# H2020-FETHPC-2014

## Coordination of the HPC strategy

## EXDCI

## European eXtreme Data and Computing Initiative

**Grant Agreement Number: FETHPC-671558**

# D2.2
# ETP4HPC Strategic Research Agenda V3

***Final***

Version: 0.3
Author(s): Michael Malms, Jean-Philippe Nominé, Marcin Ostasz (ETP4HPC)
Date: 16.02.2018

# 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:** D**2.2** | |
| | **Deliverable Nature:** Report | |
| | **Dissemination Level:** PU * | **Contractual Date of Delivery:** 28/02/2018 |
| | | **Actual Date of Delivery:** 16/02/2018 |
| | **EC Project Officer:** Evangelia Markidou | |

* - **PU** – Public, as referred to in Commission Decision 2991/844/EC.

# Document Control Sheet

| | | |
|---|---|---|
| **Document** | **Title: ETP4HPC Strategic Research Agenda V3** | |
| | **ID: D2.2** | |
| | **Version:** 0.3 | **Status:** Final |
| | **Available at:** http://www.exdci.eu | |
| | **Software Tool:** Microsoft Word 2013 | |
| | **File(s):** D2.2 | |
| **Authorship** | **Written by:** | Michael Malms, Jean-Philippe Nominé, Marcin Ostasz (ETP4HPC) |
| | **Contributors:** | Costas Bekas, IBM<br>Paul Carpenter, BSC<br>Marc Duranton, CEA<br>Thomas Eickermann, FZJ/JSC<br>Hans-Christian Hoppe, INTEL<br>Guy Lonsdale, SCAPOS<br>Sai Narasimhamurthy, SEAGATE<br>Laurent Cargemel, ATOS/BULL<br>Dirk Pleiter, FZJ/JSC<br>Pascale Rossé-Laurent, ATOS/BULL<br>Igor Zacharov, EUROTECH<br>Erwin Laure, KTH<br>Stephane Requena, GENCI<br>Mark Asch, U-PICARDIE<br>François Bodin, IRISA<br>Maike Gilliot, ETP4HPC |
| | **Reviewed by:** | Oriol Pineda, BSC<br>John Clifford, EXDCI PMO |
| | **Approved by:** | MB/TB |

# Document Status Sheet

| Version | Date | Status | Comments |
|---|---|---|---|
| 0.1 | 23/01/2018 | Draft | First full version (following internal ETP4HPC review). |
| 0.2 | 24/01/2018 | Draft | Draft send for PMO review. |
| 0.3 | 16/02/2018 | Final | Final version submitted to EC. |
| | | | |
| | | | |
| | | | |

# Document Keywords

| | |
|---|---|
| **Keywords:** | ETP4HPC, Strategic Research Agenda, Multi-Annual Roadmap, HPC Technologies, Research and Innovation |

# Table of Contents

# List of Figures



# References and Applicable Documents

[1] http://www.exdci.eu

[2] http://www.etp4hpc.eu

[3] http://www.prace-ri.eu/

[4] http://www.prace-ri.eu/prace-the-scientific-case-for-hpc/

[5] http://www.etp4hpc.eu/strategic-research-agenda/

[6] https://ec.europa.eu/digital-agenda/en/high-performance-computing-contractual-public-private-partnership-hpc-cppp

[7] http://www.etp4hpc.eu/news-archive/etp4hpc-signs-agreement-to-form-a-contractual-public-private-partnership-for-european-high-performance-computing/

[8] https://ec.europa.eu/programmes/horizon2020/en/h2020-section/high-performance-computing-hpc

[9] https://ec.europa.eu/programmes/horizon2020/en/news/overview-eu-funded-centres-excellence-computing-applications

[10] http://ec.europa.eu/research/participants/portal/desktop/en/opportunities/h2020/calls/h2020-fethpc-2014.html

[11] http://ec.europa.eu/research/participants/portal/desktop/en/opportunities/h2020/calls/h2020-fethpc-2016-2017.html

[12] High Performance and Embedded Architecture and Compilation https://www.hipeac.net/

[13] Big Data and Extreme Computing, http://www.exascale.org/bdec/

[14] http://www.bdva.eu/

[15] http://www.eurolab4hpc.eu/

[16] http://eurohpc.eu/

# List of Acronyms and Abbreviations

| | |
|---|---|
| BDEC | Big Data and Extreme-scale Computing |
| BDVA | Big Data Value Association |
| BW | Bandwidth |
| CoE | Centres of Excellence for Computing Applications |
| cPPP | contractual Public-Private Partnership |
| CSA | Coordination and Support Action |
| DX.Y | Deliverable Number X.Y |
| DoA | Description of Activity |
| EC | European Commission |
| EsD | Extreme scale Demonstrators |
| EU | European Union |
| EXDCI | eXtreme Data and Computing Initiative |
| FET | Future and Emerging Technologies |
| FET-HPC | HPC part of FET programme |
| FP7 | Framework Programme 7 |
| H2020 | Horizon 2020 – The EC Research and Innovation Programme in Europe |
| HPC | High Performance Computing |
| HiPEAC | European Network of Excellence on High Performance and Embedded Architecture and Compilation |
| IESP | International Exascale Software Project |
| IoT | Internet of Things |
| ISV | Independent Software Vendor |
| IT | Information Technology |
| LEIT | Leadership in Enabling and Industrial Technologies |
| M | Month |
| PM | Person Month |
| Q | Quarter |
| R&D | Research and Development |
| R&I | Research and Innovation |
| SKA | Square Kilometre Array |
| SRA | Strategic Research Agenda |
| TRL | Technology Readiness Level |
| WG | Working Group |
| WP | Work Package |

# Executive Summary

ETP4HPC's Strategic Research Agenda (SRA) is a multi-annual roadmap which outlines the European research priorities in the area of HPC technology. The provision of HPC technology is viewed here as one of the three pillars of European HPC, alongside infrastructure and application expertise. The importance of HPC is now widely recognised and European HPC has a good momentum due to the investments made and programmes launched. Europe is improving its position to compete with other regions in the area of HPC technology provision and matches its weight in the consumption of HPC systems.

This SRA suggests that the development of European HPC technology should continue along the seven main research lines (HPC System Architecture and Components, System Software and Management, Programming Environment, Energy and Resiliency, Balance Compute, I/O and Storage Performance, Big Data and HPC usage Models, Mathematics and algorithms for extreme scale HPC systems); it also includes the concept of prototyping ('Extreme-Scale Demonstrators') to test the readiness of the European technology projects and resulting building blocks, vendors and users to produce a globally competitive HPC system. The area of Big Data has received a special attention and work will continue to synchronise HPC technology and Big Data solutions. We also believe that two other areas, namely support mechanisms for SMEs and Start-ups and Education and Training needs to remain as priorities to help Europe compete in this complex market.

This is the third Strategic Research Agenda (SRA3) and its second fully revised issue (following the first full issue in 2013). As in the case of the previous versions, the process of writing the SRA was an open one, with all the members of ETP4HPC having an opportunity to contribute their expertise. This is the result of the collective work of nearly two hundred and thirty experts in eight technical working groups as well as some other non-technical task forces and technical and non-technical experts.

# 1 Introduction

The role of SRA3 is to define Europe's roadmap towards HPC technology provision[1] at Exa-scale and beyond (and, as a consequence, throughout the broader ICT landscape). The key motivation of ETP4HPC is to increase the global market share of the HPC technology developed in Europe, while contributing to a competitive European HPC value chain via infrastructures and applications. This means that future systems produced in Europe need to be able to compete with systems from other geographies. These systems need to meet the requirements of European (and global) scientific and industrial users and facilitate the pervasive use of HPC [5].

In this process, ETP4HPC, as part of the European HPC Ecosystem [8], issues its SRA, which is then used by the EC as a recommendation in formulating its research programme. In the next step, the EC's open calls for proposals are announced and the guidelines included in this SRA3 are expected to be used as a reference for the call objectives and assessment criteria for project proposals submitted to the EC's HPC technology Work Programme (WP) 2018-2020. Any project proposal submitted within the above Programme following the issue of the SRA should address its research recommendations, defined as research priorities and research milestones. In particular, as SRA 3 is being issued in Q4 2017, any project funded under WP18-20 (the last part of the EC's Horizon2020 Framework Programme) should cover milestones included in this document.

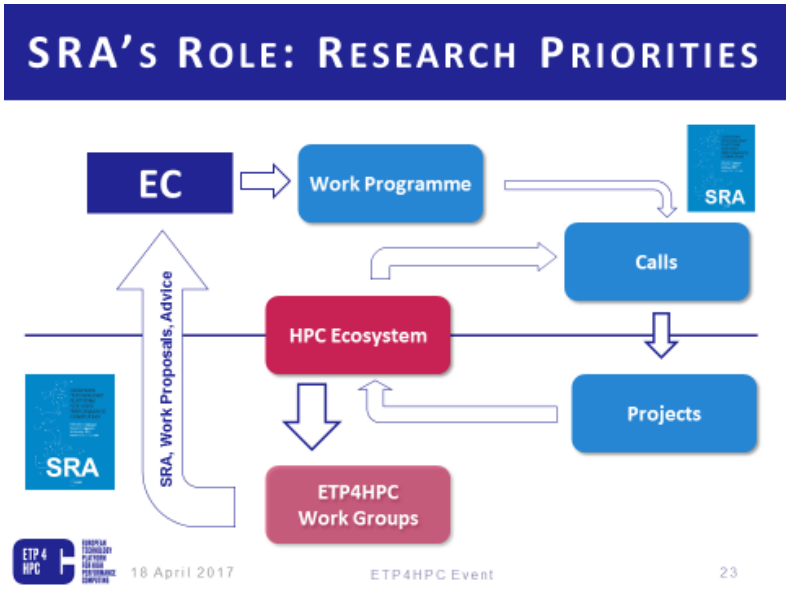


*The SRA influences the process of defining EC calls for proposals in the area of HPC technology. The guidelines included in the SRA are expected to be used as a reference for the call objectives and assessment criteria for project proposals submitted.*

**Figure 1: The role of SRA in the H2020 process**

[1] The term 'technology' is used here to denote the entire HPC system stack as per the ETP4HPC HPC Model (Section **Error! Reference source not found.** of the SRA) – hardware components and subsystems, whole hardware stack, up to services and solutions.

Throughout the H2020 timeline, there is a periodic SRA update process, validated by the European HPC ecosystem, so as to serve as a reference for the projects participating in the EC HPC research programme [10][11]: the FET-HPC[2] part of the programme concerns the development of basic HPC technology. Extreme Scale Demonstrators (EsD) build on FET-HPC building blocks and are in LEIT/ICT programmes. The CoE[3] sub-programme supports Centres of Excellence in Computing Applications, consolidating the European HPC application expertise [9]. The ecosystem development is supported by a series of Coordination and Support Actions[4], which orchestrate the European HPC strategy [1],[15]. It is important to note that some relevant elements of the European HPC effort now fall into other programme parts such as LEIT[5], e.g. afore mentioned EsDs (with higher TRL ambitions) and also joint HPC/Big Data applications (so-called 'testbeds').

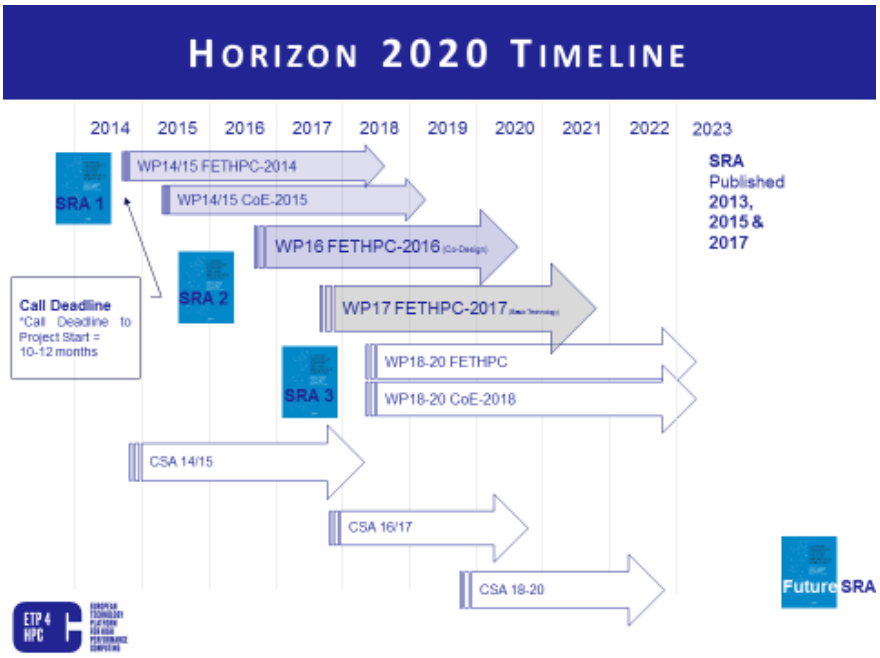


*The timings of the SRA and various Horizon 2020 EC HPC technology work programme parts- there is a valid SRA in place at any time within the programme.*

**Figure 2: SRA editions timeline and link with H2020 structure**

After Chapter 1 (Introduction), this deliverable describes the process for SRA3 elaboration in Chapter 2. Then it gives some highlights regarding the main findings and new aspects of SRA3.

[2] Future and Emerging Technologies – High-Performance Computing
[3] Centres of Excellence (in Computing Applications)
[4] CSA
[5] Leadership in Enabling and Industrial Technologies

# 2 The process of preparing SRA 3

The process of preparing SRA3 began in March 2017 with the definition of all steps required. As a first step, an analysis of the applications' requirements was identified. In this step, ETP4HPC collected input into the SRA from the following areas:

- Scientific Applications: Centres of Excellence in Computing Applications - CoEs [9] - and PRACE [3] Application (Scientific) users - through the 'Scientific Case 2017' [4]
- Industrial users,
- BDEC's recommendations [13],
- HiPEAC's vision [12]

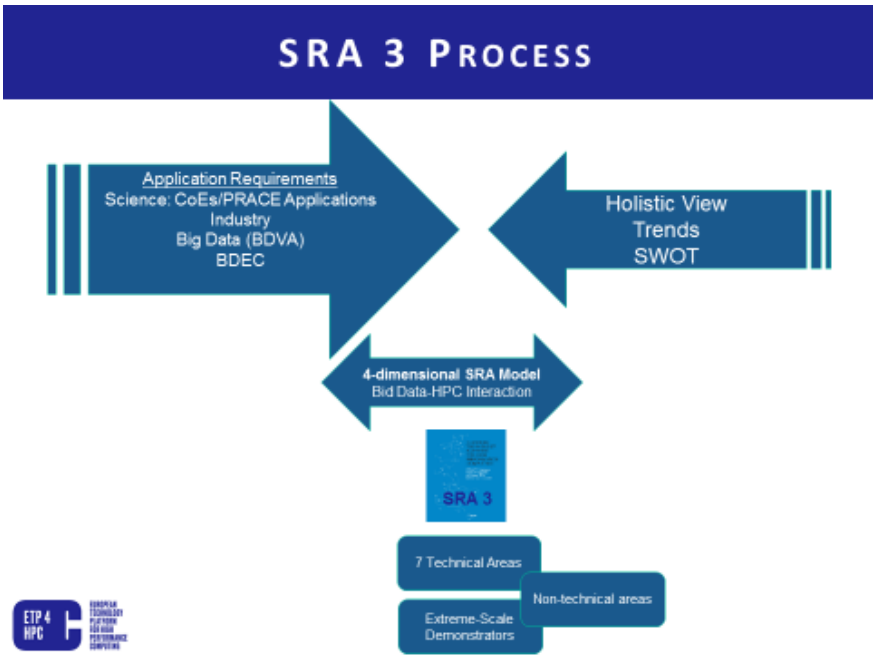


**Figure 3: The process of preparing SRA 3 - inputs used, analysis undertaken, outputs**

A special emphasis was placed on the needs of the European Big Data Community represented by BDVA [14].

Then, an analysis of the current state of European HPC technology was carried out. In this part, a holistic view of it and some recommendations are included. Global technology and HPC market trends are analysed and summarised.

The results of these analyses are fed into the following eight working areas, in line with the previous issue of the SRA:

- HPC System Architecture and Components;
- System Software and Management;
- Programming Environment;
- Energy and Resiliency;
- Balance Compute, I/O and Storage Performance;
- Big Data and HPC Usage Models;
- Extreme-Scale Demonstrators ;

and also non-technical areas: Education and Training, SMEs and Start-ups.

All ETP4HPC members were invited to participate in technical working groups (WGs) mirroring the categories above and led by selected ETP4HPC member organisations. The task of these working groups was to define the research priorities and milestones in the corresponding areas. This work took place through conference calls within the groups and a workshop involving all working group leaders. All ETP4HPC members have been able to review and comment on the last draft of this document before it reached the approval stage at the ETP4HPC Steering Board level.

Within this process a special emphasis has been placed on the concept of the Extreme-Scale Demonstrators – prototypes of European exascale supercomputers (EsDs). To facilitate the definition of the scope of these systems, a total of five[6] workshops has been held in 2016 and 17, with two of these events falling within the framework of SRA 3:

- A 'Round-table' workshop aiming to present the potential contributions of the European HPC technology projects, involving also the Centres of Excellence in Computing Applications (CoEs), system integrators and HPC Centres (May 2017, during the EXDCI European HPC Summit Week [1])
- A workshop dedicated to the industrial use of the EsDs, involving industrial users of HPC and Independent Software Vendors (ISVs). The invited parties were requested to present their application domains and describe their interest in participating in future EsD projects.

The second area of focus in this SRA is Big Data and its HPC system requirements. In order to synchronise the technology of these two areas a workshop involving ETP4HPC and BDVA (Big Data Value Association) was held in July 2017, within the scope of this SRA[7].

[6] These five EsD-related workshops have provided an opportunity for the following entities to express their comments on the concept of the EsDs and their potential contributions to them: 1/all FETHPC and other technology projects, CoE projects and system integrators, 2/all CoE projects and other application users, 3/ system integrators, and (within the scope of SRA 3) 4/all FETHPC and other technology projects and 5/industrial HPC users and ISVs.

[7] Another Big Data-related workshop was held in Sept 2016 with an objective to analyse selected Big Data use area where HPC could contribute.

# 3 Highlights from SRA3

## 3.1 Outline

SRA3 outline stems from the SRA history since the first edition in 2013 and successive amendments and add-ons, as described in Chapter 0 of this report. The overall SRA structure and scope however remained quite stable since 2013.

Chapter 1 of SRA3 introduces the objectives of the document and its context. Chapter 2 presents 'the case for European HPC technology'. It also summarises the current European HPC strategy in a short review of the current tools and mechanisms in place which serve to develop European HPC. Chapter 3 contains a review of the European HPC ecosystem and gives a short overview on the status of HPC in other geographies. Chapter 4 explains the structure of the technical Chapters (5 to 8) and contains a detailed description of the requirements for future research priorities meeting the needs of scientific and industrial users. This is followed by the technical (5 – trends, 6 – research priorities, 7 – milestones and 8 – Extreme-Scale Demonstrators) and non-technical (9) areas of this roadmap. Finally, section 10 provides the conclusions.

## 3.2 The Case for HPC Technology

SRA team worked on application requirements with CoEs/PRACE for Science, Industry, Big data (BDVA), and with both international and European networks developing global visions for computing – incl. BDEC and HiPEAC.

Extreme-scale requirements are the drivers but there is a general motivation to develop HPC at all scales, building on common components. Expansions of uses (Big Data, Machine Learning, etc.) are also important axes, as well as less technical ones like education and training and the support of SMEs.

The software aspects are taken very seriously. There is a common understanding for application requirement:

- Portability and maintainability is a key concern for applications
  ⇨ interfaces / standards are needed
- Productive coding and performance tuning
  ⇨ smart underlying system software layers are needed
- Data requirements
  ⇨ high bandwidth memory and networks + minimize data circulation are to be considered
- Energy management
  ⇨ profiling tools are needed
- Fault tolerance and resilience
  ⇨ this is to be considered on the algorithmic/solver level, as well as on the runtime/programming environment level
- Support complex end-to-end workflows, smart coupling of applications
  ⇨ not only capability based applications able to scale out on a full machine)

## 3.3 The EU HPC ecosystem

Since the first release of the SRA in 2013, the European HPC ecosystem has grown substantially in terms of players, mostly on the application and user side. The technology providers, few in number, will generate significant more revenue than in the previous years, but they still lag behind compared to the providers in other geographies. This is the stimulus for a significant initiative by the EC and thirteen EU Member States (EuroHPC [16]) to push the provision of Exascale HPC infrastructures in Europe as well as R&D in the fields of HPC technology, including processor chips, system software and applications.

Besides the known players as ETP4HPC and PRACE significant projects such as HiPEAC and the BDEC forum attract many experts providing analysis, advice and recommendations for further steps to strengthen the role of HPC provisioning in the context of industrial use (e.g. HPC in cyber physical environments, HPC-in-the-loop) and 'extreme compute-extreme data' – challenges.

New competence centres on the application side are now well established (CoEs) and they are entering their second round of project activity. These entities have provided more than ten pages of application-driven input to this release of the SRA (Chapter 4.3). Together with ETP4HPC, the CoEs are the private part of the regular cPPP meetings with the EC [6][7].

Besides EXDCI, there is another Coordination and Support Action (CSA): EuroLab4HPC [15], focusing on boosting the European HPC-related research.

## 3.4 Technical areas and non-technical dimensions

The technical chapters are organised again following the 4 dimensional-HPC vision, augmented by the subject of "Extreme scale Demonstrators" (EsDs) laid out in detail in SRA3 Chapter 8. The organisation of topics along the "north dimension - major elements of the HPC compute stack", the "east dimension - major system characteristics needing continuous innovation", the "south dimension-major new use models such as BD, ML, etc." and the "west dimension- non-technical dimensions" has proven to provide a logical structure to the recommendations made.

A highlight is this SRA3 release is an emphasis on the increasing importance of High-Performance Data Analytics and Artificial Intelligence/Machine Learning and the interaction with HPC. A side-by-side structuring of the three compute stacks lays the foundation for upcoming extended research priorities.

A special effort has been made in this release to cross-relate the Research Milestones of the seven technical domains shown in Figure 4 (see below) with each other, and the application requirements mentioned above.

The non-technical SRA3 Chapter 9 gives latest views and recommendations from SMEs on the technology provider side in the context of political and market-related issues. The EXDCI Work Package on Education and Training eventually explains its findings and suggestions in improving the availability of deep HPC technical skills (this activity took over from initial ETP4HPC efforts and reflections on these aspects).

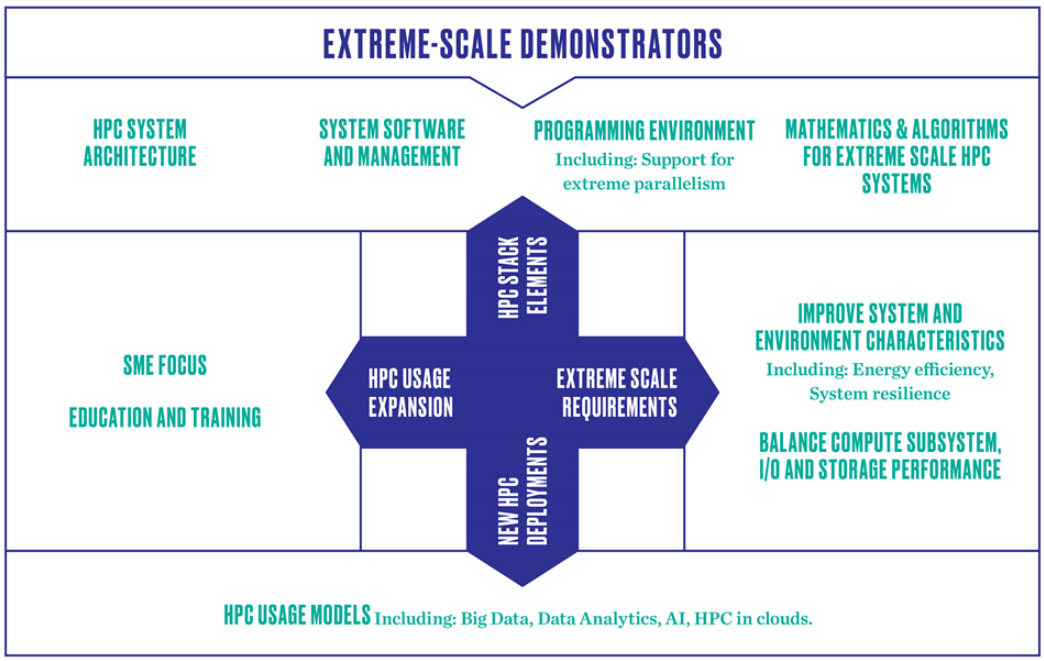


*The new modified 4-dimensional model of European HPC technology development stemming from its first version as defined in SRA 1 (2013).*

**Figure 4: The four-dimensional model of European HPC technology development**

# 4 Conclusion and next steps

This document provides a research roadmap for the remaining years of the H2020 framework programme. The awareness of the key role HPC plays in science, industry and in our everyday lives has increased dramatically over the past five years. The investments made at a pan-European public level demonstrate that HPC is not anymore a narrow niche of pure technical computing.

This effort of road-mapping Europe's HPC technology strategy needs to continue despite the existence of other tools and mechanisms aimed at producing competitive European supercomputers. ETP4HPC welcomes and supports these initiatives but also claims that there will be a persistent need for basic research, prototyping and aligning priorities with other stakeholders, all of which should take place in the context of European competitiveness in science and industry. This is one of the roles that the future SRAs should perform.

It is obvious that High Performance Computing, the vast number of Big Data use cases and fast-growing world of the "Internet of Things" cannot be seen as three separate silos. While each of the domains will keep having its own focus areas and priorities in the future, the interdependence is explicit in the case of important large projects such as SKA, autonomous driving, energy management, etc. In this context embedded and networking technologies will play a considerable role.

Therefore, at a European level, the forthcoming road-mapping efforts also need to be carried out in much closer cooperation between private/public bodies covering the domains of HPC, Big Data and IoT including expert groups representing the embedded and networking domains. The work performed lately together with BDVA and HiPEAC is a start in this direction.

# 5 Annex: Strategic Research Agenda 3 – Full text

The full version of SRA 3 is available at www.etp4hpc.eu/sra.