

H2020-FETHPC-2014

Coordination of the HPC strategy



EXDCI

European eXtreme Data and Computing Initiative

Grant Agreement Number: FETHPC-671558

D4.3

First holistic vision and recommendations report

Draft

Version: Alpha
Author(s): François Bodin, University of Rennes 1
Date: 18.09.2016

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: D4.3	
	Deliverable Nature: Report	
	Dissemination Level: PU*	Contractual Date of Delivery: October 31, 2016
		Actual Date of Delivery: DD / Month / YYYY
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: First holistic vision and recommendations report	
	ID: D4.3	
	Version: alpha	Status: draft
	Available at: http://www.exdci.eu	
	Software Tool: Microsoft Word 2013	
	File(s): EXDCI-Deliverable-Template-Word2013-1.1.docx	
Authorship	Written by:	François Bodin,
	Contributors:	Mark Asch Thierry Bidot Guillaume Colin de Verdière Maïke Gilliot Jean Gonnord Catherine Inglis Michael Malms Jean-Philippe Nominé Stéphane Requena Jean-Pierre Vilotte
	Reviewed by:	NN, Organisation
	Approved by:	MB/TB

Document Status Sheet

Version	Date	Status	Comments
alpha	18/09/2016	Draft	

Document Keywords

Keywords:	PRACE, , Research Infrastructure, Recommendations, disruptive, Operation Policies, IPCEI, Extreme Scale Demonstrator, SME, Big Data, International, Training, R&D Capitalization, Standards, Advanced Technologies
------------------	--

Copyright notices

© 2016 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

Project and Deliverable Information Sheet	i
Document Control Sheet.....	i
Document Status Sheet.....	i
Document Keywords	ii
Table of Contents.....	iii
List of Figures	iii
List of Acronyms and Abbreviations	iii
Executive Summary.....	Erreur ! Signet non défini.
1 Introduction.....	6
2 Holistic Vision Elaboration	6
3 Operation Policies and Federation toward Convergence.....	7
3.1 Recommendation	8
4 IPCEI for advanced research and innovations	8
4.1 Recommendation	8
5 Paving the way from EsD development towards applications.....	8
5.1 Recommendation	9
6 Encouraging commercial relationships between SME and industry through European projects	9
6.1 Recommendation	10
7 Big Data and extreme scale international initiative.....	10
7.1 Recommendation	11
8 Concerted approach to HPC training in Europe	11
8.1 Recommendation	12
9 Improving FETHPC and CoE results capitalization.....	12
9.1 Recommendation	12
10 Incentives to increase EU stakeholders implications in standards initiatives	13
10.1 Recommendation	14
11 Access to advanced technologies	14
11.1 Recommendation	14
12 Conclusion.....	14

List of Figures

Figure 1: recommendations and expected impacts.....	6
Figure 2: Elaboration HPC stakeholders layers (draft schema)	7

List of Acronyms and Abbreviations

Below is an extensive the 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)
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

This report proposes a set of recommendations that aims at improving the current European HPC ecosystem. We focus on three aspects of the European ecosystem:

1. Research instruments;
2. R&D efficiency;
3. Industry competitiveness.

These three aspects form a continuum from fundamental research to exploitation. In EXDCI we have identified a set of decisive areas where new actions can help to achieve the overall goal set by the EU commission President, Jean Claude Juncker: *“Our goal is for Europe to become one of the top 3 world leaders in high-performance computing by 2020.”* The proposed recommendations are summarized in Figure 1.

Recommendations	Expected impact
Encouraging commercial relationships between SME and industry through European projects	Industry competitiveness
Concerted approach to HPC training in Europe	
Incentives to increase EU stakeholders implications in standards initiatives	
IPCEI for advanced research and innovations	R&D efficiency
Paving the way from EsD development towards applications	
Improving FETHPC and CoE results capitalization	
Operation policies and federation toward convergence	Better research instrument
Big Data and extreme scale international initiative	
Access to advanced technologies	

Figure 1: recommendations and expected impacts

These recommendations take into account the major evolutions of the HPC domain and in particular the advances in data analytics and new disruptive technologies. They are complementary to technical recommendations that are provided in the PRACE Scientific Case, the Strategic Research Agenda (SRA) and other EXDCI deliverables.

The remainder of this report provides a description of the elaboration process and details the recommendations.

2 Holistic Vision Elaboration

The elaboration process has been a collective process involving numerous stakeholders of the HPC ecosystem. It is illustrated in Figure 2. The EXDCI WPs have collected feedback and information from the HPC community at large. This includes:

1. PRACE members;
2. ETP4HPC members;

3. FET HPC projects and CoE via a questionnaire¹ and EXDCI workshops;
4. The International stakeholders via the BDEC initiative²;
5. A set of interviews of startups and SMEs.




Figure 2: Elaboration HPC stakeholders layers (draft schema)

This process has delivered a set (that we have voluntarily kept small) of recommendations that aims at proposing new actions to improve the European HPC ecosystem.

3 Operation Policies and Federation toward Convergence

Recommendation issued from EXDCI WP3

Access to data – generated by large simulations, large instruments and observational systems – is changing the ways in which we think about and address problems in science and societal research, including climate change and environmental hazard and risk. Shared data is catalyzing change in global collaboration and in businesses. A user-driven sustainable data-and-compute e-infrastructure needs to be application-oriented, easily accessible, open and agile so that it can continuously adapt to changes in technology and research practice. Case studies must, therefore, include the relevant science and technology expertise while identifying gaps in missions, resources and capabilities, and also involving appropriate business models to explore how to sustain long-term stewardship.

A crucial innovation to be explored is the agile strategies that will allow to intimately coupling research thinking with technical innovations. These strategies need to be deployed in a multi-organizational context in order to discover solutions and ways forward that closely match domain researchers' requirements, which evolve rapidly as the potential of new capabilities and data are appreciated. Eventually, this will change the focus from HPC production environments and community data services, where middleware standards, security procedures and connectivity dominated, to research-focused scenarii.

¹ <https://exdci.eu/activities/questionnaire-coe-fet-hpc>

² <http://www.exascale.org/bdec/>

3.1 Recommendation

It is recommended that a task force be formed to pioneer the federation of autonomous organizations providing data, computing and data-intensive analysis resources, together with a comprehensive and operational virtual research environment and E-infrastructure devoted to the full path of data use in a research-driven context.

4 IPCEI for advanced research and innovations

Recommendation issued from EXDCI WP2

IPCEI is a new framework meant to enable large HPC projects with strong industrial leadership, in answer to the new EC HPC policy (as described in April 2016 communications on the European Cloud Initiative, confirmed by the Council of Competitiveness of May 2016, in the more global context of the Digital Single Market and with a strong willingness to accelerate industry and public services digitalization, in addition to putting in place the necessary computing and data services and infrastructures for scientific research).

IPCEI should allow to overcome some regulation and funding limitations of conventional H2020 instruments – enabling funding for much larger projects; allowing the pooling of different sources of funding from the EC and member states, and overcoming some limitations of state aids and competition rules. H2020-funded FETHPC and CoE projects already increased the industry-research interaction, within the EU HPC cPPP and with a significant support from EXDCI. ETP4HPC is now suggesting to include “Extreme Scale Demonstrators” in Work Programme 2018-2020, so as to integrate successful H2020 R&D projects outcomes into first-of-a-kind HPC systems including significant European technology. It is recommended that EXDCI considers and actively supports the articulation between the European Cloud Initiative (and its underlying European Data Infrastructure), FETHPC projects incl. Extreme Scale Demonstrators, and IPCEI. Alignment of timelines and milestones in terms of pre-exascale then exascale systems in this global perspective is a major issue. EXDCI can bring significant added value to this process, via the stakeholders mobilization, and by leveraging its outcomes in terms of global vision and recommendations from technologies to applications.

4.1 Recommendation

It is recommended that EXDCI considers and actively supports the articulation between the European Cloud Initiative (and its underlying European Data Infrastructure), FETHPC projects incl. Extreme Scale Demonstrators, and IPCEI.

5 Paving the way from EsD development towards applications

Recommendation issued from EXDCI WP2

The EsD calls will have a high dependency on the outcome of WP2014/15 and WP2016/17 projects regarding their timing, but mostly regarding contents. The portfolio of accepted projects in the work programmes must provide a sound technology basis for building EsDs,

and the accepted projects should be actively encouraged to foster cross-project interlock. The structure of the WP16/17 should support cross-project integration, with particular regard to IP visibility and licensing clarity. Little coherence between accepted projects, too many disjoint focus areas, and insufficient technology options and readiness might otherwise jeopardise the success of the EsD calls.

It is proposed that calls should be announced within the WP 2018-2019. It is proposed that the EsD project calls will have a funding envelope compatible with a spending of €20-40M (30-50% R&D and 50-70% parts costs) per EsD project for phase A and €3-6M for phase B to cover utilities, operation-manpower and maintenance. Phase A should have a duration of 18-24 months and phase B of 24 months with a validation feedback checkpoint after 9 months. Therefore, total project duration of 32-48 months is envisaged.

The EsD characteristics will need to be further refined, however, they should deliver a high enough TRL to support a stable and effective production environment in their respective Phase B. Their impact on commercial product lines is not expected before 2020. Looking at the hardware characteristics, it is expected that the EsD architectures target scalability of applications up to 200 Pflops. This and other hardware characteristics (energy efficiency, I/O bandwidth, resiliency, etc.) will be detailed in the 2017 release of the SRA, also taking into account results from the FETHPC projects and requirements from the CoEs.

ETP4HPC recommends suitable projects to involve three types of partners for EsD projects: technology providers, application owners and HPC centers.

5.1 Recommendation

Enabling development and deployment of ambitious Extreme scale Demonstrator projects paving the way to Europe's competence in RD&I of competitive High Performance Systems Technology. These projects need new implementation instruments in the area of application & technology co-design and deployment to enable solutions of a large variety of extreme complex scientific problems.

6 Encouraging commercial relationships between SME and industry through European projects

Recommendation issued from EXDCI WP4

In the frame of WP4 of the EXDCI project, we interviewed startups and SMEs in order to understand if there are intrinsic particularities of the HPC-ecosystem making it particularly difficult for startups to emerge and for SMEs to grow. **One recurrent topic in the 13 interviews that we conducted was about “finding the first client/ the reference client”.**

This reference client can be very helpful for the final product development phase and for the product launching itself. They are typically early adopters of the novel technology offered by the startup, accepting sometimes non-finalized products. As such, they help the startup to understand the customer's needs. In an ideal case, those reference clients contribute to the final development phase of the product with know-how and experience from a user's point of view. Further, the first client may help the startup to confirm (or to adjust) its pricing strategy

and to gain insight into purchase procedures. And once the product is launched, the product benefits from support via the client's *renommée*.

In HPC, this reference client has an even more particular importance for getting involved in larger HPC systems, i.e. those classified within the Top500 HPC-systems. A typical criteria for being selected for a large system (i.e. a Top500-machine) is to have already contributed to a Top500-machine³, which leads to a vicious circle: You won't get accepted for a Top500-machine as long as you have not been selected for a Top500-machine. So finding a reference client is crucial in this set-up.

To find such an early client is generally difficult for startups. However, it seems particularly difficult for HPC-startups. The first reason is that in most cases HPC-startups offer highly-technological products or services. The group of potential customers also being able to take the role of an expert-adviser is small as to the high technological level (compared for example to a startup in online-sales).

A second reasons may be the nature of the HPC market itself: public entities take a considerable part of the market, based on public procurement procedures. As pointed out in D4.2, those procedures request effort, experience, and time - assets a startup is often lacking. Thus, the collaboration with a larger industrial partner is considered useful for tenders.

Within European R&I projects, startups and larger companies get acquainted and collaborate in a less competitive context with complementary skills. European research projects are thus a way to build friendly connections between startups and SMEs on the one hand and larger companies on the other hand. This has also been pointed out by the SME- working group of ETP4HPC. In their 2015 position paper, they advocate for "co-design and co-development partnerships" [REF, recommendation No. 10]. We share the belief that the full potential of those friendly connections is not tapped yet and are convinced that they are a solid basis for commercial relations.

6.1 Recommendation

Larger companies should be strongly encouraged to extend their collaboration with start-ups and SMEs beyond European R&I actions and to engage direct business relations.

7 Big Data and extreme scale international initiative

Recommendation issued from EXDCI WP6

The convergence of scientific HPC systems towards Numerical Laboratories (extreme data combined with extreme computing) requires new insight into the manner major international scientific initiatives are organized. In particular, it is expected that the current focus on compute capabilities will shift towards data acquisition (from sensors, instruments, social media, etc.) and data analytics capabilities. Sharing such assets at a global level requires broader integration of international research.

The European Open Science Cloud (EOSC), the European Cloud Initiative (ECI) and the IPCEI (see Section 3) are three pillars of the Digital Single Market strategy of the European

³ <https://www.top500.org>
EXDCI - FETHPC-671558

Commission. EXDCI can play a role in ensuring that these initiatives maintain coherency with FET-HPC projects and, at the same time, with international roadmapping efforts that are regularly discussed in the BDEC (Big Data and Exascale Computing) forum.

The tendency is now focused on the **convergence** between “big data” and “big compute”. This can only be achieved at a global scale and requires international coordination. Initiatives such as Belmont Forum and Future Earth can provide excellent use cases for convergence and subsequently inform global policy on future investments in e- and data-infrastructures. There are also direct consequences for data policies and for human capacity development.

One example is smart cities that require data analysis, sensor processing, simulation and machine learning, all of this in a context of big data. These data need to be analyzed rapidly and hence coupled with adaptive HPC systems, workflows and technologies. There are a large number of data and data-intensive computing challenges in Earth and Universe sciences. The EU VERCE project has addressed some of these.

7.1 Recommendation

The recommendation aims at encouraging the EU to facilitate future international collaborations involving compute-intensive exploitation of scientific data, as advocated by the BDEC consortium. This should be an integrated action, within the EOSC (European Open Science Cloud) initiative, that could associate other major international initiatives in "Data Management and e-Infrastructures", notably those of Belmont Forum and Future Earth.

8 Concerted approach to HPC training in Europe

Recommendation issued from EXDCI WP5

The EC has stated the aim of attaining European leadership in the supply and use of HPC systems and services by 2020. European research is successfully delivering many technological advances, but no matter how good any new technology may be, supply cannot lead to widespread take-up and effective use unless high-quality training is available and can be easily identified by those who need it.

EXDCI aims to consolidate a synchronized European HPC Community, and one aspect of this is to establish a community of training providers. By bringing the training providers together into a single community, training gaps can be identified, duplication of effort can be spotted, and resources can potentially be pooled to develop new joint training material and a coherent course timetable, whether for “in person” courses or online teaching.

EXDCI works closely with a disparate community of European technology stakeholders and applications stakeholders, including academic and commercial HPC developers and users, as well as with experts outside Europe. Many of these stakeholders either provide training and/or have a need for it. EXDCI is therefore well-placed to match those who need training with the experts who can provide it.

As a result of work carried out under EXDCI, we propose to compile a comprehensive and easy-to-find on-line catalogue of existing courses and training material, making this easily accessible by the HPC community at large. A Training Gap Analysis will inform a Training Roadmap for interdisciplinary working. The aim is to bring together training providers and

those who need training, from both academic and commercial arenas, covering all layers of the European HPC and Big Data ecosystem, encompassing both hardware and software, including programming models and applications, and covering traditional HPC and emerging fields such as disruptive technologies, data analytics and machine learning, in order to provide complete and comprehensive training opportunities for the whole community.

8.1 Recommendation

A concerted approach to HPC training in Europe must be put in place to allow rapid uptake by developers and end-users alike of the new HPC technologies currently being developed through numerous European initiatives. This will allow the European HPC community to keep up-to-date with new technologies and techniques and to become confident in their use, thus helping to secure Europe's competitive edge in the international arena.

9 Improving FETHPC and CoE results capitalization

Recommendation issued from EXDCI WP4

FET projects are major producers of novel technologies. Most of these innovative results are taking the form of new software or hardware components. Many of them aims at being proof of concepts, solely for the sake of research, that do not need to survive beyond the projects. Other components may deserve to be pushed further to reach a community of users and/or a market. This is rarely achievable in the timeframe of the projects.

The series of Mont-Blanc projects (<http://www.montblanc-project.eu>) is an interesting example. Started in 2011 the first project established the foundation of the approach. It is only today in 2016 that the results are starting to reach the market via Bull-Atos, who took over the coordination of the third Mont-Blanc project. In this case, continuing effort in a structured manner via FET projects and industrial partnership has made it possible to potentially generate new commercial revenue. Of course, there are not one unique ways to ensure technology development on the long haul until it gets to the users.

For every innovative result the right pathway must be invented (new projects, investors, open-source communities, startups...). It is very rare that a new technology reaches the market in an organic and effortless manner. Resources and dedicated people are needed all the way and beyond.

User-reaching plan must be established very well in advance of the end of the project. Therefore we recommend that specific best practices and support actions be set up in a systematic way to organize market/community-reaching plan within early stage of projects. A large variety of stakeholders (private and public) should be involved in these actions.

9.1 Recommendation

We recommend that specific best practices and support actions be set up in a systematic way to organize market-reaching plan within early stage of FET projects.

10 Incentives to increase EU stakeholders implications in standards initiatives

Recommendation issued from EXDCI WP4

International standards de-facto-standards are strongly influencing research experiments, further developments, and industrial products. In the HPC domain they provide long-term visibility for developers, and are warrantors of a widely shared and agreed on basis for further developments. As such, they are the most used basis for large software and hardware products.

Contrary to some application domains such as represented by ETSI or the OpenFOAM initiative (now driven by the ESI group³), the European HPC community is not much involved in technological-oriented standards specifications efforts. For instance, in OpenMP⁴ only ARM is a Permanent Members of the Architecture Review Board (ARB) (13 members), the most influential committee of the association. Europe's representation is stronger at the Auxiliary Members of the ARB (14 members): BSC, Bristol University, EPCC, and RWTH Aachen University.

In the case of OpenMPI⁵, European technology providers are more present, with ARM and Bull being contributor to the initiatives.

Regarding the Big Data community, whose technology is becoming crucial to future HPC applications, the EU stakeholder influence is close to nil. For instance EU stakeholders are not present in the Hadoop Project Management Committee⁶. Overall there is a lack of involvement in technology-oriented initiatives contrary to application-oriented initiatives

This absence has multiple consequences for Europe's position in HPC: First, standards specifications are by construction best suited for the products and technologies represented by the most active members. Not being involved thus means the necessity to adapt to choices and decisions taken by parties following their interests and objectives. Second, not being involved in the specification effort introduces a delay in the acquisition of the corresponding knowledge (it is usually very clear to committee members which direction is being taken year(s) in advance). The effort to catch-up is then important and there is an inherent delay to reach the market.

The low degree of involvements is not due to a lack of competencies in the EU ecosystem. We believe that it is mainly due to a lack of incentive for actions that have long-term effect and require a strong commitment. Indeed, being influential in these bodies implies the need to allocate time and resources for highly skilled experts to attend meetings, as well as for development engineers to perform the necessary experiments to build the technical contributions.

If the current status is to be changed, i.e., to increase participation in standardization bodies, it is mandatory to define efficient incentives.

⁴ <http://openmp.org/wp/about-openmp/>

⁵ <https://www.open-mpi.org/about/members/>

⁶ <https://hadoop.apache.org/who.html>

10.1 Recommendation

Incentives for EU stakeholders to participate in international standardization in the extreme scale computing and big data must be increased. New forms of support must be invented to ensure the presence of high-profile scientists and EU industry stakeholders in existing and emerging initiatives.

11 Access to advanced technologies

Recommendation issued from EXDCI WP4

Emerging technologies are strongly influencing software and hardware design. Many researches are depending on early access to these technologies to ensure they are performed in a timely fashion. Late access is strongly degrading competitiveness of an ecosystem (lack of trained people, software not up-to-date...). Being late on a highly technological market, such as the HPC one, usually means low profits or worse disappearance from the market. For instance, when GPGPU appeared on the market, only single precision floating-point operations were available. Some groups decided to provide double precision by developing software libraries. By the time the library was operational, new GPU with double precision capabilities were on the market, making *de facto* the libraries obsolete. Roadmap knowledge and access to advanced technologies are of paramount importance when planning R&D activities.

Unfortunately, access to advanced technologies is dependent only on the developer/provider wishes and strategy. Most HPC technologies being originated in the USA, the EU, especially the academic community, has many difficulties to get access (even under NDA) to new innovations. The result is a research that is based on off-the-shelf products and by essence the leadership is biased towards non EU researchers. While some topics are insensitive to such considerations, research in HPC is by essence very linked to the most advanced technology.

A given research team is rarely in position to have enough incentive (except of course if the innovation is intended to go the other way) or funding to be provided access to the most advanced technology. Only a heavyweight stakeholder (e.g. PRACE) can have an influence great enough to change the balance of power. Therefore we recommend that the EU commission creates a body able to facilitate access (technical and legal) to advanced technologies to research groups.

11.1 Recommendation

We recommend that the EU commission ecosystem creates a body able to facilitate access (technical and legal) to advanced technologies to research groups.

12 Conclusion

This report presents a set of recommendations aiming at increasing the European Research instruments, the R&D efficiency as well as the industry competitiveness. They have been elaborating after collecting information from numerous stakeholders of the HPC ecosystem.

This is the first version of this EXDCI deliverable. It is intended to go back the community, in particular the FET HPC projects and the CoEs to complement this version and produce second version of this deliverable.