H2020-FETHPC-2014

Coordination of the HPC strategy

EXDCI

European eXtreme Data and Computing Initiative

Grant Agreement Number: FETHPC-671558

D6.4
Report on the BDEC workshop (EU)

Final version

Version: 1.4
Author(s): Mark ASCH, ANR and Université de Picardie Jules Verne
Thierry BIDOT, Neovia Innovation
Corentin LEFEVRE, Neovia Innovation
Date: 25.08.2016

Project and Deliverable Information Sheet

EXDCI Project	**Project Ref. №: FETHPC-671558**	
	Project Title: European eXtreme Data and Computing Initiative	
	Project Web Site: http://www.exdci.eu	
	Deliverable ID: D6.4	
	Deliverable Nature: Report	
	Dissemination Level: PU	**Contractual Date of Delivery:** 31 / August / 2016
		Actual Date of Delivery: 30 / August / 2016
	EC Project Officer: Beatrice MARQUEZ-GARRIDO	

* - The dissemination level are indicated as follows: **PU** – Public, **CO** – Confidential, only for members of the consortium (including the Commission Services) **CL** – Classified, as referred to in Commission Decision 2991/844/EC.

Document Control Sheet

Document	**Title: Report on the BDEC Workshops (EU)**	
	ID: D6.4	
	Version: 1.4	**Status:** Final version
	Available at: http://www.exdci.eu	
	Software Tool: Microsoft Word 2016	
	Files: EXDCI_D6.4_Report on the BDEC Workshops (EU)_V1.4.docx	
Authorship	**Written by:**	Mark ASCH, ANR and Université de Picardie Jules Verne Thierry BIDOT, Neovia Innovation Corentin LEFEVRE, Neovia Innovation
	Contributors:	
	Reviewed by:	Renata GIMENEZ (BSC) Natale CURATOLO (PRACE)
	Approved by:	MB/TB

Document Status Sheet

Version	Date	Status	Comments
1.1	05/08/2016	Draft	Preliminary version
1.2	09/08/2016	Draft	First full version
1.3	11/08/2016	Draft	For internal review
1.4	25/08/2016	Final version	Final version

Document Keywords

Keywords:	ETP4HPC, PRACE, Exascale, Research Infrastructure, BDEC, Big Data, Convergence, Machine Learning, Numerical Laboratories, International Cooperation, Simulation

Table of Contents

List of Figures

References and Applicable Documents

[1] https://exdci.eu
[2] http://www.exascale.org/bdec/

List of Acronyms and Abbreviations

ANR	Agence Nationale de la Recherche (National Research Agency – France)
BDEC	Big Data and Extreme-scale Computing
BSC	Barcelona Supercomputing Center
CEA	Commissariat à l'Energie Atomique (Atomic Energy Commission – France)
DE	Deutschland (Germany)
EC	European Commission
EU	European Union
EXDCI	European eXtreme Data and Computing Initiative
HPC	High Performance Computing
INRIA	Institut National de la Recherche en Informatique et en Automatique (National Institute for Computer Sciences and Applied Mathematics – France)
IPCEI	Important Projects of Common European Interest
ISC	International Supercomputing Conference
LWK	Lightweight Kernel
NASA	National Aeronautics and Space Administration – USA
NSF	National Science Foundation - USA
PRACE	Partnership for Advanced Computing in Europe
RIKEN	Rikagaku Kenkyusho (Institute of Physical and Chemical Research – Japan)
US/USA	United States of America

Executive Summary

Since 2003, Big Data and Extreme-scale Computing workshops are organised as a world-wide international coordination between the HPC ecosystem and the Big Data community.

In 2016, the workshop was organised in Frankfurt (DE) under the supervision of the EXDCI project. 70 participants from 10 countries attended the conference.

Through global and inspirational talks (keynotes), state-of-the-art presentations (white papers on scientific cases, roadmaps on policies) and collective thinking (brainstorming sessions), the meeting achieved the following objectives: *(1) giving an overview of the current status of the convergence between HPC and Big Data, (2) determining the next challenges to tackle, (3) initiating the writing of a multi-authored, international document provisionally entitled "Pathways to Convergence"*.

1 Introduction

In the past years, the United States, the European Union and Japan have each moved aggressively to develop their own plans for achieving exascale computing in the next decade. Such concerted planning by the traditional leaders of HPC speaks eloquently about both the substantial rewards that await the success of such efforts, and about the unprecedented technical obstacles that block the path upward to get there.

While the exascale initiatives have understandably focused on the big challenges of exascale for hardware and software architecture, the relatively recent emergence of the phenomena of Big Data in a wide variety of scientific fields represents a tectonic shift that is transforming the entire research landscape on which all plans for exascale computing must play out.

Since 2013, the Big Data and Extreme-scale Computing (BDEC) workshops have been held to systematically map out and account for the ways in which the major issues associated with Big Data intersect with those of HPC and extreme scale computing.

These workshops have a triple role:

- **Coordination**. By sharing state of the art thoughts and projects, both within their scientific disciplines and the national or international frameworks, the participants aim to detect and/or initiate international coordination. As a side effect, this contributes to the dissemination of EU projects and policies.
- **Prospective**. Brainstorming is at the heart of the BDEC process. The goal is to identify the future challenges, breakthroughs and potential brakes between the HPC and Big Data communities. The result is to achieve a common view on scientific challenges to tackle.
- **Networking**. BDEC is a unique place where international researchers can meet and network.

In 2016, the BDEC workshop was held in Frankfurt, Germany, at the initiative of the EXDCI project.

During two days, international experts from 10 countries exchanged on the state of the art of the convergence between HPC and Big Data, and drew paths for the future through intensive brainstorming sessions.

Finally, main presentations and conclusions have been presented at ISC16 trough a dedicated open workshop session.

2 Setting and Agenda

The meeting leader was Mark ASCH from EXDCI. A specific group composed by EU and US representatives (Mark ASCH and Thierry BIDOT, Terence MOORE and Jack DONGARRA) supervised the meeting with an active participation of Asian colleagues.

The BDEC workshop took place in Frankfurt from the 15th to the 17th of June 2016, at the hotel NH Frankfurt Airport West. This location was chosen for its proximity with the International Supercomputing Conference (ISC), also held in Frankfurt, a week after the BDEC meeting.

A subcontractor (Barcelo Congresos), which already worked on the organisation of the BDEC meeting held in Barcelona in 2015, took partly in charge the logistic organisation of the meeting.
For the registration to the workshop, a web page was prepared at:
https://barcelo.eventsair.com/bdec16/registration/Site/Register
The registration form took into account special dietary requirements for the attendees.

Mark ASCH sent a call for papers by e-mail. All contributions went through a scientific committee.

The number of participants in the meeting was around 70, of which 8 were women. The names and affiliations of the participants can be found in the annex.

Figure 1 - BDEC group photo

Program of the workshop

The program of the workshop is given below. Compared with the original schedule, only some modifications on the orders of the speakers occurred.

Wednesday, June 15th, 2016	
6:00pm to 8:00pm	Reception and Registration
Thursday, June 16th, 2016	
INTRODUCTION AND KEYNOTES	
9:00am to 9:15am	**Welcome, Overview and Goals of the Meeting** Sergi Girona Turell, Partnership for Advanced Computing in Europe Mark Asch, French National Research Agency
9:15am to 10:00am	**Numerical Laboratories on Exascale**

	Alex Szalay, Johns Hopkins University, USA.
10:00am to 10:45am	**AI Platform and Challenges: Perspectives from AI Research Center (AIRC)** Junichi Tsujii, AIST, Japan.
10:45am to 11:15am	Coffee Break
BDEC PATHWAYS TO CONVERGENCE - FLASHES	
11:15am to 1:00pm	**White Paper Presentations - Session 1** 5-slide Presentations of Contributions, 10 Minutes per Presentation *Using an LWK in a Multi-Kernel Environment to Improve Containers* Balazs Gerofi, Yutaka Ishikawa, Rolf Riesen, Robert Wisniewski Presented by: Robert Wisniewski *OpenHPC: a Collaborative Environment for Development of a Cohesive Comprehensive HPC Software Stack Suitable for Cloud Integration* Robert Wisniewski *Pathways to Convergence* Osamu Tatebe *Integrative Multi-Scale Imaging, Simulation and Precision Cancer Therapy* Joel Saltz *Toward the converge of in-situ data analysis and deep-learning methods for efficient pre/post processing of pertinent structure among massive amount of scientific data* Sandro Fiore, Giovanni Aloisio, Philippe Ricoux, S. Brun, Jean-Michel Alimi, Mathis Bode, R. Apostolov, Stephane Requena Presented by: Jean-Michel Alimi *Names Matter: Eliminate the False Dichotomy Between Big Data and HPC* Dan Reed *Framework for Development of Data-Movement Centric Applications* Kengo Nakajima *Merging Big Data and HPC for Large-scale Analysis/Analytics at the NASA Advanced Supercomputing (NAS) Division* Piyush Mehrotra
1:00pm to 2:00pm	Lunch Break

PATHWAYS TO CONVERGENCE	
2:00pm to 4:00pm	**White Paper Presentations - Session 2** 5-slide Presentations of Contributions 10 Minutes per Presentation *Why convergence? A contrarian view and a path to convergence enabling specialization* Barney Maccabe *Current status and future prospects of optical communications technology and possible impact on future BDEC systems* Tomohiro Kudoh, Kiyo Ishii, Shu Namiki Presented by: Tomohiro Kudoh *Numerical Algorithms, Libraries, and Software Frameworks for Future HPC Systems (Towards the Post Moore Era)* Takeshi Iwashita *Big Data, Simulations and HPC Convergence* Geoffrey Fox, Judy Qiu, Shantenu Jha, Supun Kamburugamuve, Saliya Ekanayake Presented by: Geoffrey Fox *Big Data Analytics and High Performance Computing Convergence Through Workflows and Virtualization* Ewa Deelman *Scalable Compression as a Fundamental Pattern for HPC–Big Data Convergence* Franck Cappello, Sheng Di Presented by: Franck Cappello *On-Demand Data Analytics and Storage for Extreme-Scale Simulations and Experiments* Franck Cappello, Katrin Heitmann, Gabrielle Allen, Salman Habib, Ed Seidel, Brandon George, Brett Bode, Tim Boerner, Maxine D. Brown, Michelle Butler, Randal L Butler, Kenton G. McHenry, Athol J Kemball, Rajkumar Kettimuthu, Ravi Madduri, Alex Parga, Roberto R. Sisneros, Corby B. Schmitz, Sean R. Stevens, Matthew J. Turk, Tom Uram, David Wheeler, Michael J. Wilde, Justin M. Wozniak Presented by: Franck Cappello *Big Data for climate and air quality* Pierre-Antoine Bretonnière, Francesco Benincasa, Francisco Doblas-Reyes, Kim Serradell Presented by: Francesco Benincasa *Enablement of multi-scale simulation, analytics and visualization workflows*

	Marc Casas, Miquel Moretó, Rosa M. Badia, Raul Sirvent, Eduard Ayguade, Jesus Labarta, Mateo Valero Presented by: Rosa Badia *Toward large scale distributed experiments for climate change data analytics in the Earth System Grid Federation (ESGF) eco-system* Sandro Fiore, Dean Williams, Valentine Anantharaj, Sylvie Joussaume, Davide Salomoni, Stephane Requena, Giovanni Aloisio Presented by: Sandro Fiore *SAGE: Percipient Storage for Exascale Data Centric Computing* Malcolm Muggeridge, Sai Narasimhamurthy Presented by: Malcolm Muggeridge *Data Intensive and High Performance Computing: The View from High Energy Physics* Anshu Dubey, Salman Habib *CEA: Connecting Co-design Capabilities with the Operation of Large HPC and Data Processing Facilities* Edouard Audit, Christophe Calvin, Jean Gonnord, Jacques-Charles Lafoucrière, Jean-Philippe Nomine Presented by: Edouard Audit
4:00pm to 4:30pm	Coffee Break
BREAKOUT GROUPS	
4:30pm to 4:50pm	**Pathways to Convergence - Where Do We Go From Here?** Jack Dongarra, University of Tennessee, USA. Pete Beckman, Argonne National Laboratory, USA. Terry Moore, University of Tennessee, USA.
4:50pm to 6:30pm	Breakout **Breakouts on Pathways**
8:00pm	Dinner
Friday, June 17th, 2016	
ROADMAPS AND OVERVIEWS	
9:00am to 9:30am	**Convergence: What it Means to Me** Rick Stevens, Argonne National Laboratory, USA.
9:30am to 9:50am	**EU Overview** Jean Gonnord, Atomic Energy and Alternative Energies Commission (CEA, France)
9:50am to 10:10am	**USA-NSF Overview** Rajiv Ramnath, National Science Foundation, USA.
10:10am to 10:30am	**Japan Overview** Yutaka Ishikawa, RIKEN Advanced Institute for Computational

	Science, Japan. Satoshi Matsuoka, Tokyo Institute of Technology, Japan.
10:30am to 11:00am	**China Big Data and HPC Initiatives Overview** Xuanhua Shi, Huazhong University Of Science And Technology, China.
11:00am to 11:30am	**Data and Data-intensive computing challenges in Earth and Universe Sciences** Jean-Pierre Vilotte, Paris Institute of Earth Physics, France.
11:30am to 12:15pm	Breakout Report **Breakout Review and Reporting**
12:15pm to 1:30pm	Lunch Break
KEYNOTES 2	
1:30pm to 2:00pm	**Big Data and Extreme Scale Computing: a Storage-based Pathway to Convergence** Gabriel Antoniu, INRIA, France.
BREAKOUT GROUPS 2	
2:00pm to 4:00pm	Breakout **Breakouts on Pathways and Reporting**
4:00pm to 4:30pm	Coffee Break
BREAKOUT REPORTS AND CLOSING	
4:30pm to 5:30pm	Breakout Report **Breakout Reports** 4 x 15 Minutes
5:30pm to 5:45pm	**Conclusion and Next Steps**

The presentations were published online during the workshop on the dedicated website www.exascale.org thanks to the Argonne National Laboratory (ANL – US) team.

3 Report on Day 1

On the first day, the participants had a fresh view of the finished and upcoming projects in the BDEC field. Two keynotes opened the session with the purpose to give a global vision. Finally, the first brainstorming session was held with 4 breakout groups: Operations, Architecture, Software, Applications.

3.1 Keynotes #1

Alexander Szalay and Junichi Tsujii presented the introductory keynotes. The first presentation focused on the increased role of data in simulations, and, therefore, their computational needs.

Alex Szalay plead the idea of creating "numerical laboratories" as simulations are becoming more and more important in various scientific and industrial fields.
Not only should HPC and Data Analysis evolve together but, more importantly, innovative access patterns to this new kind of laboratories have to be found.

Junichi Tsujii, on his side, explored the increased role of Machine Learning and Artificial Intelligence (AI), as a third player between HPC and Big Data. Machine Learning is now a well-developed branch of computational sciences with wide range potential. Machine Learning is closely related to data exploitation, and the interaction between the two has consequences on computational design and infrastructure.
In this domain as well, new business model should flourish, especially regarding the ownership of data and should take into account that the current ecosystem has very different backgrounds, should it be in the US or in the EU/Japan.

Figure 2 - Junichi Tsijii concluding its keynote

3.2 White papers presentation

Twenty-one white papers were presented in two sessions, covering a wide range of topics and disciplines at the forefront of HPC and Big Data. At a glance, the presentations highlighted:

Big Data
All speakers emphasized the continuous increase of data for scientific purposes, and anticipated that the trends will continue.
Data explosion challenges the community in two particular ways:

- Storage;
- Access and Analysis.

Several solutions and experiments were presented on this side, including propositions to design joint infrastructures with HPC.
In particular

- On Data Storage, several projects and approaches complemented each other,
    - By working on compression strategy, analyse and elimination of "bad" data sets.
    - By working on "on the fly" or "intermediate" analysis, avoiding useless data storage.
- On Access and Analysis,
    - By proposing new models based on cloud technology and mutualisation ("Data Community", for example linked to already globalised research infrastructures).
    - By working on solutions to access data for multiple users.

HPC
The exascale challenge is reinforced by the increased complexity of scientific simulations, based on both the complexities of models and increased data availability.
Scalability is one of the most complex issues to address as well as the adaptation of the HPC centres to users' real needs.
Regarding the hardware part, the need for HPC to access the increased amount of data requires that the infrastructures are designed in common.

The convergence between HPC and Big Data will be restrained by specialisation and the address of specific communities' needs and/or history. However, as they are both integrated in the same process, their development will need:

- Co-design methods: one cannot define a system alone;
- Shared infrastructure: mutualisation at a global level;
- Cross-fertilisation: some disciplines have already worked on problems that may impact other scientific communities.

A newcomer: Machine Learning
Machine Learning has emerged during this BDEC session as a potential breakthrough and an important pillar of the development of HPC and Big Data.

3.3 First breakout sessions

The breakout sessions were around four main pillars: Operations, Architecture, Software and Applications.
The participants were randomly distributed into four groups. Each group participated to each subject, building upon the results of the earlier contributions.

Two spokespersons were committed to each group with two functions:

- Organise the discussion and the capitalisation.
- Organise and present the results at the end of the process.

Figure 3 - First breakout session

4 Report on Day 2

The morning of the second day of the workshop dedicated a moment to have a look at the upcoming policies of the major players in the field: EU, US, Japan and, notably China, which

once again maintained its leadership in the Top500 list. Then two keynotes gave another global view on convergence and prospectives between HPC and Big Data.
Ultimately, the participants collectively discussed the main challenges to tackle in the *pathways to convergence.*

4.1 Roadmaps presentations

The speakers respectively presented the current roadmaps of the EU, US, Japan and China.

A new EU initiative driven by the European Commission is set-up through the European Cloud Initiative, the European Open Science Cloud and the IPCEI project. The efforts on the European side are also coordinated through the EXDCI project with the reinforcement of the EU HPC ecosystem.
The ambition on the EU side is high, with a focus to reduce the market asymmetry.

Figure 4 - EU: a new ambition

The NSF in the US has adopted a user-centric viewpoint. The idea is to provide many interfaces and shared services for the scientists, within the development of a cyberinfrastructure, also called "cyberinfrastructure ecosystem". The National Strategic Computing Initiative (NSCI) drives the US strategy.
The Japanese roadmap focuses on the anticipation of the Post-Moore and Post-K era. Several projects will be supported, distributed among the well-established laboratories and research centres in Japan.
Chinese projects are under a roadmap for the coming years. They intend to reinforce their supercomputing capacity with two 100PFlops machines and they are already working on several projects combining Big Data and HPC challenges – such as Data Storage.

Each initiative is not equal in terms of budget and global cooperation intensity is also different when looking at the players involved (research infrastructures, industrials…).

4.2 Keynotes #2

At the beginning of the second day, Rick Stevens gave an inspirational talk with his analysis of the HPC/Big Data ecosystem.
Starting with a prediction that Machine Learning will change the HPC landscape, he analysed the concept of convergence between HPC and Big Data through their model and the actual alternative, which is co-location.

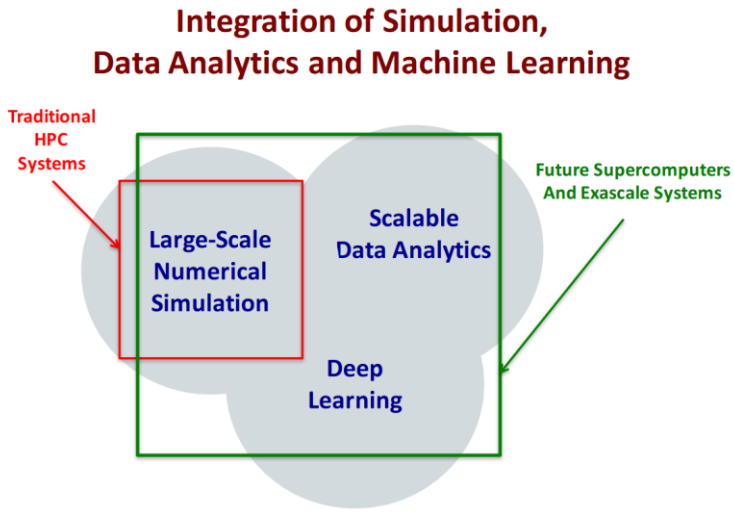

Figure 5 - Integration of simulation, data analytics and machine learning

The two models carry their own specifications, advantages or inconveniences. Rick Stevens highlighted that the current usage models are different in the Big Data and HPC ecosystem. If the goal is to connect these two worlds, new policies have to be set-up to create *a common ecosystem* between data and calculation capacity.
The creation of this kind of new ecosystem in itself (from co-location to convergence) requires solving complex problems such as data replication, user navigation of the resources, and management of complexity…
The conclusion is that convergence *may* be desirable but is not the only model to explore. If we have to do so, an incremental approach by learning from small number of projects focused on the next generation drivers is the main thing to do.

Jean-Pierre Vilotte took the examples of the Earth and Universe Sciences to assess the need of improving the current Big Data/HPC solutions, with an emphasis that the systems will be science-driven.
This implies that the scientists will gain control over the data (data storage) and that the HPC services are optimised.
Another development is the creation of numerical laboratories, an infrastructure at the forefront of convergence between Big Data and HPC. These laboratories will need the HPC to merge with Big Data in data-intensive platforms. Their linkage will depend on resolving such challenges as the movement of data in HPC systems.

The two keynotes mirrored their counterpart of the first day with Machine Learning and Numerical Laboratories.

4.3 Second breakout sessions

The session spokespersons synthetized the main ideas in post-its and summarized in a short oral presentation.
Then, all participant were invited to mix the post-its by creating clusters of challenges and concepts. The movement of creation was free and collective: the session only stopped when the group had reached a consensus.

Figure 6 - Second breakout session

At the end of the day, after an ultimate walkthrough of the ideas, people were asked to develop small chapters on the identified topics/challenges/clusters of ideas, in order to provide a more complete report to the community, the "Pathways to Convergence" document that will be presented (in draft form) at SuperComputing'16 in Salt Lake City, November 2016.

5 Conclusion of the meeting

Machine Learning and Simulation (E-science) are the key drivers for the convergence between the HPC ecosystem and the Big Data community.

This convergence may lead to multiple solutions, as specialisation and specific communities' needs will lead to develop and propose *ad hoc* infrastructures.

These new infrastructures – Numerical Laboratories or Cyberinfrastructure ecosystem must be user-driven and will need co-design approaches.

Building such infrastructure will require international cooperation and depend on the current structuration of given communities.

Data management (storage, access) is a current key-issue with undergoing projects proposing different approaches. The issue of access to and storage of ever increasing volumes of observation and sensor data is also important.

Integrative models between HPC, Big Data and the newcomer Machine Learning will require tackling new challenges and an integrative roadmap between each community.

Annex – Workshop participants

First Name	Last Name	Affiliation
James	Ahrens	Los Alamos National Laboratory
Jean-Michel	Alimi	Observatoire de Paris
Giovanni	Aloisio	Euro-Mediterranean Center on Climate Change / University of Salento
Jean-Claude	Andre	Jca Consultance & Analyse
Gabriel	Antoniu	INRIA
Mark	Asch	French National Research Agency
Edouard	Audit	Atomic Energy and Alternative Energies Commission (CEA)
Rosa	Badia	Barcelona Supercomputing Center
Pete	Beckman	Argonne National Laboratory
Costas	Bekas	IBM Research
Francesco	Benincasa	Barcelona Supercomputing Center
Thierry	Bidot	Neovia Innovation
Mathis	Bode	RWTH Aachen University
Francois	Bodin	Institute for Research in Computer Science and Random Systems / University of Rennes 1
Taisuke	Boku	University of Tsukuba
Christophe	Calvin	Atomic Energy and Alternative Energies Commission (CEA)
Franck	Cappello	Argonne National Laboratory / University of Illinois
Marc	Casas	Barcelona Supercomputing Center
Barbara	Chapman	Stony Brook University
Susumu	Date	Osaka University
Bronis	de Supinski	Lawrence Livermore National Laboratory
Ewa	Deelman	University of Southern California Information Sciences Institute
Jack	Dongarra	University of Tennessee
Sudip	Dosanjh	Lawrence Berkeley National Laboratory
Anshu	Dubey	Argonne National Laboratory
Giovanni	Erbacci	CINECA
Hugo	Falter	ParTec Cluster Competence Center
Sandro	Fiore	Euro-Mediterranean Center on Climate Change / University of Salento
Geoffrey	Fox	Indiana University
Fabrizio	Gagliardi	Barcelona Supercomputing Center
Adriano	Galano	Fujitsu
Sergi	Girona Turell	Partnership for Advanced Computing in Europe / Barcelona

First Name	Last Name	Affiliation
		Supercomputing Center
Jean	Gonnord	Atomic Energy and Alternative Energies Commission (CEA)
Jim	Harrell	Cray
William	Harrod	United States Department of Energy
Michael	Heroux	Sandia National Laboratories
Vasant	Honavar	Pennsylvania State University
Hans-Christian	Hoppe	Intel
Yutaka	Ishikawa	RIKEN Advanced Institute for Computational Science
Takeshi	Iwashita	Hokkaido University
Kate	Keahey	Argonne National Laboratory
Alison	Kennedy	Partnership for Advanced Computing in Europe
Jamie	Kinney	Amazon
Masaaki	Kondo	University of Tokyo
Tomohiro	Kudoh	National Institute of Advanced Industrial Science and Technology
Jesus	Labarta	Barcelona Supercomputing Center
Jean-Francois	Lavignon	Atos
Jysoo	Lee	Korean Institute of Science and Technology Information
Corentin	Lefevre	Neovia Innovation
Barney	Maccabe	Oak Ridge National Laboratory
Michael	Malms	IBM Research
Satoshi	Matsuoka	Tokyo Institute of Technology
Piyush	Mehrotra	NASA Ames Research Center
Marek	Michalewicz	Agency for Science, Technology, and Research (A*STAR)
Bernd	Mohr	Jülich Supercomputing Centre
Terry	Moore	University of Tennessee
Malcolm	Muggeridge	Seagate Technology, LLC
Kengo	Nakajima	University of Tokyo
Hiroshi	Nakashima	Kyoto University
Jean-Philippe	Nomine	Atomic Energy and Alternative Energies Commission (CEA)
Marcin	Ostasz	Barcelona Supercomputing Center
Tracy	Rafferty	University of Tennessee
Rajiv	Ramnath	National Science Foundation
Dan	Reed	University of Iowa
Joel	Saltz	Stony Brook University
Mitsuhisa	Sato	RIKEN Advanced Institute for Computational Science / University of Tsukuba
Thomas	Schulthess	ETH Zurich

First Name	Last Name	Affiliation
Satoshi	Sekiguchi	National Institute of Advanced Industrial Science and Technology
Xuanhua	Shi	Huazhong University Of Science And Technology
Rick	Stevens	Argonne National Laboratory
Martin	Swany	Indiana University
Alex	Szalay	Johns Hopkins University
Osamu	Tatebe	University of Tsukuba
Rajeev	Thakur	Argonne National Laboratory
Panagiotis	Tsarchopoulos	European Commission
Akira	Ukawa	RIKEN Advanced Institute for Computational Science
Jean-Pierre	Vilotte	Paris Institute of Earth Physics
Robert	Wisniewski	Intel
Kathy	Yelick	Lawrence Berkeley National Laboratory
Rio	Yokota	Tokyo Institute of Technology
Igor	Zacharov	Eurotech