

## **H2020-FETHPC-2014**

### **Coordination of the HPC strategy**



### **EXDCI**

### **European eXtreme Data and Computing Initiative**

**Grant Agreement Number: FETHPC-671558**

### **D6.3**

### **Report on the BDEC workshop (US)**

*Final*

Version: 1.21  
Author(s): Mark ASCH, ANR and Université de Picardie Jules Verne  
Date: 17.04.2018

## Project and Deliverable Information Sheet

<b>EXDCI Project</b>	<b>Project Ref. №: FETHPC-671558</b>	
	<b>Project Title: European eXtreme Data and Computing Initiative</b>	
	<b>Project Web Site: <a href="http://www.exdci.eu">http://www.exdci.eu</a></b>	
	<b>Deliverable ID: D6.3</b>	
	<b>Deliverable Nature: Report</b>	
	<b>Dissemination Level:</b> PU	<b>Contractual Date of Delivery:</b> 28 / February / 2018
		<b>Actual Date of Delivery:</b> 18 / April / 2016
<b>EC Project Officer: Evangelia MARKIDOU</b>		

\* - 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

<b>Document</b>	<b>Title: Report on the BDEC Workshops (EU)</b>	
	<b>ID: D6.3</b>	
	<b>Version: 1.21</b>	<b>Status: Final version</b>
	<b>Available at: <a href="http://www.exdci.eu">http://www.exdci.eu</a></b>	
	<b>Software Tool: Microsoft Word 2016</b>	
	<b>Files: EXDCI_D6.3_Final.docx</b>	
<b>Authorship</b>	<b>Written by:</b>	Mark ASCH, ANR and Université de Picardie Jules Verne
	<b>Contributors:</b>	
	<b>Reviewed by:</b>	John CLIFFORD, PRACE AISBL Michael MALMS, IBM Zurich
	<b>Approved by:</b>	TB / MB

## Document Status Sheet

<b>Version</b>	<b>Date</b>	<b>Status</b>	<b>Comments</b>
1.1	05/04/2018	Draft	For internal review
1.2	13/04/2018	Final version	Final version
1.21	17/04/2018	Uploaded version to PP	Some minor changed by PMO

## Document Keywords

<b>Keywords:</b>	Exascale, Future Cyber Infrastructure, BDEC, Big Data, Convergence, International Cooperation.
------------------	--

### Copyright notices

© 2018 EXDCI Consortium Partners. All rights reserved. This document is a project document of the EXDCI project. All contents are reserved by default and may not be disclosed to third parties without the written consent of the EXDCI partners, except as mandated by the European Commission contract FETHPC-671558 for reviewing and dissemination purposes.

All trademarks and other rights on third party products mentioned in this document are acknowledged as own by the respective holders.

# Table of Contents

- Executive Summary ..... 6**
- 1 Introduction ..... 7**
- 2 Setting and Agenda..... 7**
- 3 Report on Day 1 .....10**
- 4 Report on Day 2.....10**
- 5 Conclusion of the meeting.....11**
- Annex – Workshop participants .....12**

## List of Figures

Figure 1 - BDECng group photo ..... 8

## References and Applicable Documents

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

## List of Acronyms and Abbreviations

ANL	Argonne National Laboratory (USA)
ANR	Agence Nationale de la Recherche (National Research Agency – France)
BDEC	Big Data and Extreme-scale Computing
BSC	Barcelona Supercomputing Center (Spain)
DOE	Department of Energy (USA)
EC	European Commission
EU	European Union
EXDCI	European eXtreme Data and Computing Initiative
HPC	High Performance Computing
NSF	National Science Foundation - USA
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 March 2018, the closing workshop was organised in Chicago (US) with the participation of the EXDCI project. Nineteen participants from US, Asia and Europe attended this planning meeting.

The objective of the meeting was to agree upon, and set up, a planning process for the next stage of BDEC, that we currently name **BDECng** (ng = next generation). The outcome of the meeting will be summarized in a 2-page prospectus, a calendar for the period September 2018 to September 2020, and a process.

### **Important Note:**

Since this event took place after the official end of EXDCI, a detailed report will be prepared and submitted during the EXDCI-2 project.

## 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 2018, the final BDEC workshop was held in Chicago, USA, with the collaboration of the EXDCI project.

During two days, a core-group of international experts from Europe, USA and Asia exchanged on the state of the art of the convergence between HPC and Big Data, and drew paths for the future execution of the BDEC initiative, that we have provisionally named BDECng (ng = next generation).

## 2 Setting and Agenda

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

The BDEC planning workshop took place in Chicago from the 26<sup>th</sup> to the 28<sup>th</sup> of March 2018, and was held in the Law School of Northwestern University.

The number of (invited only) participants in the meeting was 19, of which 5 were women. The list was composed by the US partners. The names and affiliations of the participants can be found in the annex.



Figure 1 - BDECng group photo

### Program of the workshop

The program of the workshop is given below.

<b>Monday, March 26th, 2018</b>	
6:00pm to 8:00pm	Reception and Registration
<b>Tuesday, March 27th, 2018</b>	
INTRODUCTION AND STATE OF PLAY	
9:00am to 9:20am	<b>Welcome, Overview and Goals of the Meeting</b> Jack Dongarra (ICL/UTK) and Pete Beckman (ANL)
9:20am to 09:40am	<b>The Big Picture</b> Dan Reed, U. Iowa, USA.
09:40am to 10:10am	<b>US stakeholders update</b> Pete Beckman (DOE), Manish Parashar (NSF).
10:10am to 10:30am	Coffee Break
INTERNATIONAL UPDATES	
10:30am to 11:30am	<b>China</b> Yutong Lu (NSCC)



	<p><b>Japan</b> Yutaka Ishikawa (RIKEN)</p> <p><b>EU</b> Rosa Badia (BSC)</p>
11:30am to 12:30pm	<p><b>Possible Architecture Visions</b> Pete Beckman (ANL) François Bodin (IRISA)</p>
12:30pm to 1:30pm	Lunch Break
1:30pm to 3:00pm	<p><b>Discussion Part 1</b></p> <p><i>Terry Moore:</i> Shaping strategy: How to map shaping into a community process?</p> <p><i>Michela Tauffer:</i> Assumptions: What assumptions are we making about the environment that we have to address?</p> <p><i>Alex Szalay:</i> Broader Community: Who is and how do we engage the broader community: scientific and cyberinfrastructure</p> <p><i>François Bodin:</i> Convergence architecture</p>
3:00pm to 3:30pm	Coffee Break
3:30pm to 5:30pm	<p><b>Discussion Part 2.</b></p> <p><i>Geoffrey Fox:</i> Commercial Providers / Cloud interoperability</p> <p><i>Sergi Girona:</i> Business: Business model/Government involvement</p> <p><i>Micah Beck:</i> Artifacts: What should outcomes of the BDECng community process be?</p> <p><i>Yutong Lu:</i> Transnational Collaboration: What will the major challenges be in trying to work across national boundaries.</p> <p><i>Mark Asch:</i> Meeting Structures, Methods (When/Which communities)</p>
7:30pm	Dinner
<b>Wednesday, March 28th, 2018</b>	
<b>DISCUSSIONS AND DRAFTING</b>	
9:00am to 9:10am	<p><b>Agenda Overview</b> Terry Moore.</p>
9:10am to 10:15am	<p><b>Discussions</b></p> <ul style="list-style-type: none"> <li>• New community (20min)</li> </ul>

	<ul style="list-style-type: none"> <li>• Process (20min)</li> <li>• Schedule (20min)</li> <li>• Pick 4 breakout leads</li> </ul>
10:15pm to 10:30am	Coffee Break
DRAFTING	
10:30am to 11:30am	<b>Drafting of elements for Prospectus.</b> <ul style="list-style-type: none"> <li>• Tech challenge</li> <li>• New community: Who and when</li> <li>• Potential Artifacts</li> <li>• Process</li> </ul>
SUMMARIZE AND CLOSING	
11:30am to 12:00pm	Summarize and plan future videoconferences.

### 3 Report on Day 1

On the first day, the participants took stock of the actual situation and discussed the elements for going forward. The assumption is that BDECng will use a “shaping strategy” approach that is well suited to our context: how to engage a community effort for structuring a ground-breaking domain, when a roadmap is no longer indicated.

### 4 Report on Day 2

On day two, we discussed the elements of planning the next 2 years. A basic calendar was agreed upon. There will be 3 meetings per year, alternating between large “open” meetings and small “closed” meetings.

	EU	Asia	US
2018			Oct: (Small) Geoffrey Fox
2019	May (small): Mark Asch	Feb. (large) Yutaka Ishikawa	Sept/Oct (large)
2020	May (large): Mark Asch	Feb/March (small)	

## 5 Conclusion of the meeting

We have reached an international consensus on how to pursue the BDEC initiative over the next 2.5 years. A strong rhythm of meetings (see table above) is needed to keep pace with the rapidly changing context. The meetings will follow a “shaping strategy approach” that is made up of 3 stages: a shaping platform, a shaping view, and finally, shaping acts and assets. For each meeting, appropriate “friends”, from outside our community, will be invited to contribute.

In the months leading up to the first BDECng meeting (US, October 2018), we will prepare a short document describing the process that will be distributed to stakeholders and funders.

This meeting’s objective was to create a plan for an international community planning process to shape community convergence on a common, next generation cyberinfrastructure paradigm in which stakeholders can invest.

Our long-term objective is to create a plan for a common, pervasively distributed computing platform capable of supporting a global, yotta-scale data ecosystem for science. This will consist of the following deliverables:

- Progress reports.
- Prototypes and demos for specific contexts.
- Outreach and diffusion of the results and conclusions.

## Annex – Workshop participants

First Name	Last Name ▲	Affiliation
Mark	Asch	University of Picardy
Rosa	Badia	Barcelona Supercomputing Center
Micah	Beck	University of Tennessee
Pete	Beckman	Argonne National Laboratory
Francois	Bodin	Institute for Research in Computer Science and Random Systems / University of Rennes 1
Alok	Choudhary	DOE
Jack	Dongarra	University of Tennessee
Nicola	Ferrier	Argonne National Laboratory
Geoffrey	Fox	Indiana University
Sergi	Girona	Partnership for Advanced Computing in Europe / Barcelona Supercomputing Center
Yutaka	Ishikawa	RIKEN Advanced Institute for Computational Science
Kate	Keahey	Argonne National Laboratory
Yutong	Lu	National Supercomputer Centre in Guangzhou
Terry	Moore	University of Tennessee
Jean-Pierre	Panziera	ATOS/Bull/ETP4HPC
Manish	Parashar	National Science Foundation
Dan	Reed	University of Iowa
Alex	Szalay	Johns Hopkins University
Michela	Taufer	University of Delaware