H2020-FETHPC-2014

Coordination of the HPC strategy

EXDCI

European eXtreme Data and Computing Initiative

Grant Agreement Number: FETHPC-671558

D4.7
Second holistic vision and recommendations report

Final

Version: 1.0
Author(s): François Bodin, University of Rennes 1
Date: 10.08.2017

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: D4.7	
	Deliverable Nature: Report	
	Dissemination Level: PU*	**Contractual Date of Delivery:** August 31, 2017
		Actual Date of Delivery: August 10, 2017
	EC Project Officer: Evangelia Markidou	

* - 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: Second holistic vision and recommendations report**	
	ID: D4.7	
	Version: 1.0	**Status:** final
	Available at: http://www.exdci.eu	
	Software Tool: Microsoft Word 2013	
	File(s): D4.5-final.docx	
Authorship	**Written by:**	François Bodin, University of Rennes 1
	Contributors:	
	Reviewed by:	Michael Malms Stéphane Réquena Iñigo Yenes (PMO)
	Approved by:	MB/TB

Document Status Sheet

Version	Date	Status	Comments
alpha	16/06/2017	Draft	
Beta	09/07/2017	Draft	
Beta2	12/07/2017	Draft	
Final	10/08/2017	Final	

Document Keywords

Keywords:	High Performance Computing (HPC), Extreme computing, Exascale, ETP4HPC, PRACE, European Commission, Research Infrastructure, BDEC, Recommendations, disruptive, Operation Policies, IPCEI, Extreme-Scale Demonstrator, SME, Big Data, International, Training, R&D Capitalisation, Standards, Advanced Technologies

Table of Contents

List of Figures

List of Acronyms and Abbreviations

Below is an extensive list of Acronyms used in previous deliverables. Please add additional ones specific to this deliverable and delete unrelated ones.

AISBL	Association Internationale Sans But Lucratif (International Non-for-Profit Association)
BDEC	Big Data and Extreme-scale Computing
BDV	Big Data Value
CoE	Centres of Excellence for Computing Applications
cPPP	contractual Public-Private Partnership
CSA	Coordination and Support Action
D	Deliverable
DG	Directorate General
DoW	Description of Work
EC	European Commission
ECMWF	European Centre for Medium-range Weather Forecasts
EESI	European Exascale Software Initiative
ENES	European Network for Earth System modelling
EPOS	European Plate Observing System
EsD	Extreme scale Demonstrators
EU	European Union
FET	Future and Emerging Technologies
FP7	Framework Programme 7
GDP	Growth Domestic Product
H2020	Horizon 2020 – The EC Research and Innovation Programme in Europe
HPC	High Performance Computing
IDC	International Data Corporation
IESP	International Exascale Software Project
INVG	Istituto Nazionale di Geofisica e Vulcanologia (National Institute of Geophysics and Volcanology)
IPCEI	Important Project of Common European Interest
ISV	Independent Software Vendor
IT	Information Technology
KPI	Key-Performance Indicator
M	Month
OS	Operating System
PM	Person Month
Q	Quarter
R&D	Research and Development

R&I	Research and Innovation
RFP	Request for Proposal
ROI	Return On Investment
SHAPE	SME HPC Adoption Programme in Europe
SHS	Social and Historical Sciences
SME	Small and Medium Enterprise
SRA	Strategic Research Agenda
SWOT	Strengths, Weaknesses, Opportunities and Trends
TRL	Technology Readiness Level
US	United States
WG	Working Group
WP	Work Package

1 Introduction

EXDCI's goal is to "coordinate the development and implementation of a common strategy for the European HPC Ecosystem in order to achieve its global competitiveness within the Horizon 2020 Programme". This requires developing a common European HPC Strategy and helping the European HPC Community to conduct the transition from today's Petascale systems to the Exascale era. In this context, EXDCI has been contributing to this understanding and to widely spreading a vision of the present challenges at a European level.

Petascale to Exascale transition is a very complex transition, especially as it is not happening in isolation. At the same time, data deluge is happening. Exascale definition, limited to producing a machine capable of a rate of 10^{18} flops, is of interest to very few scientific domains. As illustrated in Figure 1, the main issue is dealing with data generated by sensors as well as numerical simulations themselves. Whilst in EESI it was clear that the data issue would be of crucial importance, during the EXDCI time frame the focus has been shifted to the convergence of extreme data and computing with new considerations such as edge computing, in-transit computing, etc. This new focus brings many new capabilities for science (e.g. machine learning) and connections to the Big Data Market, however, at the same time, also with an enlargement of the HPC ecosystem. It is admitted in the EXDCI community that this new challenge is shaping industry and science today and will do so even more in the future.

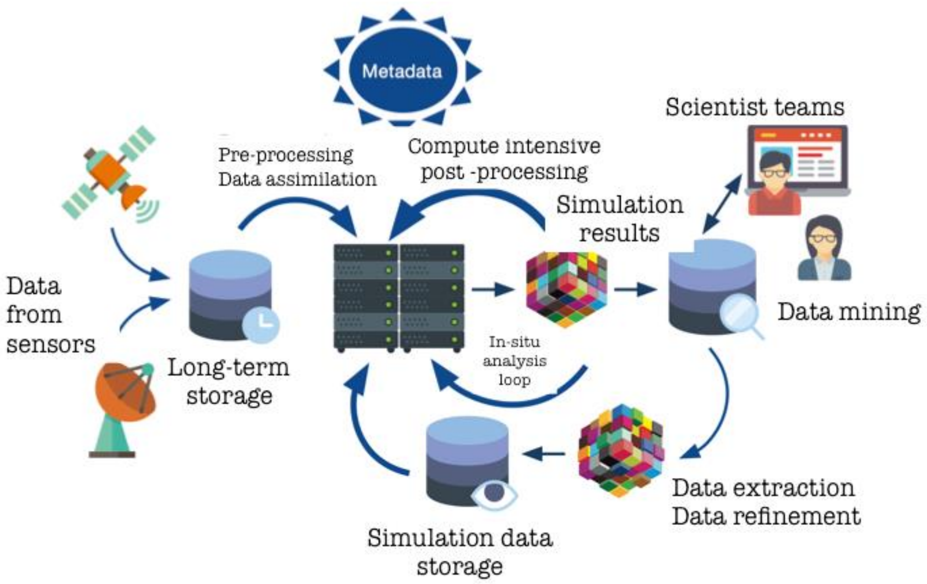

Figure 1: Typical scientific workflow structure (From *« Les big data à découvert »*, CNRS Éditions, 2017)

Despite all these considerations, for part of the scientific community, the Exascale project is perceived as a long–term and abstract issue, whilst the need for computing power is a short-term important question. A shortsighted point of view could consist into buying the next generation of supercomputers from the market. However, the Peta-Exa transition is not an incremental change: it requires the community to adapt to a new way of using computing and storage, and the later that this adaptation starts; the more expensive it is, with negative impact on competitiveness and hiring of skills, as well as the kind of science that can be done. Furthermore, industry is dependent on the academic community to produce trained PhD and engineers for cutting-edge technologies. This can only happen with a community-focused effort and access to a world class research HPC infrastructure based on advanced technologies.

This document does not aim at giving a complete overview of all the activities of EXDCI. Many topics such as KPIs and SMEs are presented in the other deliverables.

This report is organised as follows: section 2 illustrates why moving to the Exascale era is not an incremental step. Section 3 summarises the convergence between data and compute that is at the core of the paradigm change. Section 4 draws the current view of the international competition that the EU is facing in this endeavor. Section 5 gives an overview of the Extreme-scale Demonstrator (EsD) proposal, which is the first concrete step towards the next generation of systems. Section 6 presents the ecosystem and how EXDCI has been creating a network of stakeholders in order to achieve co-design capabilities. Finally, section 7 concludes this report and proposes an additional recommendation to the one already presented in Deliverable D4.3 [9].

2 Exascale is a Paradigm[1] Change

The Petascale to Exascale transition is happening in a context where data needs to be handled differently. The volume of generated data (from scientific instruments, sensors, simulations, etc.) is so large and its localisation is so wide-spread that it affects all system design parameters up to the discovery scientific process. In this respect, the Peta-Exa transition is a much greater disruption than the Tera-Peta transition. Indeed, the Tera-Peta has largely been helped by the parallelisation of applications for the cluster based supercomputers performed during the Giga-Tera transition. As an illustration of such a difference, Figure 2 shows the two periods by comparing workflows of a scientific computation as described in Figure 1 (extract from [4]).

	Tera-to-Peta	Peta-to-Exa
Complexity	Code coupling	More multiphysics, multiphase models, data assimilations, data analytics, edge computing, ...
Heterogeneity	Mostly homogeneous	**Mix of data analytics and simulation, heterogeneous bricks**
Localization	All in one system	Data may come from large scientific instruments, or a large number of small instruments
I/O constrained	Solvable issue	Cannot move the data around, not sure it can be solved
Allocation	Batch mostly	Batch, **interactive (guided simulation and analysis)**, (soft) real-time* (visualization, ...)

Figure 2: Exascale Transition Impact on Workflows. *Big data assimilation for Extreme-scale NWP, Takemasa Miyoshi

[1] Webster definition: *a world view underlying the theories and methodology of a particular scientific subject*.

Tera-Peta transition was smooth and with minimum (side-) effects for most HPC users. Peta-Exa requires profoundly reviewing the way in which scientists use computing and storage resources. As a consequence, there is a long list of questions to consider, any one of these questions having a potentially great impact on scientists and industry practice. Here are some questions raised by the Peta-Exa transition:

- Can one platform fit all?
- How to benefit from the "*Cloud*"?
- How to integrate with edge-computing, in-transit computing?
- How to adapt the discovery process to make advantage of the new data analytics and machine learning?
- How to manage urgent Extreme computing (e.g. immediate access to supercomputers for emergency computations)?

The complexity of designing new systems is furthermore increased by the weight of legacy codes. They represent a huge technical capital and cannot be discarded or adapted without a huge cost (rewriting or *re-architecturing* an application is usually almost unfeasible). So we need to design "revolutionary" systems addressing a deluge of data, whilst at the same time, providing an evolutionary path of least resistance to legacy codes.

A last consideration is the emergence of disruptive technologies such as new non-volatile memory (e.g. Intel 3D Xpoint [1]), quantum accelerators (e.g. IBM qubits,[2]), etc. Many of these technologies can be mainstream at the time of delivery of the Exascale sustainable systems. Currently, it is extremely complex and difficult to anticipate the integration of these technologies in any design.

The next section considers the question of Compute and Data Convergence as a path to the new generation of HPC systems.

3 The Compute and Data Convergence

Exascale systems are not systems of a kind. The scientific community (from Academia as well as Industry) cannot afford the development cost of many specific technologies dedicated to specific domains. The solution to be developed must be sustainable and address a larger market than that of HPC. Furthermore, the Big Data technology may provide many solutions to the scientific community. This is well summarised in the USA National Strategic Computing Initiative (NSCI) [3] "*The NSCI vision is to establish public and private sector collaborations for developing and broadly deploying next generation HPC paradigms that will drive economic growth, enable scientific discovery, and foster innovation*." In EXDCI this question has been addressed in two ways.

The first initiative is a worldwide action known as the Big Data and Extreme Scale initiative [5]. It is a worldwide international initiative that allows scientists from Asia, EU, and USA to exchange on road mapping efforts in order to understand the paradigm shift underlying extreme data and computing. Previous efforts, which started with the IESP (International Exascale Software Project) initiative, have been very successful in helping the Exascale roadmap definition.

The second approach has been to establish links [10] with the Big Data Value Association [7]. This association has a strong industrial basis and includes data users, data providers, data technology providers and researchers.

Both ways are providing a very complementary view of this convergence. The BDEC work has a strong emphasis on the technical aspects such as the “convergence” of the software stacks. An illustration is provided in Figure 3 where ultimately we need to combine both Big Data/AI and HPC capacities on the same platform. Besides the technical aspects, BDVA is providing a strong connection to the Big Data market and technology providers.

The convergence analysis has led to very interesting tracks where HPC is not restricted to numerical simulation but is integrated in a larger context, which includes smart-cities and Internet of Things, where computing at the edge is necessary to avoid expensive and slow data movements and the Cloud becomes part of the HPC infrastructure. For more information the reader can refer to the Report “Pathways for convergence” [6].

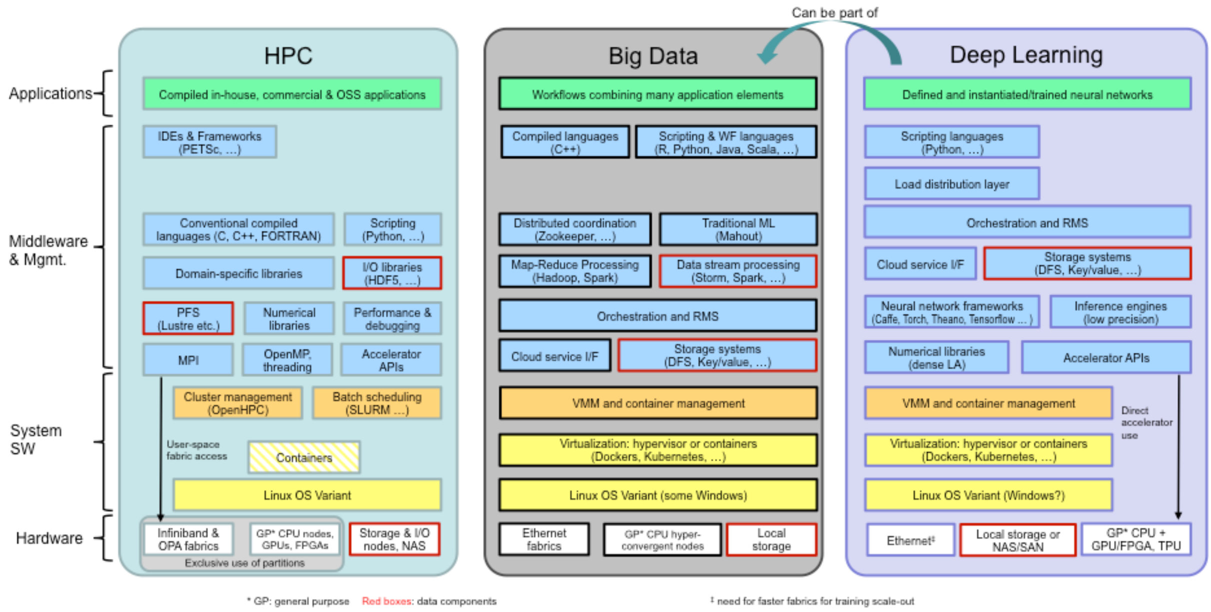

Figure 3: HPC versus Big Data/Deep Learning software stack (Extract from Hans-Christian Hoppe’s presentation at the EXDCI Technical Meeting 2017).

4 International Competition

International competition is becoming stronger as we get closer to the delivery of Exascale systems. Figure 4 shows the current estimate. The remainder of this section gives a brief overview of the status as it was perceived at this year’s edition of the BDEC conference on March 9th and 10th 2017 organised by the Chinese team of the National Supercomputing Center in Wuxi.

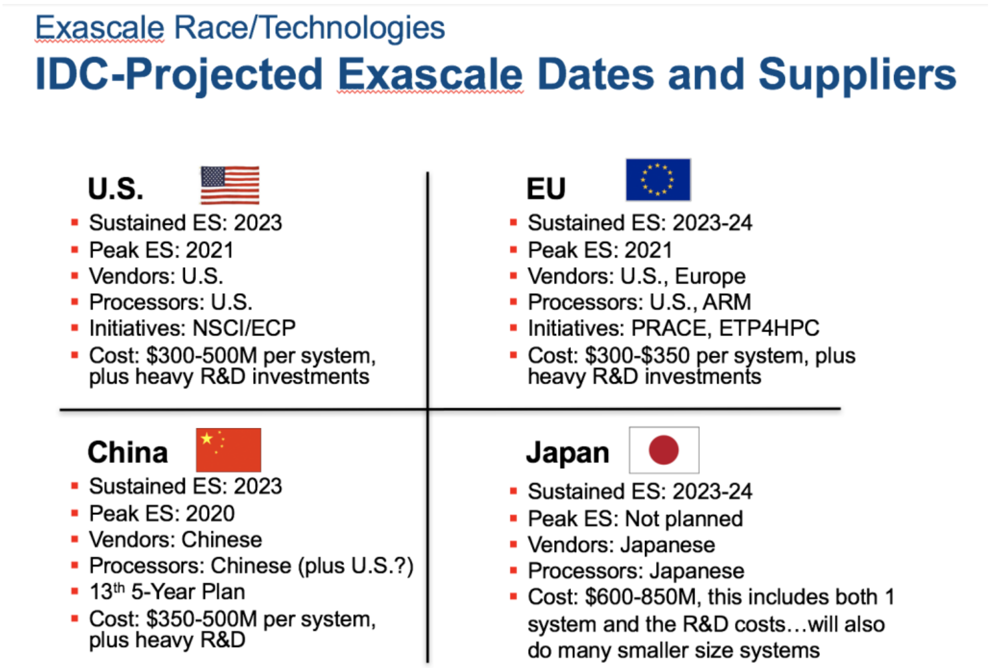

Figure 4: IDC analysis (Hyperion report, https://www.hpcwire.com/2017/04/20/hyperion-idc-paints-bullish-picture-hpc-future/)

This year's edition of BDEC is a milestone in the implementation of convergence. Whilst previous editions of BDEC had concentrated on roadmapping, this year's edition has given a crisp picture of American, Japanese and Chinese implementations and work plans (all talks are available at http://www.exascale.org/bdec/agenda/china).

On the US side, Rob Ross from Argonne National Laboratory presented the Exascale computing project. There are two remarks to be made. Firstly, as shown in Figure 5, all critical software technologies for next generation systems are more or less controlled / developed by American teams. Secondly, the American strategy includes coordination with vendor efforts to achieve sustainability.

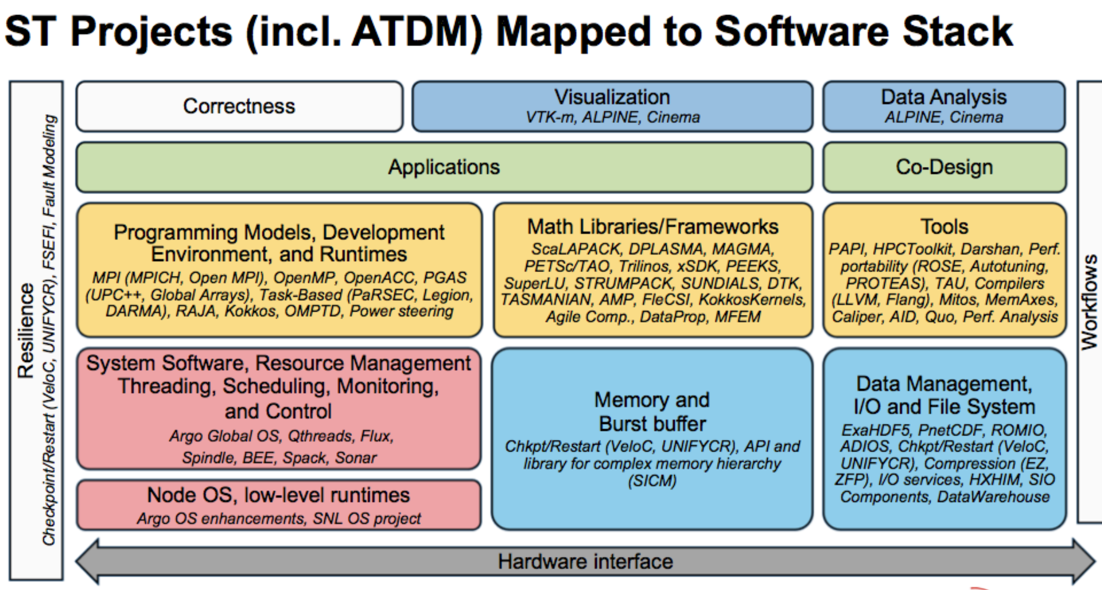

Figure 5: American ECP Software Technology Overview

On the Japanese side, the Exascale Post-K architecture, based on ARM, is already underway (since Q1 2016). It is interesting to note that CEA in France has initiated collaboration with RIKEN (Japan). Besides this flagship project, Japan continues to invest in large systems that will result in a fast convergence of HPC and AI (new AI Research Center, systems Oakforest, Tsubame 3 and ABCI, a new Deep Learning Processor from Fujitsu).

China has been making steady progress on all levels of BDEC related technologies. Not only have they been able to build a 100 Petaflop machine based on Chinese technology, but recent progress on the software stack and applications (a Chinese team won the 2016 ACM Gordon Bell Prize) has been stunning. After an effort, mainly driven by academic players, China starts to build an industry like Sugon, which has been selected as one of the three teams working on a Exascale architecture project.

5 The Extreme Scale Demonstrator concept

From the ETP4HPC SRA: "*The Extreme scale Demonstrators (EsDs) are vehicles to optimise and synergise the effectiveness of the entire HPC H2020 Programme through the integration of isolated European Commission funded R&D outcomes into fully integrated HPC system prototypes; It is a key step towards establishing European exascale capabilities and solutions*."

This is the experimental step that brings together the 3 HPC pillars: Technology providers, infrastructure providers and user's communities (co-design). Figure 6 gives an overview of the current EsD proposal [11].

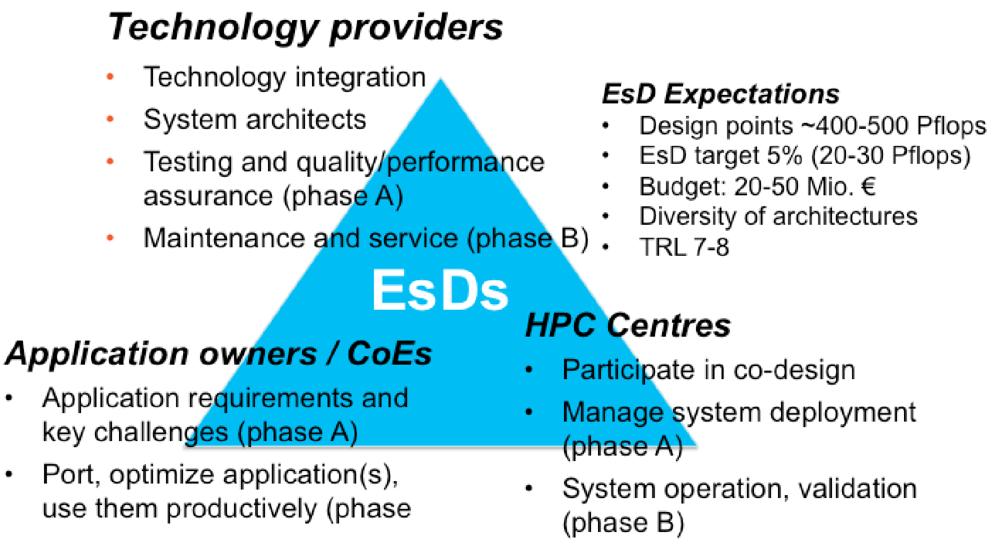

Figure 6: EsD overview.

The first set of EsDs is by far the most critical action to come. The capability to combine R&D results inside a concrete system will provide very precious feedback for the next steps toward a sustainable Exaflop machine. EsD projects will also be an organisational challenge where a R&D and roadmapping effort is transformed into an industrial mission oriented effort. Furthermore, as important as this EsD concept is, it mainly covers the design of the core supercomputer system. The integration in the Edge and in the overall EU infrastructure (e.g. Géant, IPCEI, EOSC [8]) will have to be understood and supported as well.

A set of application challenges have been set for the Extreme-scale Demonstrator by the SRA and EXDCI workgroups on applications. This is illustrated in Figure 7. In particular, with the help of the POP center of excellence [17] a first set of eight applications has been identified to serve as a basis for defining the co-design objectives.

	WG3.1 Industrial & engineering apps	WG3.2 WCES	WG3.3 Fundamental Sciences	WG3.4 Life Sciences & Health
HPC System Architecture and Components	move to GPU or others accelerators virtualisation of resources	co design, GPU/MIC move hierarchical storage	GPU, MIC, FPGA, specialised HW ok with opt. libs, memory per core (astrophysics, fusion), mem & network BW	Large width vector units, low-latency networks, high-bandwidth and large memory; fast CPU<-> accelerator transfer rates, Heterogeneous acceleration, floating-point
System Software and Management	smart runtime systems, scalable couplers, data-aware schedulers	scalable couplers, smart scheduling	scalable couplers, mix of capacity/capability	urgent computing, link with instruments, co-location of compute & data, dynamic scheduling, support for workflows
Programming Environment	use of standards, sustainability	use of DSL, solution & performance portability	task programming / runtimes, use of DSL, standards	portability, fast code driven by python interfaces
Energy and Resiliency	scalable monitoring tools, data compression, application based FT	mixed precision	reduced precision, data compression	Distributed computing techniques to handle resiliency/fault tolerance
Balance Compute, I/O and Storage Performance	in-situ/in transit post processing	in-situ/in transit post processing active storage techniques	dynamic load balancing in mesh refinement and spectral element techniques	I/O driven, in memory database, Data-focused workflows, handling lots of small files in bioinformatics
Big Data and HPC usage Models	compression of data, remote viz, UQ/ optimisation, ML/DL, mix of capacity/ capability	UQ, compression of data, multi-site experiments to support multi-model analysis experiments, in-memory analytics, HPC-through-the cloud, ML/DL	remote data viz, ML/DL	convergence HPC-HTC, integration of HPDA tools (Hadoop, Spark, …) inc ML/DL Cloud Computing security / privacy
Mathematics and algorithms for extreme scale HPC systems	ultra-scalable solvers, // in time, automatic/adaptive meshers, model order reduction , meshless and particle simulations, coupling stochastic and deterministic methods	parallelisation in time, ensemble simulations	scalable solvers, adaptive meshers, // in time, FMM and h-matrices	multiscale/physics workflows tools, ensemble simulation, model order reduction

Figure 7: EXDCI WP3 and SRA3: 5yr - specific challenges per communities (extract from EXDCI Technical Workshop, 2017, Bologna, 3-4 of July 2017, EXDCI WP3 Overview, Stéphane Requena).

6 The European HPC Ecosystem

The EU HPC Ecosystem is rich and complex. It gathers hundreds of organisms, research labs, universities, SME and larger companies involved in numerous European initiatives and projects: FET HPC [12], CoEs [13], BDVA [7], ICT R&I, EOSC [8], EuroHPC [14], European low-power microprocessor technologies [15], etc.

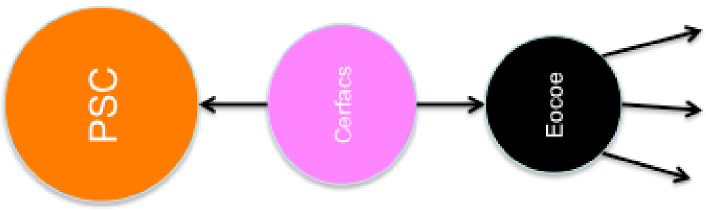

Figure 8: Extract from the graph shown in Figure 9.

Figure 9 shows a representation of the Ecosystem related to EXDCI. An extract is given in Figure 8. Larger orange nodes are EXDCI events gathering organisms around a common task. Black nodes are CoEs, FET HPC projects are in blue, research centers in pink. Other nodes are SMEs, funding agencies, universities, companies, etc. This includes about 260 entities that have been contributing into building a common vision, thanks to EXDCI.

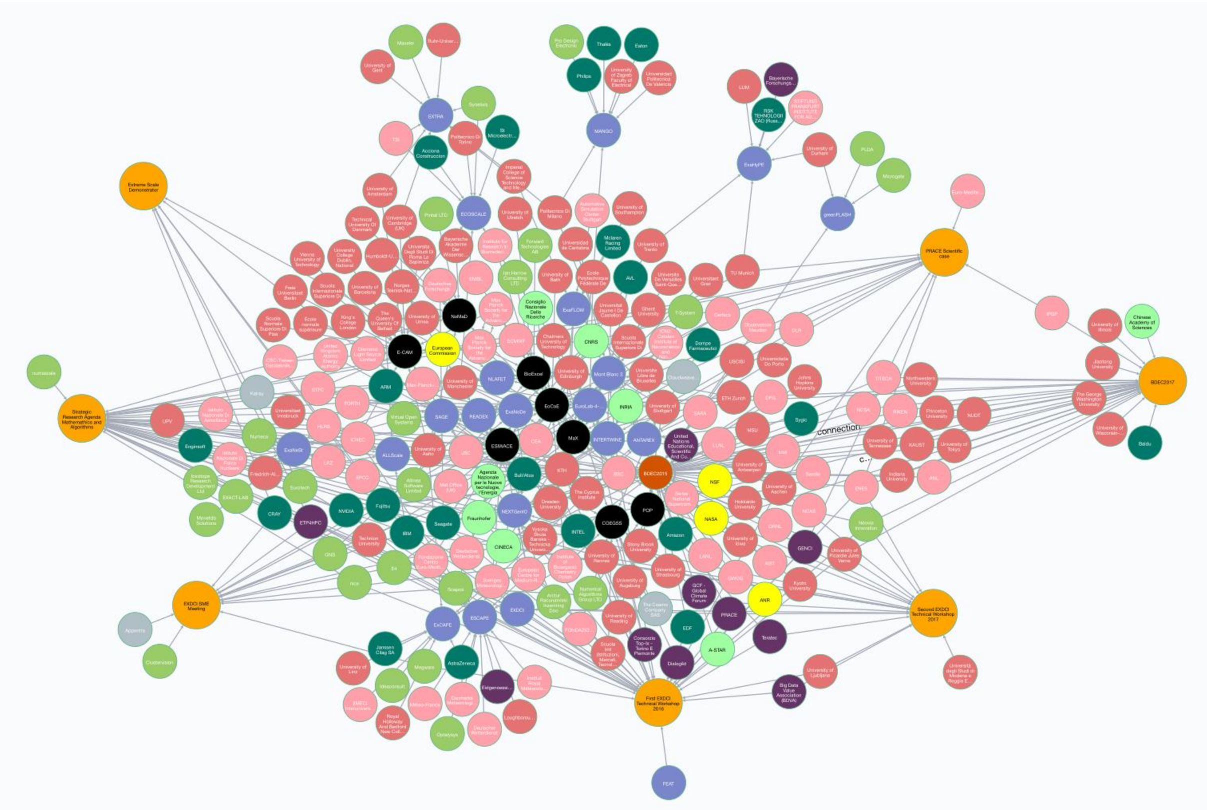

Figure 9: A map of the ecosystem as connected by EXDCI related activities (in orange, larger node).

EXDCI addresses the Ecosystem coordination using two approaches. The first method is hierarchical via PRACE and the EPT4HPC, whilst the second is transversal.

The hierarchical method aims at producing roadmaps and establishing the scientific challenges that need to be addressed. In EXDCI, this is the Strategic Research Agenda and in Prace this is the Scientific Case, lead by ETP4HPC and PRACE respectively. These two main pieces of work contribute to a common effort for the Extreme-scale Demonstrator proposal.

In EXDCI, the transversal approach has been implemented through a set of events.

The first major event setup by EXDCI has been the European HPC Summit Week, which was held in Prague (2016) and in Barcelona (2017). These events have gathered together hundreds of academics, FET HPC projects, CoEs, and company people.

The second major action has been to represent the EU partner in the international Big Data and Extreme Scale Computing (BDEC) initiative. EXDCI has also participated in the organisation of a set of BoF and workshops at major events such as Supercomputing (USA) and ISC (Germany).

Two EXDCI technical workshops were opportunities for establishing stronger links with HPC technology providers and the Big Data Value Association (BDVA).

Within EXDCI, the activities related to start-ups and SMEs focused on how the start-ups and the SMEs in HPC perceive their situation and what are the hurdles that they face. We broadened the scope of this study including the point of view of the other stakeholders. This led to a set of recommendations promoting the partnerships between SME, start-ups, large companies and HPC centers.

The Figure 10 shows the entities (about 80) directly involved in EXDCI activities. These entities are relays to the other stakeholders of the ecosystem via their participation in CoEs or FET HPC projects. The center of the graphs contains institutions that are involved in many initiatives. A live graph version of this data will be available on EXDCI website.

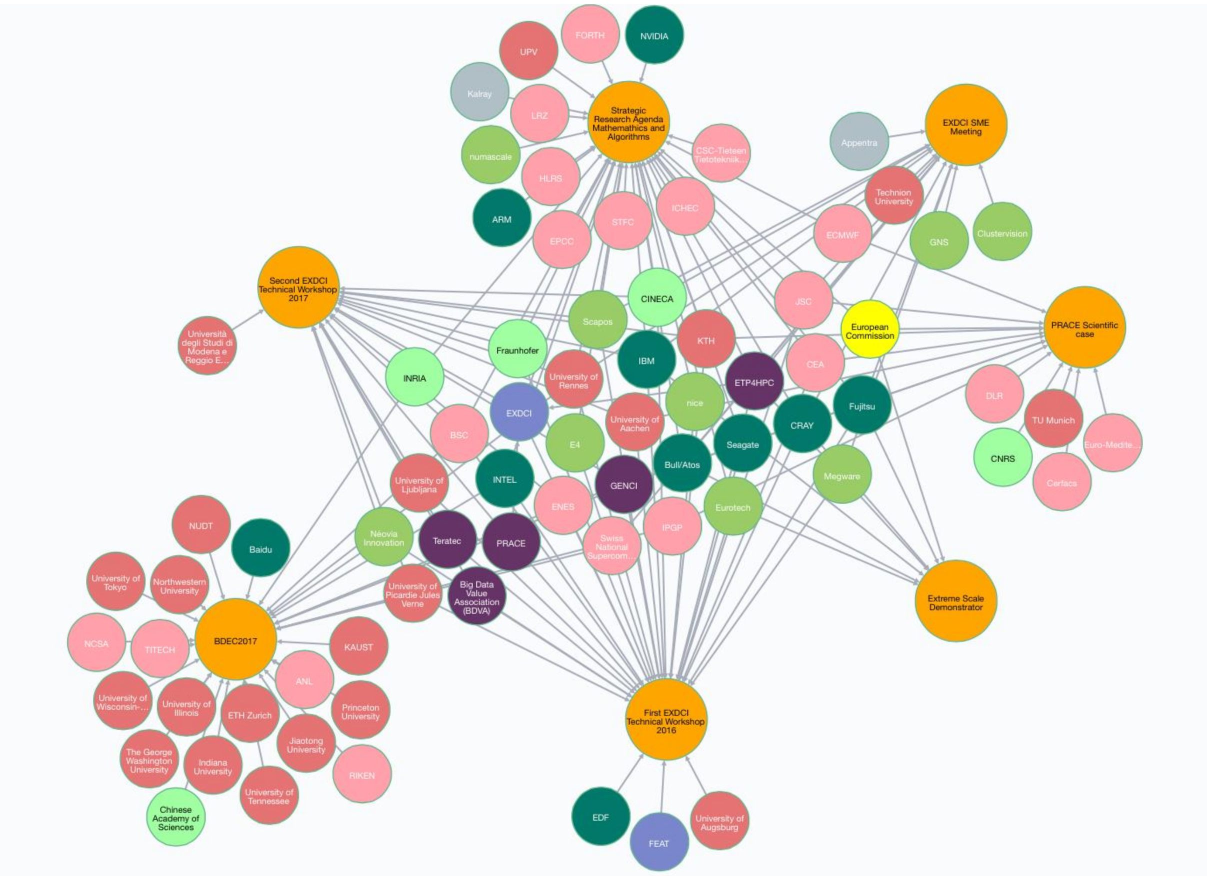

Figure 10: Entities directly involved in EXDCI related activities.

Other transversal actions include monitoring the ecosystem by establishing a Balanced Scorecard to capture its evolution, synchronising with Eurolab4HPC and promoting HPC by publishing Career Case Studies. The later aims at promoting the various interesting and exciting career opportunities that HPC can offer to young people [16].

This EXDCI transversal work has led to a set of "ecosystem" level recommendations summarised below [9]. They address three domains:

- **Better research instruments** - The following recommendations aim at improving current research instruments, both computing resources and deployment of new technologies in order to better support applications and researcher discovery processes.
    - **R1**: Design new operation policies and federations towards convergence.

  - **R2**: Reinforce Big Data and extreme scale international initiatives.
  - **R3**: Improve access to advanced technologies.
- **R&D efficiency -** The purpose of the following recommendations is to ensure that the public and private investment in R&D is carried out in a coherent manner, maximising the impact of research.
  - **R4**: IPCEI for advanced research and innovation.
  - **R5**: Paving the way from EsD development towards applications.
  - **R6**: Improving FETHPC and CoE result capitalisation.
- **Industry competitiveness** - The ambition of the following recommendations is to leverage R&D excellence and translates its output into industry competeveness:
  - **R7**: Encouraging commercial relationships between SMEs and industry via European projects
  - **R8**: Concerted approach to HPC training in Europe
  - **R9**: Incentives to increase EU stakeholders implications in standard initiatives

They are complementary to the technical recommendations proposed in the Strategic Research Agenda and in the Deliverable D3.1 “First set of recommendations and reports toward applications”. Details can be found in EXDCI Deliverable 4.3 “First holistic vision and recommendations report” published in the summer of 2016.

7 Conclusion and Final Recommendation

We are in the middle of a “paradigm shift” driven by the data and computing *convergence*. It carries many opportunities, however, it requires that the European community adapts and invents new ways for science and industry. EXDCI has contributed in better setting these new goals and challenges for the HPC community. This large-scale effort is a first step in bridging applications and technology, whilst encouraging the dialog between all stakeholders.

The full stack of software and hardware technologies must be addressed to create an efficient technological ecosystem. This set of resources will not be produced without a high level, long-term effort. Whilst other countries such as China have been able to develop state of the art technologies, the EU has seen its key suppliers such as ARM and Allinea move under Japanese control. The lack of some parts of technology, software or hardware, is very likely going to create larger issues. The European Low Power processor project is an action in the right direction.

There is a need to focus and organise current efforts in a way that is closer to an integrated industrial project rather than a set of loosely coupled research projects. These are necessary in setting the background, however, they cannot be the tip of the spear. The current situation is likely to produce a severe brain drain, since our best scientists are very likely to turn their efforts to currently available technologies, reinforcing EU competitors. There is therefore an urgent need of a more mission-oriented programme with industrial impact.

The report adds to previous ecosystem level EXDCI recommendations as a single but paramount one:

Focus and organize the current efforts in a way that is closer to an integrated industrial project in order to ensure a successful delivery of European Exascale level sustainable systems capable of serving a convergence based scientific discovery process.

Indeed, thanks to the EU community efforts and collaborations, we have now the understanding of the challenges, as well as an organised dialog between the stakeholders, to accelerate the development of the EU next generation HPC infrastructure.

8 References

[1] Intel 3D Xpoint, https://en.wikipedia.org/wiki/3D_XPoint

[2] IBM Q, https://www.research.ibm.com/ibm-q/

[3] White House National Strategic Computing Initiative Workshop Proceedings Summary of a Workshop, October 20-21, 2015, https://www.nitrd.gov/nsci/files/NSCI2015WorkshopReport06142016.pdf

[4] "Convergence Calculs-Données, Une Vision Informatique », François Bodin, April 2017, http://www.maisondelasimulation.fr/semmod.php

[5] Big Data and Extreme-scale Computing (BDEC), http://www.exascale.org/bdec/

[6] BDEC "Pathways to Convergence" Report, online http://www.exascale.org/bdec/, July 9th, 2017.

[7] Big Data Value Association AISBL (BDVA), http://www.bdva.eu/, July 9th, 2017

[8] European Cloud Initiative - Building a competitive data and knowledge economy in Europe https://ec.europa.eu/digital-single-market/en/news/communication-european-cloud-initiative-building-competitive-data-and-knowledge-economy-europe

[9] D4.3 First holistic vision and recommendations report, EXDCI report, 2016. https://exdci.eu/exdci-at-a-glance

[10] ETP4HPC – BDVA cPPPs Position Paper https://exdci.eu/sites/all/themes/exdci_theme/images/BDVA-ETP4HPC-paper-ETP4HPC-Final.pdf, 2016.

[11] The Extreme-Scale Demonstrators concept: Current State of Definition, 13 November 2016, https://exdci.eu/sites/all/themes/exdci_theme/images/EsD-Concept-Current-State-22-Jul-2016.pdf

[12] 21 new H2020 High Performance Computing projects https://ec.europa.eu/programmes/horizon2020/en/news/21-new-h2020-high-performance-computing-projects

[13] Eight new Centres of Excellence for computing applications https://ec.europa.eu/programmes/horizon2020/en/news/eight-new-centres-excellence-computing-applications

[14] EU ministers commit to digitising Europe with high-performance computing power (EuroHPC) https://ec.europa.eu/digital-single-market/en/news/eu-ministers-commit-digitising-europe-high-performance-computing-power

[15] ICT-42-2017: Framework Partnership Agreement in European low-power microprocessor technologies, http://ec.europa.eu/research/participants/data/ref/h2020/wp/2016_2017/main/h2020-wp1617-leit-ict_en.pdf

[16] HPC Career Case Studies https://exdci.eu/jobs-and-training/hpc-career-case-studies

[17] Performance Optimisation and Productivity Centre of Excellence in Computing Applications https://pop-coe.eu