H2020-FETHPC-3-2017 - Exascale HPC ecosystem development

EXDCI-2

European eXtreme Data and Computing Initiative - 2

Grant Agreement Number: 800957

D5.1
Report on the first international workshop

Version: Final
Authors: Mark Asch - U. de Picardie Jules Verne, Vera Ovcharenko - Neovia Innovation
Date: 30/09/2019

Project and Deliverable Information Sheet

<table>
<tr><td rowspan="6">EXDCI Project</td><td colspan="2">Project Ref. №: FETHPC-800957</td></tr>
<tr><td colspan="2">Project Title: European eXtreme Data and Computing Initiative - 2</td></tr>
<tr><td colspan="2">Project Web Site: http://www.exdci.eu</td></tr>
<tr><td colspan="2">Deliverable ID: D5.1</td></tr>
<tr><td colspan="2">Deliverable Nature: Report</td></tr>
<tr><td>Dissemination Level:
PU</td><td>Contractual Date of Delivery:
31 / August / 2019</td></tr>
<tr><td></td><td>Actual Date of Delivery:
23/ July / 2019
(Revised version resubmitted 30/09/2019)</td></tr>
<tr><td colspan="2">EC Project Officer: Evangelos Floros</td></tr>
</table>

* - 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

<table>
<tr><td rowspan="5">Document</td><td colspan="2">Title: Report on the first international workshop</td></tr>
<tr><td colspan="2">ID: D5.1</td></tr>
<tr><td>Version: 1.3</td><td>Status: Final</td></tr>
<tr><td colspan="2">Available at: http://www.exdci.eu</td></tr>
<tr><td colspan="2">Software Tool: Microsoft Word 2016</td></tr>
<tr><td colspan="3">File(s): EXDCI_D5.1_Report on the first International event (EU)_V1.3.docx</td></tr>
<tr><td rowspan="4">Authorship</td><td>Written by:</td><td>Mark Asch – Université de Picardie Jules Verne
Vera Ovcharenko – Neovia Innovation</td></tr>
<tr><td>Contributors:</td><td>Thierry Bidot, Neovia Innovation</td></tr>
<tr><td>Reviewed by:</td><td>M. Gillot (ETP4HPC), O. Beaumont (INRIA)</td></tr>
<tr><td>Approved by:</td><td>MB/TB</td></tr>
</table>

Document Status Sheet

Version	Date	Status	Comments
0.3	01/July/2019	Draft	Preliminary version
0.2	05/July/2019	Draft	Preliminary version
1.0	12/July/2019	Final	Reviewed version
1.1	16/July/2019	Final	Reviewed version
1.2	17/July/2019	Final	Reviewed version
1.3	30/09/2019	Final	Revised version after Interim Review.

Document Keywords

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

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

AI	Artificial Intelligence
AIOTI	The Alliance for the Internet of Things Innovation
BDEC	Big Data and Extreme-scale Computing
BDV	Big Data Value
BDVA	Big Data Value Association
BSC	Barcelona Supercomputing Centre
CEA	Commissariat à l'Energie Atomique (Atomic Energy Commission – France)
CNES	Centre National d'Etudes Spatiales
DCP	Digital Continuum Platform
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)
KAUST	King Abdullah University of Science and Technology – USA
LHC	Large Hadron Collider (CERN)
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)
SKA	Square Kilometer Array
US/USA	United States of America

Executive Summary

The Big Data and Extreme-scale Computing (BDEC) community addresses the subject of HPC and Big Data convergence with a global prospective view. Taking part in the BDEC workshops series represents a unique opportunity for Europe to participate in and contribute to the global definition of the future of HPC, based on a shaping strategy approach, which aims to affirm the European engagement in providing strategic directions for the profitable use of HPC in the context of HPDA and Big Data in general.

Since 2003, BDEC workshops are organized as a world-wide international coordination between the HPC ecosystem and the Big Data community. Europe is part of BDEC since its inception, and the European participation in the BDCEC meeting has been supported by EESI, EESI2 and EXDCI, and now, EXDCI-2 projects. From the 14th to the 16th of May 2019, a workshop was organized in Poznań (Poland) under the supervision of the EXDCI-2 project and in the framework of the EuroHPC Summit Week. The event gathered 51 participants from 15 countries.

Through global and inspirational talks (keynotes), state-of-the-art presentations (demonstrators on scientific cases, roadmaps on policies) and collective thinking (brainstorming and breakout sessions), the meeting achieved the following objectives: giving an overview of the current status of the convergence between HPC and Big Data, studying Machine Learning and Artificial Intelligence as key drivers for the convergence between HPC and Big Data, determining the next challenges for BDEC demonstrators (a *proof-of-concept* platform designed to demonstrate some common capabilities- see Section 2) to tackle.

The major outcome of the workshop was the choice of 2 initial demonstrators that will pave the way for more numerous application-specific exemplars. These will be further elaborated in the upcoming BDEC workshop, to be held in October 2019 in the US.

1. Introduction

In the past years, the United States, the European Union, China 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 phenomenon of Big Data (and more recently, Machine Learning) 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.

Over the past years, the transition to the Exascale has prompted the creation of various actions – the International Exascale Software Project (IESP) and the European Exascale Software Initiative (EESI) to name but two. The realization of the importance of High-Performance Data Analysis (HPDA), for both traditional supercomputing and large-scale cloud environments, caused the overhaul of the IESP community into that of the Big Data and Extreme-scale Computing (BDEC), an academic-driven think tank, with a large international involvement in Europe, USA and Asia.

The BDEC project has organized a series of international workshops – in Europe, Asia and the U.S. 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:

1. **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.
2. **Prospective**. Brainstorming is at the heart of the BDEC process. The goal is to identify future challenges, breakthroughs and potential barriers between the HPC and Big Data communities. The result is to achieve a common view on scientific challenges to tackle.
3. **Networking**. BDEC is a unique place where international researchers and experts can meet and network.

In the first phase of BDEC (2013-2018), the focus was on the convergence between HPC and HPDA via the software stacks. The main outcome summarizing the BDEC effort is the "Pathways to Convergence" paper[1]. Building on the positive feedback of these previous editions, BDEC2 has planned six further international workshops, over a 2-year period (2018-2020), to enable transnational research communities in a wide range of disciplines to converge on a common approach to supporting big data scientific research consistent with exascale initiatives. As part of this, the May 2019 BDEC workshop was held in Poznań, Poland, as part of the EuroHPC Summit Week, which is organized by the EXDCI-2 project. The next meeting, organized by the US in October 2019, will be the fourth in the series.

[1] Asch, M., Moore, T., Badia, R., Beck, M., Beckman, P., Bidot, T., … Zacharov, I. (2018). Big data and extreme-scale computing: Pathways to Convergence-Toward a shaping strategy for a future software and data ecosystem for scientific inquiry. The International Journal of High Performance Computing Applications, 32(4), 435–479. https://doi.org/10.1177/1094342018778123

During three days in Poznań, international experts from 15 countries exchanged on the state of the art of the convergence between HPC and Big Data, and drew paths for the future through intensive breakout and brainstorming sessions. The major focus of this workshop was on "demonstrators" that we define as: *a proof-of-concept platforms designed to demonstrate some common capabilities that some of our BDEC applications and application communities need.*

2. Setting and Agenda

The meeting leader was Mark ASCH from University of Picardie Jules Verne, co-leader of the work package on international collaboration. A specific group composed of EU and US representatives (Mark ASCH, Rosa BADIA, Sergi GIRONA, François BODIN, Jean-François LAVIGNON, Thierry BIDOT, Pete BECKMAN and Terence MOORE) supervised the meeting with an active participation of Asian colleagues.

The BDEC workshop took place in Poznań from the 14th to the 16th of May 2019, at the hotel Mercure Centrum Poznań. This location was chosen for its proximity with the main event that BDEC was a part of, the EuroHPC Summit Week, also organized by the EXDCI-2 project.

After some keynotes giving insight on the state of the art (cf. Objective 1 of the BDEC series, as indicated in Section 1), the majority of the workshop was dedicated to the demonstrators. BDEC's aim is to produce some *proof-of-concept* platforms designed to demonstrate some common capabilities that some of the BDEC applications and application communities need. The objective is to produce a working version of an international, federated, continuum-spanning[2] demonstrator that can be cooperatively operated and managed and that engages stakeholders at all levels.

Mark ASCH sent out a call for demonstrators prior to the workshop. Eleven responses were received, covering all the major regions: Asia, US, Europe. All contributions were evaluated by a scientific committee and were accepted for presentation during the workshop. During the workshop, discussions were held to condense the 11 proposals into a small number of demonstrators (initially, 2) that capture as much as possible of the commonalities among all the others.

The number of registered participants in the meeting was 51, of which 13 were women. The names and affiliations of the participants can be found in the annex. During the workshop, additional PRACE representatives, who came over from the EuroHPC venue, also participated.

[2] By continuum spanning, we mean a complex workflow of computations and data exchanges that is executed on resources at the edge (sensor, IoT, scientific instrument), in the cloud and at some HPC system.

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.

Tuesday, May 14th, 2019	
7:00pm to 9:00pm	Reception and Registration at Mercure Hotel
Wednesday, May 15th, 2019	
INTRODUCTION AND KEYNOTES	
8:30am to 8:40am	**Welcome Address** Pete Beckman, ANL Mark Asch, EXDCI
8:40am to 9:40am	**Opening Plenary Session at EuroHPC** **Plenary Talks:** *The Convergence of Big Data and Large-scale Simulation Leveraging the Continuum* David Keyes, KAUST *Workflows for continuum computing platforms (from edge to HPC computing)* Rosa Badia, BSC

9:40am to 10:30am	**Opening Plenary Session at EuroHPC** **Regional Updates** China *"Big Data and Extreme Computing: Software Ecosystem for Sunway Processors"* Jingheng Xu, National Supercomputing Centre Wuxi Japan *"Japanese HPC Infrastructure Update"* Masaaki Kondo, Riken R-CCS/Univ. of Tokyo (on behalf of Satoshi Matsuoka, Director, Riken R-CCS) USA *"USA Update"* Pete Beckman, ANL EU *"EuroHPC and the European HPC"* Leonardo Flores Añover, DG CONNECT, EC
10:30am to 11:15am	Coffee Break
BDEC WORKSHOP	
11:15am to 11:45am	**Keynote** *Polish e-Infrastructure used for Grand Challenges* N. Meyer, PSNC
11.45am to 12.15pm	**Platform Demonstrators** Pete Beckman, ANL
12:15pm to 1:30pm	Lunch Break
DEMONSTRATORS	

1:30pm to 3:30pm	**Demonstrators Presentations** 5-slide Presentations of Contributions-10 Minutes per Presentation *Towards a demonstrator of the Sigma Data Processing Architecture for BDEC 2* Gabriel Antoniu *Personalized healthcare: workflow orchestration for cyberinfrastructure platforms* Rosa Badia *A Global Data Logistics Network for the Digital Continuum* Micah Beck *Twister2 Demonstrations for BDEC* Geoffrey Fox *Demonstrators for HEP* Maria Girone/ Gavin McCance *Supercomputing in the Era of "Society 5.0" by Integration of Simulation, Data, Learning* Toshihiro Hanawa *DATA TERRA: toward a fully integrated earth sciences data distributed platform* Richard Moreno *Innovation Driven Computing platform* Ariel Oleksiak *Stream processing for the continuum* Martin Swany *ABCI Distributed and containerized AI/Bigdata processing environment* Ryousei Takano *Multi SPMD (mSPMD) programming model for HPC, Big Data and Edge Computing* Miwako Tsuji *Data Processing in Radio Astronomy - Preparing for the Square Kilometre Array* Michiel van Haarlem / Raymond Oonk

3:30pm to 3:45pm	Coffee Break
BREAKOUT SESSIONS	
3:45pm to 5:45pm	**Breakout session 1 (first group)**
3:45pm to 5:45pm	**Breakout session 1 (second group)**
7:00pm – 10.00pm	Gala Dinner (EuroHPC Summit Week)
Thursday, May 16th, 2019	
ROADMAPS AND OVERVIEWS	
9:00am to 9:30am	**Day 1 Breakout Reports**
9:30am to 11:30am	**Breakout session 2 (groups merged)**
11:30am to 12:30pm	Reporting of Breakout Sessions I and II
12:15pm to 1:30pm	Lunch Break
1:30pm to 2:00pm	**Panel discussion**
2:00pm to 2:30pm	Wrap up of Demonstrator sessions and next steps
2:30am to 3:00pm	Organizing for next BDEC workshop (US, autumn 2019)
3.00pm – 3.30pm	Conclusion – Chairs of BDEC Steering Committee
3.30pm	Visit of PNSC Living Labs

The presentations were published online during the workshop on the dedicated website www.exascale.org thanks to the University of Tennessee team.

3. Report on Day 1

On the first day, BDEC speakers opened the EuroHPC Summit Week keynote session with two keynotes and four regional presentations to give a global vision of the community to the broader public. We then moved on to the closed sessions, beginning with a presentation by our Polish hosts (N. Meyer, PSNC) on "*Polish e-Infrastructure used for Grand Challenges".* This opening talk was followed by the presentation of the workshop's goals and organization (P. Beckman, T. Moore, M. Asch).

3.1. Keynotes

BDEC provided a European and international view on different needs of Big Data projects and the increasing demands of HPC. The two keynotes given by David Keyes (KAUST) and Rosa Badia (BSC) introduced the challenges, expectations and potential solutions for the "Digital continuum".

David Keyes pleaded the idea of the importance of the convergence of Big Data and Large-scale simulation, which in the exascale computational era will give humanity predictive tools to overcome grand natural and technological challenges. Moreover, there are several motivations for the convergence: the scientific and engineering advances (predictive, execution performances), economy of data centre operations and development of a competitive workforce. He also highlighted the roles for Artificial Intelligence (AI) and Machine Learning (ML) in the application, computational infrastructure as well as machine learning at the edge.

Figure 2 - David Keyes presenting for the BDEC plenary at the EuroHPC Summit Week 2019 in Poznań

Rosa Badia, on her side, explored workflows for continuum computing platforms, providing such relevant examples as: mF2C[3] – Smart Fog Hub System installed in the Cagliari airport; CLASS – the use-case in Automotive Smart Area in the city of Modena (Italy); ELASTIC[4] – Florence tramway network. The common challenges for these solutions are the following: complex infrastructures, large amount of data from multiple sources, workflow orchestration with complex applications in complex environments. To tackle these challenges the Python programming language associated with a task based runtime (PyCOMPSs/COMPSs) seems to be the most appropriate as it allows the integration with persistent memory, support for elasticity and interactivity and also integrates smoothly with Machine Learning packages.

Figure 3 - Rosa Badia presenting for the BDEC plenary at the EuroHPC Summit Week 2019 in Poznań

The keynotes were followed by Regional updates. The speakers respectively presented the current roadmaps of China, Japan, the US and EU.

Jingheng Xu started with the presentation of the Software Ecosystem for Sunway Processors with the overview of the Exascale Supercomputers in China, development of Performance Tools and Libraries on Sunway Processors and HPC & Big Data Applications on Sunway-TaihuLight.

Masaaki Kondo gave an update on the Japanese HPC Infrastructure, sharing the preliminary performance evaluation results on the largest and fastest supercomputer to be built by 2020 (Post-K now named Fugaku).

Pete Beckman unveiled the Exascale Computing project (ECP) activities in the USA. The project is part of the broader DOE Exascale Computing Initiative (ECI) and its objective is to

[3] mF2C (Fog-to-Cloud) is an European project supported by Horizon 2020 under the grant agreement n° 730929
[4] ELASTIC (A Software Architecture for Extreme-ScaLe Big-Data AnalyticS in Fog CompuTIng ECosystems) is a European project supported by Horion 2020 under the grant agreement Nº 825473

maintain the high standards of AI in the USA, promoting sustained investment in AI R&D, reducing barriers to the use of AI technologies to promote their innovative applications, ensuring that technical standards minimize vulnerability to attacks, training the next generation of American AI researchers etc.

Leonardo Flores from the European Commission shared the mission and objectives of The EuroHPC Joint Undertaking Action for the next eight years, which aims to establish an integrated world-class supercomputing and data infrastructure and support a highly competitive and innovative HPC and Big Data ecosystem in the EU. The ambition on the EU side is high, with a focus to reduce the market asymmetry.

Figure 4 - Leonardo Flores presenting for the BDEC plenary at the EuroHPC Summit Week 2019 in Poznań

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…).

After the keynotes session ended, Norbert Meyer from the Poznań Supercomputer and Networking Centre (PSNC), the co-organizer and the local host of the EuroHPC Summit Week, presented the PSNC and the Polish e-Infrastructure to the BDEC workshop participants.

3.2. Demonstrators' presentation

As stated above (Section 2), BDEC's aim is to produce some *proof-of-concept* platforms designed to demonstrate some common capabilities that some of the BDEC applications and application communities need. The objective is to produce a working version of an international, federated, continuum-spanning demonstrator that can be cooperatively operated and managed and that engages stakeholders at all levels.

As existing cyberinfrastructure was not designed to adequately deal with edge-to-cloud/HPC workflows, especially not extremely data intensive ones, the BDEC community proposed to

launch a series of international *workshops*, bringing together computer scientists, application scientists, Big Data, IoT, AI and other stakeholders who are focused on achieving this goal.

Twelve demonstrator proposals were presented in Poznań (including the Polish host's project), covering a wide range of topics and disciplines at the forefront of HPC and Big Data (see Workshop Agenda above).

Each demonstrator was requested to provide a workable solution for several challenges:

- an end-to-end problem, spanning the continuum, from Edge to HPC in the Cloud;
- a software stack problem;
- a resource allocation problem to address an on-demand, shared infrastructure;
- a data movement and logistics problem (in both directions- edge/instrument to HPC, HPC to instrument/edge;
- a robustness of large, interlinked, composed infrastructures;
- a support of monitoring and control (based on AI).

Five main criteria for a good demonstrator were mentioned:

1. Could evolve to support multiple application domains.
2. Reveals a programming model from edge to cloud.
3. Shows global workflow- data, resources, users, etc.
4. Should be based on a reusable architecture across multiple scales.
5. Could evolve to run across several different composed infrastructures.

Domains of application of the demonstrators were quite various, from healthcare to data processing in Radio Astronomy, logistics network and the multi SPMD programming model for HPC, Big Data and Edge computing, satellite data capture and analysis.

Figure 5 - Miguel Vázquez presenting the "Personalized healthcare: workflow orchestration for cyberinfrastructure platforms" demonstrator at BDEC workshop 2019 in Poznań

3.3. First breakout sessions

Breakout sessions are the traditional BDEC tool for exchanging, elaborating and writing up conclusions. Following the white-paper presentations, we always regroup into 3-4 breakouts (depending on the theme and the goals of the workshop), where the issues are discussed, noted down, and then reported in post-breakout, plenary sessions.

The breakout sessions in Poznań were gathered around seven main questions related to the demonstrators:

- Describe the data at the instrument/edge.
- Describe how data is moved and buffered/stored.
- Describe how the global workflow is specified.
- Describe how resources are allocated.
- Describe how the continuum is programmed.
- Describe how/where the (big) data analytics are done.
- Describe the simulation/model.

Two broad themes for the demonstrators were agreed upon by the participants to be discussed in the breakout sessions:

- "Learning Everywhere" - use ML (AI) in conjunction with Simulations and Big Data to learn Science Discoveries everywhere – "MLaroundHPC.
- A global data logistics network. An international, cooperative network of nodes deployment. A network that will store, process and move images from satellites and microscopy, and from major instruments in physics and astronomy, as well as their secondary products as necessary, to make them seamlessly available worldwide regardless of source or current location or format, referenced by community specified metadata.

Two spokespersons were allocated to each group with two functions:

- Organize the discussion and the capitalization.
- Organize and present the results at the end of the process.

4. Report on Day 2

4.1. Breakout Reports

The session spokespersons synthesized the main ideas and summarized in a short oral presentation, François Bodin, Mark Asch and Terry Moore moderated the reports session.

After merging the results of each group, several innovations were identified in data storing, resources allocation and continuum programming.

Innovations identified:

- The need to generalize data locality and data management optimization.
- The lack of adequate, unified (but supported by a variety of infrastructures) authentication / authorization/credentials management through the system, and that persists (in time, duration).
- The need to improve portability across time and space, i.e. across the continuum.

It was decided to discuss some topics at the next workshop session:

- The necessity of compelling demonstration applications, to apply for funding.

- New Machine Learning algorithms need to be addressed from an HPC viewpoint.

4.2. Second breakout session and wrap-up

For the second breakout session the moderators proposed to unify two working groups in one room in order to receive an immediate feedback from each other during the discussion. The moderators chose a different methodology to frame this session – DARPA's Heilmeier catechism[5] - a set of questions that anyone proposing a research project or product development effort should be able to answer.

Heilmeier catechism:

- What are you trying to do? Articulate your objectives using absolutely no jargon.
- How is it done today, and what are the limits of current practice?
- What is new in your approach and why do you think it will be successful?
- Who cares? If you are successful, what difference will it make?
- What are the risks?
- How much will it cost?
- How long will it take?
- What are the mid-term and final "exams" to check for success?

These questions are intended to be addressed in the inter-workshop working groups. These groups meet remotely, on a regular basis, and will present reports at the next meeting (US, October 2019).

5. Conclusion of the meeting

The BDEC workshop in Poznań emphasized the challenges and expectations, and potential solutions for the "Digital continuum" on which the next generation cyberinfrastructure platforms will rely, dealing with heterogeneity of data and computing resources.

Our main objective is to enable transnational research communities in a wide range of disciplines to converge on a common "digital continuum platform" (DCP), a next generation network computing platform for creating distributed services in a world permeated by devices and saturated by digital data. The BDEC community intends to tackle this problem by pursuing three complementary objectives:

- Draft a design for a distributed services platform for science to serve as shared software infrastructure for the growing continuum of computing devices and data sources on which future science will rely
- Organize and develop an international demonstration of the feasibility and potential of the DCP (prototype implementation)
- Develop a corresponding "shaping strategy" addressing all relevant stakeholders and moving the community toward convergence on a standard DCP specification.

In addition to the work of the international meetings themselves, inter-meeting working groups are intended to accelerate community-wide discussion and collaborative activities needed to address the multi-dimensional challenges of the emerging digital continuum. By achieving its goals, this project aims to supply the different stakeholder communities with the kind of well-defined vision and consensus building strategy necessary to realize a common, open, and interoperable DCP for digital continuum era.

[5] https://www.darpa.mil/work-with-us/heilmeier-catechism

The international BDEC consortium agreed during its meeting in Poznań on two initial demonstrators of such platforms, involving a panel of experts coming from both the major HPC regions in the world, but also the entire value chain: from machines to users, including participants from the large international instruments LHC, SKA and CNES. The choice of these demonstrators is by no means fixed or final. They are to play the role of *initiating a focused effort* that will enable BDEC to solicit funding from national research funding agencies, around the globe. These initial demonstrators aim to address generic issues, and develop *common tools* that can then be adapted and adopted by the numerous application-specific demonstrators.

The meeting was broadly acknowledged as a big success. We have shown that Europe can lead the way in bringing together major application communities (high energy physics, radio astronomy, satellites, climate, etc.) together with computer scientists and architects. This unprecedented convergence guarantees the place of Europe in the future international cyber infrastructure endeavors. These will enable us to be leaders in addressing major societal challenges, such as climate change, clean energy, personalized health, autonomous vehicles, etc.

The efforts of the BDEC think tank are well integrated in ongoing efforts at European scale: BDEC continues to be an active contributor to the preparation of the 4th Strategic Research Agenda (SRA-4), and via EXDCI-2, we are in contact with the following European Technology Platforms: ETP4HPC, BDVA, and AIOTI. This effectively “closes the loop” by ensuring that the SRA is aligned with the international state-of-the-art in converged and continuum e-infrastructures.

Annex – Workshop participants

First name	Last name	Affiliation
Gabriel	Antoniu	INRIA
Mark	Asch	University of Picardie, Amiens
Rosa	Badia	Barcelona Supercomputing Centre
Micah	Beck	University of Tennessee
Peter	Beckman	Argonne National Laboratory
Thierry	Bidot	Neovia Innovation
François	Bodin	University of Rennes 1
David	Carrera	Barcelona Supercomputing Centre
Carlos	Costa	IBM Research
Sam	Crawford	University of Tennessee
Ewa	Deelman	University of Southern California Information Sciences Institute
Mary	Dzielski	Argonne National Laboratory
Athanasia	Evangelinou	European Commission
Geoffrey	Fox	Indiana University
Todd	Gamblin	Lawrence Livermore National Laboratory
Sergi	Girona	Barcelona Supercomputing Centre
Maria	Girone	European Organization for Nuclear Research
Toshihiro	Hanawa	University of Tokyo
Ren	Hu	National Supercomputing Centre in Wuxi
Kate	Keahey	Argonne National Laboratory
David	Keyes	King Abdullah University of Science and Technology (KAUST) / Columbia University
Christian	Kniep	QNIB Solutions
Masaaki	Kondo	University of Tokyo / RIKEN Centre for Computational Science
William	Kramer	National Centre for Supercomputing Applications / University of Illinois
Jean-François	Lavignon	TS-JFL

Violaine	Louvet	UMS GRICAD
Michael	Malms	IBM Research
Satoshi	Matsuoka	RIKEN Centre for Computational Science
Gavin	McCance	European Organization for Nuclear Research
Piyush	Mehrotra	NASA Ames Research Centre
Takemasa	Miyoshi	RIKEN Advanced Institute for Computational Science
Terry	Moore	University of Tennessee
Richard	Moreno	National Centre for Space Studies
Ariel	Oleksiak	Poznań Supercomputing and Networking Centre
Raymond	Oonk	SURFsara
Marcin	Ostasz	Barcelona Supercomputing Centre
Michel	Perault	French National Centre for Scientific Research
Maria	Pérez	Universidad Politécnica de Madrid
Tracy	Rafferty	University of Tennessee
Kentaro	Sano	RIKEN Centre for Computational Science
Joan	Snoderly	University of Tennessee
Martin	Swany	Indiana University
Ryousei	Takano	National Institute of Advanced Industrial Science and Technology
William	Tang	Princeton University / Princeton Plasma Physics Lab
Osamu	Tatebe	University of Tsukuba
Miwako	Tsuji	RIKEN Centre for Computational Science
Miguel	Vazquez	Barcelona Supercomputing Centre
Jean-Pierre	Vilotte	Paris Institute of Earth Physics
Robert	Wisniewski	Intel
Jinghen	Xu	Tsinghua University
Ramin	Yahyapour	GWDG