes;
- Collective communication aspects;
- Hierarchical algorithms (like multigrid) may suffer on coarse levels.

4.6 Next meeting

The next IESP workshop meeting has been planned in Kyoto, Japan, April 6-7, 2011.

4.7 Conclusions

The Maui IESP meeting was the fifth workshop in a row which started in Santa Fe in April 2009. Building on the findings of the earlier workshops, the discussions during the meeting in Maui focused on the start of the practical implementation of Co-Design Vehicles (CDV's). The main part of the second day has been used for trying to identify as clearly as possible which kind of information was needed for the successful collaboration between application, hardware and software groups. This information has been characterized and should be used for the implementation of Co-Design Centers.

With respect to these Co-Design Centers, the DoE in the US has taken steps forward in selecting Co-Design Centers. Calls for proposals have been made, leading to 21 candidates, which are in the process of being reviewed. It seems that the US has taken the lead here.

With respect to setting up an Exascale Software Center (ESC), there are planning efforts underway, which foresee in delivering high quality system software for exascale platforms in the 2015-2018 range, to enable increased productivity and capability and reduce the risk of exascale deployments. The initial staffing effort of the ESC is anticipated as 10-20 distributed teams of 3 to 7 people each for applied Research & Development, and additional staff to ensure quality assurance and verification. The launch is foreseen in the first half of 2011.

Another important topic which has not been discussed yet in detail is the relationship between vendors on one side and applications/software groups on the other side. There must be a model in which applications/software groups are able to anticipate on new technology developments at vendors, which may be confidential and/or strategic for those vendors. A clearly defined Application Programming Interface (API), including responsibilities on maintaining will be required. An additional

dimension to this issue may be the aspect of national security: many vendors are US-based, which could mean limited access to application/software groups outside the US.

After this fifth workshop it has become apparent that IESP is able to build a broad community that is prepared and able to mobilize the stakeholders and others concerned to achieve the goals set forth by the IESP and prepare for a large scale redesign effort for software/applications in close cooperation with technology developers. It is also clear that the US strongly will go ahead whatever happens outside of the US. Only if the EC realises its own stakes in this matter (as up to now seems to be the case if one looks at the present FP7 and IDC efforts), Europe will be ready to deploy next generation HPC tools and equipment to Europe's own case and advantage.

5. References

[1] *The International Exascale Software Project: A Call to Cooperative Action by the Global High Performance Community*, Dongarra, J., Beckman, P., Aerts, P., Cappello, F., Lippert, T., Matsuoka, S., Messina, P., Moore, T., Stevens, R., Trefethen, A., Valero, M. Volume 23, Number 4, Winter 2009, *International Journal of High Performance Computer Applications*, pp 309-322, ISSN 1094-3420.

[2] International Exascale Software Project Roadmap, version 1.1, www.exascale.org.

6. Appendix A – Attendees of the IESP workshops

6.1 Attendees of previous Workshops

SC08, Austin, TX, USA, November 17-20, 2008
Santa Fe, NM, USA, April 7-8, 2009
Paris, France, June 28-29, 2009
Tsukuba, Japan, October 19-21, 2009
Oxford, UK, April 12-14, 2010

		Affiliation	Country	SC08	Sante Fe	Paris	Japan	Oxford
Patrick	Aerts	NWO	NL		x	x	x	x
Giovanni	Aloisio	Euro-Mediterranean Centre for Climate Change	Italy			x	x	x
Dong	Ahn	LLNL	US	x				
Yutaka	Akiyama	Tokyo Tech	Japan				x	
Jean-Claude	Andre	CERFACS	France			x	x	x
Phil	Andrews	UT	US	x				
Mutsumi	Aoyagi	U Kyushu	Japan			x	x	
Mike	Ashworth	Daresbury	UK			x	x	x
Franck	Barbier	ANR	France			x		
Venkat	Balaji	Princeton	US					X
David	Barkai	Intel	US	x	x	x	x	x
Sanzio	Bassini	CINECA	Italy			x		
Kyriakos	Baxevanidis ●[3]	EC	EU			x		x
Pete	Beckman	ANL	US	x	x	x	x	x
Jean-Yves	Berthou	EDF	France	x	x	x	x	x
Richard	Blake	Daresbury	UK		x			x
Jay	Boisseau	TACC	US	x				
Taisuke	Boku	U of Tsukuba	Japan		x	x	x	x
Bertrand	Braunschweig	ANR	France		x	x	x	
Bill	Camp	Intel	US		x			
Franck	Cappello	INRIA	France	x	x	x	x	x
Charlie	Catlett	ANL	US					x
Ruay-Shiung	Chang	National Dong Hwa University	Taiwan					x
Barbara	Chapman	U of Houston	US		x	x	x	x
Xuebin	Chi	CAS	China				x	x
Alok	Choudhary	NWU	US		x	x	x	x
Iris	Christadler	LRZ	Germany			x		
Almadena	Chtchelkanova	NSF	US		x			

[3] Attendees marked with "•" are official observers.

			Affiliation	Country	SC08	Sante Fe	Paris	Japan	Oxford
Guillaume	Colin de Verdière		CEA	France			x		
Frederica	Darema		NSF	US		x			
Bronis	de Supinski		LLNL	US	x				
Peter	Coveney		U College London	UK					x
David	Dean	●	ORNL/DOE	* US				x	x
Mik	Dewer		NAG	UK					x
Jack	Dongarra		U of Tennessee	US	x	x	x	x	x
Sudip	Dosanjh		SNL	US		x	x	x	x
Thom	Dunning		NCSA	US	x				
Iain	Duff		Rutherford Lab	UK					x
Hugo	Falter		ParTec	Germany		x	x	x	
Teresa	Finchum		U of Tennessee	US				x	x
Leonardo	Flores	●	EC	EU					x
Fabrizio	Gagliardi		Microsoft	US			x		
Alan	Gara		IBM	US		x	x		
Al	Geist		ORNL	US		x			
Luc	Giraud		CERFACS	France		x	x		
Kostas	Glinos	●	EC	EU			x		x
Jean	Gonnord		CEA	France	x		x		x
Robert	Graybill		ISI	US		x			
Bill	Gropp		UIUC	US	x		x	x	x
Jim	Hack		ORNL	US	x				
Jean-Francois	Hamelin		EDF	France		x		x	
Robert	Harrison		ORNL	US				x	x
Bill	Harrod		DAPRA	US	x				
Stefan	Heinzel		Max Planck DEISA	Germany		x	x	x	x
Barb	Helland	●	OS	US	x				x
Mike	Heroux		Sandia	US	x	x	x	x	x
Ryutaro	Himeno		RIKEN	Japan		x		x	x
Kimihiko	Hirao		Riken	Japan				x	
Dan	Hitchcock		OS	US	x				
Thuc	Hoang		NNSA	US		x			
Adolfy	Hoisie		LANL	US				x	
Charlie	Holland		DARPA	US	x				
Koh	Hotta		Fujitsu	Japan				x	x
Herbert	Huber		LRZ	Germany					x
Yuichi	Inoue		MEXT	Japan				x	
Yutaka	Ishikawa		U of Tokyo	Japan		x	x	x	x
Satoshi	Itoh		MEXT	Japan				x	
William	Jalby		U of Versailles	France		x			
Jean-Pascal	Jégu		Teratec	France			x		
Zhong	Jin		CAS	China				x	x

		Affiliation	Country	SC08	Sante Fe	Paris	Japan	Oxford
Fred	Johnson	DOE	US	x		x	x	x
Andrew	Jones	NAG	UK			x	x	
Emma	Jones	EPSRC	UK					x
Laxmilkant	Kale	UIUC	US			x		
Richard	Kenway	EPCC	UK		x		x	x
David	Keyes	Columbia U.	US		x	x		x
Moe	Khaleel	PPNL	US				x	
Kimmo	Koski	CSC	Finland			x		
Bill	Kramer	NCSA	US		x	x	x	x
Dimitri	Kusnezov	NNSA	US	x				
Jesus	Labarta	BSC	Spain		x	x	x	x
Jean-Francois	Lavignon	Bull	France		x	x	x	x
Alain	Lichnewsky	Genci	France		x	x		x
Volker	Lindenstruth	Heidelberg U	Germany			x		
Thomas	Lippert	Juelich	Germany	x		x	x	x
Bob	Lucas	ISI	US		x	x		x
Rusty	Lusk	ANL	US					x
Barney	Maccabe	ORNL	US		x	x	x	x
Satoshi	Matsuoka	TiTech	Japan	x	x	x	x	x
Simon	McIntosh-Smith	U. Bristol	UK					x
Bob	Meisner	NNAS	US	x				
Paul	Messina	ANL	US	x		x	x	
Peter	Michielse	NWO	NL		x	x		x
Kazunori	Mikami	Cray	Japan				x	
Leighanne	Mills	U of Tennessee	US				x	
Bernd	Mohr	Juelich	Germany		x	x	x	x
Terry	Moore	U of Tennessee	US	x	x	x	x	x
Hervé	Mouren	Teratec	France			x		
Jean-Michel	Muller	CNRS	France			x		
Matthias	Müller	Dresden	Germany				x	
Wolfgang	Nagel	Dresden	Germany		x	x	x	x
Kengo	Nakajima	U of Tokyo	Japan				x	x
Hiroshi	Nakashima	Kyoto U.	Japan			x	x	x
Mamoru	Nakono	Cray	Japan				X	
Jeff	Nichols	ORNL	US		x		x	x
Jane	Nicholson	EPSRC	UK				x	x
Jean-Philippe	Nominé	CEA	France			x		x
Nick	Nystrom	PSC	US		x			
Per	Oster	CSC	Finland	x	x			
Mike	Papka	ANL	US					x
Abani	Patra	NSF	US		x	x	x	
Rob	Pennington ●	NSF	US	x	x			x
Serge	Petiton	CNRS	France			x		x

			Affiliation	Country	SC08	Sante Fe	Paris	Japan	Oxford
Claude	Puech		INRIA	France		x	x		
Tracy	Rafferty		U of Tennessee	US	x	x	x	x	x
Dan	Reed		Microsoft	US		x	x		
Michael	Resch		HLRS Stuttgart	Germany		x			x
Catherine	Rivière		GENCI	France		x	x		
Ralph	Roskies		PSC	US	x				
Faith	Ruppert		ANL	US		x			
Christian	Saguez		Teratec	France			x		
Vivek	Sarkar		Rice	US		x			
Stef	Salvini		Oxford	UK					x
Mitsuhisa	Sato		U of Tsukuba	Japan	x	x	x	x	x
Stephen	Scott		ORNL	US		x			
Mark	Seager		LLNL	US		x			
Ed	Seidel		NSF	US	x		x	x	x
Akiyuki	Seki		MEXT	Japan				x	
Satoshi	Sekiguchi		AIST/METI	Japan				x	
Hideo	Sekino		Toyohash Inst Tech	Japan				x	
John	Shalf		LBNL	US			x	x	
Horst	Simon		LBNL	US	x	x			
David	Skinner		LBNL	US		x	x	x	x
Marc	Snir		UIUC	US	x				
Mary	Spada		ANL	US	x				
Thomas	Sterling		LSU	US		x	x	x	
James	Southern		Fujitsu	UK					x
Rick	Stevens		ANL	US	x	x		x	x
Michael	Strayer	●	DOE OS	US	x	x			
Fred	Streitz		LLNL	US				x	x
Bob	Sugar		UCSB	US				x	
Shinji	Sumimoto		Fujitsu	Japan				x	
Makoto	Taiji		Riken	Japan			x	x	x
Toshikazu	Takada		Riken	Japan				x	x
Hiroshi	Takemiya		JAEA	Japan					x
Bill	Tang		PPPL	US				x	x
John	Taylor		CSIRO	AU				x	x
Rajeev	Thakur		ANL	US				x	x
Anne	Trefethen		Oxford	UK		x	x	x	x
Akira	Ukawa		U of Tsukuba	Japan				x	
Mateo	Valero		BSC	Spain	x		x		x
Aad	van der Steen		NCF	NL				x	
Jeffrey	Vetter		ORNL	US		x	x	x	x
Vladimir	Voevodin		Moscow State U	Russia			x		
Andy	White		LANL	US	x	x			
Peg	Williams		Cray	US		x	x	x	x

		Affiliation	Country	SC08	Sante Fe	Paris	Japan	Oxford
Robert	Wisniewski	IBM	US				x	x
Felix	Wolf	Juelich	Germany					x
Kathy	Yelick	LBNL	US	x	x	x		
Akinori	Yonezawa	U Tokyo	Japan				x	
Thomas	Zacharia	ORNL	US	x				x

6.2 Maui, HI, USA, October 18-19, 2010

The attendees of the Maui IESP Workshop were:

Patrick Aerts	NWO	NL
Giovanni Aloisio	Euro-Mediterranea	Italy
Jean-Claude Andre	CERFACS	France
Michael Ashworth	Daresbury	UK
David Barkai	Intel	US
Pete Beckman	ANL	US
Jean-Yves Berthou	EDF	France
Taisuke Boku	U of Tsukuba	Japan
Bertrand Braunschweig	ANR	France
Ron Brightwell	SNL	US
Franck Cappello	INRIA	France
Charlie Catlett	ANL	US
Barbara Chapman	U of Houston	US
Xuebin Chi	CAS	China
Andrew Chien	Self	US
Alok Choudhary	NWU	US
Tim Cornwell	CSIRO	AU
Bronis de Supinski	LLNL	US
David Dean	ORNL/DOE	US
Jack Dongarra	U of Tennessee	US
Sudip Dosanjh	SNL	US
Stephane Ethier	PPPL	US
Hugo Falter	ParTec	Germany
Teresa Finchum	U of Tennessee	US
Karl Fuerlinger	LRZ	Germany
Al Geist	ORNL	US
Carlo Graziani	U Chicago	US
Bill Gropp	UIUC	US
Robert Harrison	ORNL	US
Stefan Heinzel	Max Planck/DEISA	Germany
Barb Helland	DOESC	US
Mike Heroux	Sandia	US
Ryutaro Himeno	RIKEN	Japan

Thuc Hoang	DOE	US
Adolfy Hoisie	PNNL	US
Koh Hotta	Fujitsu	Japan
Yutaka Ishikawa	U of Tokyo	Japan
Zhong Jin	CAS	China
Fred Johnson	DOE	US
Lennart Johnsson	UH	US
Larry Kaplan	Cray	US
David Keyes	Columbia U.	US
Moe Khaleel	PPNL	US
Alice Koniges	LBNL	US
Bill Kramer	NCSA	US
Jesus Labarta	BSC	Spain
Jean-Francois Lavignon	Bull	France
David Lombard	Intel	US
Bob Lucas	ISI	US
Rusty Lusk	ANL	US
Barney MacCabe	ORNL	US
Satoshi Matsuoka	TiTech	Japan
Bronson Messer	ORNL	US
Peter Michielse	NWO	NL
Bernd Mohr	Juelich	Germany
Terry Moore	U of Tennessee	US
Wolfgang Nagel	Dresden	Germany
Hiroshi Nakashima	U of Kyoto	Japan
Jeff Nichols	ORNL	US
Ross Nobes	Fujitsu	UK
Rob Pennington	NSF	US
Walt Polansky	DOESC	US
Tracy Rafferty	U of Tennessee	US
Dan Reed	Microsoft	US
Rob Ross	ANL	US
Faith Ruppert	ANL	US
Robert Schreiber	HP	US
Mark Seager	LLNL	US
John Shalf	LBNL	US
Andrew Siegel	ANL	US
Horst Simon	LBNL	US
David Skinner	LBNL	US
Thomas Sterling	LSU	US
Rick Stevens	ANL	US
Fred Streitz	LLNL	US
Sriram Swaminarayan	LANL	US
Makoto Taiji	Riken	Japan
John Taylor	CSIRO	AU
Rajeev Thakur	ANL	US
Jeffrey Vetter	ORNL	US

Andy White	LANL	US
Peg Williams	Cray	US
Marcus Wilms	DFG	Germany
Bob Wisniewski	IBM	US