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

EXDCI-2

European eXtreme Data and Computing Initiative - 2

Grant Agreement Number: 800957

D2.4
Report on coordination of the technology research action in Europe

Final

Version: 1.1
Authors: JF Lavignon, TS-JFL; Marcin Ostasz, ETP4HPC; Thierry Bidot, Neovia Innovation
Date: 26/05/2020

Project and Deliverable Information Sheet

EXDCI Project	**Project Ref. : FETHPC-800957**	
	Project Title: European eXtreme Data and Computing Initiative - 2	
	Project Web Site: http://www.exdci.eu	
	Deliverable ID: < D2.4>	
	Deliverable Nature: <Report>	
	Dissemination Level: PU	**Contractual Date of Delivery:** 31 / 05 / 2020
		Actual Date of Delivery: 28 / 05 / 2020
	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: Report on coordination of the technology research action in Europe**	
	ID: D2.4	
	Version: <1.1>	**Status:** Final
	Available at: http://www.exdci.eu	
	Software Tool: Microsoft Word 2013	
	File(s): EXDCI-2-Deliverables-D2-4 v1.1	
Authorship	**Written by:**	JF Lavignon, TS-JFL; Marcin Ostasz, ETP4HPC; Thierry Bidot, Neovia Innovation
	Contributors:	
	Reviewed by:	Sacha BrunCEA; Maïke Gilliot, ETP4HPC.
	Approved by:	MB/TB

Document Status Sheet

Version	Date	Status	Comments
0.1	4/5/2020	Draft	Full document draft
0.5	11/05/2020	Draft	Draft for internal review
1.0	19/05/2020	Version for Technical Board	Integration of reviewers' comments
1.1	26/05/2020	Final version	Editing changes

Document Keywords

Keywords:	PRACE, , Research Infrastructure, High Performance Computing, Technology project, FET-HPC projects, European Processor Initiative, Strategic Research Agenda

Table of Contents

List of Figures

List of Tables

References and Applicable Documents

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

[2] http://www.prace-project.eu

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

List of Acronyms and Abbreviations

Below is a list of acronyms used within the EXDCI-2 project. The acronym specific to this report are explained inside the text where they are used.

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
EPI	European Processor Initiative
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
IESP	International Exascale Software Project
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

Executive Summary

The objective of this work has been to assess the results or on-going work of the first HPC technology projects initiated within the Horizon Europe programme, namely within calls FETH-HPC-2014, 2016 and 2017. In a way, this is an assessment of the effectiveness of that part of the Work Programme, as we have also carried out a 'gap analysis' to compare the project outcomes with the recommendations of the European HPC Technology roadmap, the Strategic Research Agenda[1] of ETP4HPC. We have also analysed the projects' impact on the application domain and their potential opportunities to collaborate with the European Processor Initiative project.

The results of the projects in question have been compiled (which has led to an impressive database of project results) and analysed (with an objective to determine how their results can be re-used and what mechanisms are needed to achieve that). The SRA 'gap analysis' has given rise to recommendations in relations to the continuation of the Work Programme.

We have established that the FET-HPC projects have delivered a large set of IPs, which represents an important asset for the entire European HPC community. A comparison of these results with the SRA provided a tool to structure the future calls. The impact of the work of the project extends beyond the domain of technology and impacts applications and thus the work of the Centres of Excellence. Finally, the results developed by the FET-HPC projects can serve as the basis of a rich software ecosystem that could help the EPI[2] project penetrate the market.

While the coverage of the topics included in the SRA is good, some gaps exist, the most worrying being: Integration of security in HPC systems, Support of virtualisation to open HPC system usage, and Emergence of performance metrics for the new applications. Also, most of the milestones of the SRA are too ambitious to be achieved within a single FET-HPC project.

We recommend the following:

- Continue assessing the results of the HPC technology projects to build a global vision and to use it to improve the HPC strategy;
- Set up calls with higher TRL objectives for the technologies to enter the virtuous circle of continuous improvements by their user bases; These calls could be:
    - Integration projects with the objective to deliver a complete HPC solution;
    - Horizontal projects with the objective to develop a layer that could be used by several European computing centres or application developer communities
- Develop a programmatic approach of the research programme with a strong focus on the strategic axes as the EPI;
- Selects projects based on their contribution to the overall objectives; and implement a framework to facilitate strong cooperation between the projects selected.

[1] https://www.etp4hpc.eu/sra-2017.html
[2] European Processor Initiative https://www.european-processor-initiative.eu/

1 Introduction

The objective of task 2.4 of the EXDCI-2 project is to coordinate research actions targeting HPC technologies in Europe. In order to achieve this objective a global vision of the HPC technology research going on in Europe has been built. Then, this global vision has then been used to work on:

- Increasing the impact of the projects by creating synergies between the projects and by facilitating the reuse of project results;
- Understanding the position of the European ecosystem in the international HPC landscape;
- Providing recommendations for shaping the future R&D HPC technology programme.

During the first phase, the task focused on an analysis of the FET HPC projects issued from the call FETHPC-1-2014 which had started in 2015 and ended between September 2018 and May 2019. A dialogue has been established with the 19 projects originating from this call. Then the same kind of dialogue has been established with the second and third sets of FET-HPC projects funded through calls issued in 2016 and 2017. The outcomes of this activity are presented in Section 2 of this report. We provide an analysis of each FET-HPC project including either their main results and exploitation paths if the project has been completed or the expected outcomes if the project is still ongoing.

In Section 3, we present a gap analysis between the SRA-3[3] milestones and the contributions of the FET-HPC projects (i.e. either the results for finished projects or the objectives in the case of the ongoing projects). This exercise serves to assess the coverage of the current HPC technology research effort compared to the whole field and to determine how the current methodology, in combination with ETP4HPC Strategic Research Agenda and the related EC calls for proposals, performs as a well-structured research effort.

During the analysis of the FET-HPC projects, it appeared that beyond their contributions to develop European HPC technologies, these projects also work on improving the state of the art of some applications. Most of them deploy a co-design approach or at least test their technology developments against some use cases and thus provide improved versions of some of the HPC applications. The contributions of FET-HPC projects to applications has been analysed to find potential synergies between the projects working on the same applications. It also outlines potential cooperation between the technology pillar and the application pillar[4] of the European HPC strategy. This analysis is presented in Section 4.

During EXDCI-2, some activities initially planned on the Extreme Scale Demonstrators[5] have been refocused to deal with the exploration of cooperation between the FET-HPC project and the European Processor Initiative. The findings of this analysis are presented in Section 5.

In Section 6, we present some recommendations for the future HPC technology research programme. These recommendations deal mainly with the organisation of a research programme which should be more structured and result in more direct impact on the upcoming European exascale and post-exascale HPC infrastructures.

[3] The third Strategic Research Agenda issued by ETP4HPC see https://www.etp4hpc.eu/sra-2017.html
[4] Represented by the Centres of Excellence initiative
[5] Extreme Scale Demonstrator was a concept envisioned in 2018. An ESD would have been a large integration project with the objective to develop a complete solution including a hardware platform and a software stack

2 FET-HPC projects

The support of research on HPC technologies is organised within three calls from the Future and Emerging Technologies programme of the Horizon 2020 Excellence in Science pillar.

Call reference	Call title		Selected projects	Status of the projects	Budget
FETHPC-1-2014	HPC Core Technologies, Programming Environments and Algorithms for Extreme Parallelism and Extreme Data Applications	RIA	19	Finished	94 M€
FETHPC-01-2016	Co-design of HPC systems and applications	RIA	2	Running	35 M€
FETHPC-02-2017	Transition to Exascale Computing	RIA	11	Running	35,5 M€

Table 1: List of FET-HPC calls

In this report, we will refer to these three sets by respectively using the terms FET-HPC-2014, FET-HPC-2016 and FET-HPC-2017[6].

Within EXDCI-2, an analysis has been conducted to assess the technology content of these projects with the objectives to:

- Increase the reuse of the results of the projects,
- Facilitate the dissemination of project results,
- Detect and foster synergies between the projects.

As the FET-HPC-2014 projects are finished, it has been possible to carry out a survey[7] of their activities and generate a database[8] of their results. This approach has allowed us to identify the potential "users" of the results and to promote the relevant technologies through adapted channels. "Users" in this context refers to people, organisations or projects that can benefit from the result. The main findings are presented below, together with first a global view and then an analysis project by project.

In the case of the FET-HPC-2016-2017 projects, we have focused on their objective rather than on their results as they are not completed yet. Nevertheless, we have a good knowledge of what could be achieved by these research actions, especially in the case of the two large FET-HPC-2016 projects targeting co-design for exascale systems. An analysis of these different projects is provided in this Section.

The reader interested in the FET-HPC projects can also look at the handbooks that have been issued on a regular basis by ETP4HPC with the support of EXDCI-2. These handbooks can be access at https://www.etp4hpc.eu/european-hpc-handbook.html.

6 The complete list of the projects is given in Annex 1
7 The questionnaire used for the survey is presented in Annex 2
8 https://exdci.eu/activities/fethpc-results

2.1 FET-HPC-2014 projects

2.1.1 *Global perspective*

Most of these European FET-HPC projects started in September 2015 ended from September 2018 to mid-2019. The survey conducted by EXDCI-2 has used the questionnaire presented in Annex 2, which was completed during the second half of 2018. A real dialogue was established between the project's teams and the EXDCI-2 team. All projects have been very supportive and have participated in this work.

To better understand the analysis of the results produced by the projects, it is important to acknowledge the broad spectrum of technical domains addressed by these research actions. The following classification illustrates the range of the topics tackled:

HPC system focused projects		
	From package to system	ExaNoDe, EXANEST, ECOSCALE
	ARM based HPC	MONT BLANC
	Reconfigurable systems	EXTRA, MANGO, Green Flash
	IO	SAGE, nextgenio
HPC stack and application oriented projects		
	Energy efficiency	ANTAREX, READEX
	Programming model	INTERTWINE
	Multiscale	COMPAT, AllScale
	Generic applications: Hyperbolic PDE, Machine learning, Fluid dynamics, Numerical linear algebra, Weather models	EPiGRAM, EXCAPE, ExaFLOW, NLAFET, ESCAPE

Table 2: Typology of FET-HPC-2014 projects

All projects have taken part in this dedicated survey out of which a list of the most relevant project outputs (i.e. 'IPs': intellectual property elements) has been produced. A simple quantitative analysis shows that most of the results are in the field of software. Out of the 171 IPs listed, 114 (two third) are software and 20 are hardware related. The other types of results are APIs, applications optimisations, benchmark suites, trainings and demonstrators.

Interestingly, most of the IPs produced could be exploited. They address the needs of different types of users or stakeholders that can be represented in the following table:

	API	application optimisation	benchmark suite	demonstrator	hardware	report	software	training	Total
application developer	7	1		8			53	5	74
computing centre		1	4		1		6		12
end user		5				2	34	2	43
HPC system provider			2		15		15		32
integration project					4		5		9
processor provider							1		1
Total	**7**	**7**	**6**	**8**	**20**	**2**	**114**	**7**	**171**

Table 3: Users of the FET-HPC-2014 project results

Clearly, the FET HPC 2014 projects have generated a lot of IPs that can be useful for end users and application developers. Some of them can also be reused by HPC system providers, computing centres or could provide valuable technologies to be applied in integration projects (such as ESDs[9], Extreme Scale Demonstrators).

Qualitatively, in the hardware area, one should note the development of several processor or FPGA boards, active interposer technology, interconnect technologies (one using photonics) and cooling technology.

The system-oriented projects have developed 8 demonstrators, most of them open to experiment by external teams. The larger ones in terms of computing power come from ExaNest/EcoScale, Montblanc and Mango. The IO-related projects, Sage and NextGenIO have also produced demonstrators that can be used for testing new storage hierarchy or object file systems.

Some APIs have been proposed by the projects in domains such as FPGA management, object file system, energy efficiency and interaction between runtimes.

In the software area, besides the enhancement of several applications or application kernels, there are results in domains such as FPGA programming, file systems, runtime, energy efficiency, time constrained computing, tuning/debugging tools, etc.

The complete set of results is broad and diverse. The exploitation of this basis could be enabled by the new FET-HPC-2017 or other projects as they will continue some of the work (e.g. EuroExa, Montblanc2020, Sage2, Escape 2 or Recipe). In addition to these projects, it would be good to launch integration projects that could use and further develop some of the IPs generated. This vertical integration was one of the objectives of the ESD proposed by the HPC ecosystem. Based on the current results, we also suggest horizontal projects that will integrate some of the results that are closely related and complementary in one common framework. At least in domains such as FPGA programing, runtime and energy efficiency, it would be valuable to have open environments for the European HPC ecosystem.

To summarise, the FET-HPC-2014 call has produced an impressive set of IPs that could be reinforced by vertical or horizontal integration projects, thus pushing these new technologies toward industrialisation.

9 ESD stands for Extreme Scale Demonstrator - a concept developed in 2018. An ESD would have been a large integration project with an objective to develop a complete solution including a hardware platform and a software stack. This kind of project has been once included in the HPC work programme but it has been removed and replaced in the EuroHPC work programme by projects linked to the European Processor Initiative.

Techno Targeted user	Computing node/board	Interconnect Memory hierarchy	Storage/file system	Tools for FPGA	Software stack	Programming model/ tool	Optimization tools	Library	Application
Integration projects (ESD)	Ecoscale Exanode MontBlanc3	NextGenIO ExaNest Ecoscale	SAGE	MANGO EXTRA	Greenflash EcoScale Antarex	Exanode InterTwine	Ecoscale		
HPC system provider	Greenflash ExaNest Ecoscale MontBlanc3 Exanode	NextGenIO ExaNest	SAGE	MANGO Ecoscale	MANGO Greenflash MontBlanc3 Readex	AllScale InterTwine			NextGenIO COMPAT Antarex
Computing centre			SAGE		MANGO NextGenIO MontBlanc3 COMPAT Readex Antarex	InterTwine			NextGenIO COMPAT Antarex
Application developer			ExaNest SAGE NextGenIO	MANGO Ecoscale EXTRA	Readex Antarex	MANGO AllScale Greenflash MontBlanc3 Exanode InterTwine Antarex	Greenflash MontBlanc3 EXTRA Readex Antarex	ExaFlow ExCAPE NLAFET Readex Antarex	
End user						ExaFlow	NLAFET	NLAFET	ExaNext ExaFlow ESCAPE ExHype ExCAPE NLAFET Readex Antarex

2.1.2 *Dissemination and Project results data base*

A dissemination effort has also been implemented to share the analysis of the FET-HPC-2014 projects with the HPC community:

- A presentation at the Brühl PRACE-CoEs-FET HPC-EXDCI Workshop in October 2018 (see https://events.prace-ri.eu/event/750/overview);
- A presentation during the European HPC ecosystem BoF at SC2018 (see https://sc18.supercomputing.org/presentation/?id=bof104&sess=sess367);
- An article in ETP4HPC Annual Report 2018 (see https://www.etp4hpc.eu/pujades/files/ETP4HPC_Annueal-Report-2018_web.pdf);
- A presentation during the ETP4HPC General Assembly ;
- Some information was shared during the panel of the mini-symposium organised by the EC on co-design together with the European Processor Initiative at EHPCSW 2019 in Poznan
- A presentation at Poznan during the HPC Ecosystem Workshop Plenary Session on Thursday May 16th 2019.

Moreover a data base of all Project results has been built under the EXDCI website (https://exdci.eu/). This web site is the main tool to disseminate information about the project and communicating its activities and collaborations as much as possible.

In that respect, a sub-page titled "FETHPC project's results" was created in order to present the rich set of results that were produced by the FETHPC 2014 projects as they are important assets for the European HPC ecosystem. These outputs have been listed and organised into a database, easily searchable on the page (See Figure 1 opposite). This database allows the community to search through the results according to projects, domain of results (API, demonstrator, report software, training, etc.) or targeted users (developers, providers, computing centres, projects, end-users).

Figure 1 - Design of the FET HPC projects results page

2.1.3 *Analysis by project*

The following includes a short description of the main results and a point of view on the exploitation path of some of the results of each project. This analysis has been carried out by the authors and should not been viewed as the position expressed by the projects themselves.

Montblanc3

Main results

The most noticeable IPs of the project are:

- A prototype based on the Cavium (now Marwel) ThunderX2 processor
- A software stack to support HPC on ARM platform
- Some performance tools adapted to an ARM platform
- Some application optimizations for ARM platform
- A training for ARM based platform

Exploitation path

Most of the results have been industrialised by Bull/Atos which has an offering including the ThunderX2 processor of its Sequana HPC platform and a complete dedicated software stack. Some of the result are also the basis of further research projects as Montblanc 2020 and EPI.

ExaNode

Main results

The most noticeable IPs of the project are:

- Technology for an active interposer
- Integration of different chiplets in a package with the interposer technology

Exploitation path

The main exploitation path within the EuroHPC HPC strategy is to use the results of this project within the EPI project. This is achievable since the CEA, the main IP owner, is a partner of the EPI project.

ExaNeSt

Main results

The most noticeable IPs of the project are:

- FPGA boards
- Interconnect technology
- Software stack for the UNIMEM architecture
- Demonstrator
- Applications optimization for the ExaNeSt platform

Exploitation path

Most of the results of the ExaNeSt project have provided the basis of the EuroEXA co-design project (see next section). The natural exploitation path is industrialisation of this UNIMEM architecture and of the software stack to offer a very high-end HPC system mainly based on FPGA.

ECOSCALE

Main results

The most noticeable IPs of the project are:

- Software stack to program a large system of FPGAs
- CAD tools to use FPGAs system
- API to reconfigure at runtime FPGAs system

Exploitation path

Again, one of the main users of the results of this project is the EuroEXA project. The exploitation is linked to the exploitation of EuroEXA results.

EXTRA

Main results

The most noticeable IPs of the project are:

- Platform to develop the implementation of applications on FPGAS (part is the CAOS environment)
- Debugging tools
- Performance tools

Exploitation path

The natural exploitation of EXTRA would be the development of a FPGA programming environment. This could be the purpose of a horizontal project that could use other FET-HPC-2014 results to provide to the European application developers community a rich and stable programming environment for FPGA. This kind of project can be valuable in the context of EPI as this project plans to integrate FPGA in some of the chips that they are working on.

MANGO

Main results

The most noticeable IPs of the project are:

- FPGAs boards
- Many core architecture on FPGAs
- Interconnect implemented in FPGAs
- Programming environment for FPGAs (including the BarbequeRTRM software)

Exploitation path

The results of the MANGO projects are exploited by the FET-HPC-2017 project RECIPE. In the case of EXTRA, a natural exploitation path would be to set up a horizontal project providing the European application developers community with a rich and stable programming environment for FPGA.

Greenflash

Main results

This project is different than other projects because it was managed by a research team from the astrophysics domain. It has produced valuable IPs for the HPC community:

- FPGAs boards
- Mechanisms to organise data transfers with time constraints
- Mechanism to run kernel on GPU with time constraints

Exploitation path

The results of Greenflash could be integrated with the ones of EXTRA and MANGO in a horizontal project on FPGA programming environments. Some elements could also be used by projects working on HPC in the loop, urgent HPC or real time HPC.

SAGE

Main results

The most noticeable IPs of the project are:

- Object file system called MERO
- API for this object file system called CLOVIS
- HSM tools for computing centres
- Benchmark suite well suited for new IO oriented applications

Exploitation path

The SAGE project has a continuation in the FET-HPC-2017 call with the SAGE2 project. The exploitation path is the industrialisation of the MERO and CLOVIS IPs and the development of an application developer community for this software.

NEXTGenIO

Main results

The most noticeable IPs of the project are:

- A software stack to take benefit of the new memory hierarchy introduced by NVRAM
- A slurm adaptation to benefit for NVRAM
- A demonstrator

Exploitation path

The natural exploitation path is to integrate NEXTGenIO results in a software stack for HPC systems using NVRAM. The market access could be through an HPC system vendor.

ANTHAREX

Main results

The most noticeable IPs of the project are:

- DSL (Domain Specific Language) for code analysis, transformation and exploration
- Source-to-source compiler to enable the use of the DSL
- Autotuning capability implemented with runtime to optimize energy consumption

- Application implementations taking advantage of the autotuning capability developed by the project

Exploitation path

The most natural exploitation path for ANTHAREX results would be a horizontal project to industrialise an environment dealing with the energy efficiency optimisation of applications.

READEX

Main results

The most noticeable IPs of the project are:

- Extension of Score-P to measure energy efficiency
- Software stack to analyse and optimise energy efficiency of applications
- API for profiling libraries

Exploitation path

As for ANTHAREX, the most natural exploitation path for READEX results would be a horizontal project to industrialise an environment dealing with the energy efficiency optimisation of applications. The two projects are very complementary and the integration of the two approaches could be very interesting.

AllScale

Main results

The most noticeable IPs of the project are:

- Failure detector and checkpoint restart functionality
- Monitoring framework with online performance introspection capabilities
- Application implementations over the resilient software stack

Exploitation path

The exploitation of AllScale results would require an additional industrialisation effort to target a resilient HPC software stack. This could be achieved through a horizontal high-TRL software project.

COMPAT

Main results

The most noticeable IPs of the project are:

- Execution environment over multiple HPC systems
- A simulator of larges computing centres

Exploitation path

The COMPAT result could be used in the context of very large and complex workflows. Some of the project concepts could be reused in order to deploy applications ranging from edge to HPC centres.

InterTwine

Main results

The most noticeable IPs of the project are:

- Task aware MPI[10] library
- Task aware GASPI[11] library
- Runtime managing different programming model (based on StarPU[12])

Exploitation path

The project results are reused by the EPiGRAM-HS project (FET-HPC-2017). The natural exploitation path would be to industrialise an application development framework for heterogeneous architectures using InterTwine results and other projects results (mainly ASPIDE, EXA2PRO, EPiGRAM-HS, EPEEC).

ESCAPE

Main results

The most noticeable IPs of the project are:

- Software representing weather forecast model components
- Benchmark suite
- Different implementation of weather dwarfs[13]

Exploitation path

The ESCAPE project results have been consolidated by a project continuation ESCAPE-2 in the FET-HPC-2017 call and the Centre of Excellence EsiWACE. The availability through this effort of an exascale ready weather and climate model is the best exploitation path for this project.

ExCAPE

Main results

The most noticeable IPs of the project are:

- Machine learning and various types of Bayesian Matrix Factorisation applications
- Workflow execution framework
- Framework to implement various type of Neural Networks

Exploitation path

The natural exploitation plan of Escape would be the provision of tuned libraries on European exascale platforms. Some interaction with EPI could be valuable in order to develop a library ready for EPI systems.

ExaHYPE

Main results

[10] Message Passing Interface : one of the main programming models in HPC
[11] Global Address Space Programming Interface see http://www.gaspi.de/gaspi/
[12] A Unified Runtime System for Heterogeneous Multicore Architectures http://starpu.gforge.inria.fr/
[13] The dwarf notion was introduced by Phillip Colella in his 2004 presentation “Defining Software Requirements for Scientific Computing”. He gave his list of the now-famous “Seven Dwarfs” of algorithms for high-end simulation in the physical sciences that was later on complemented by other contributions.

The most noticeable IPs of the project are:

- An engine to solve hyperbolic systems of partial differential equations using special high-order discontinuous Galerkin schemes
- Applications using the library

Exploitation path

The exploitation path could be the provision of tuned libraries on European exascale platforms. Also, some interaction with EPI could be valuable in order to develop a library ready for EPI systems.

NLAFET

Main results

The most noticeable IP of the project are:

- Parallel numerical linear algebra software for the solution of dense and sparse linear systems of equations and eigenvalue problems,
- Different application optimized by using the above library

Exploitation path

Similarly to ExaHYPE, NLAFET has developed an optimised library. The same recommendations for the exploitation can be made.

ExaFLOW

Main results

The most noticeable IPs of the project are:

- Optimised version of Nek5000, an open-source code base on the spectral element method
- Different application optimized by using the above code

Exploitation path

The same case as ExaHYPE and NLAFET. Similarly, some interaction with EPI could be recommended to increase the application spectrum of HPC system based on EPI.

2.2 FET-HPC-2016 projects

In 2016, the European Commission issued a call for large projects aiming at "innovative and ground-breaking approaches to system architectures targeting extreme scale, power-efficient and highly resilient platforms with emphasis on balanced compute and data access characteristics".

Two projects were selected and started in the third quarter of 2017, scheduled to finish in the first quarter of 2021. EXDCI-2 started a dialogue with these projects in 2019 and continue to follow their activities with an objective to facilitate the dissemination of their results.

2.2.1 *DEEP-EST*

The DEEP-EST project is a continuation of two FP7 projects: DEEP and DEEPER. All these projects have been by FZJ Forschungszentrum Jülich. DEEP-EST started in July 2017 and it is expected to end in March 2021.

As explained on Cordis, the project focuses on the Modular Supercomputer Architecture (MSA) and demonstrate its benefits. The MSA integrates compute modules with different performance characteristics into a single heterogeneous system. Each module is a parallel, clustered system of potentially large size. A federated network connects the module-specific interconnects. MSA brings substantial benefits to heterogeneous applications/workflows: each part can be run on an exactly matching system, improving time to solution and energy use. It can be used in supercomputing centres running heterogeneous application mixes (higher throughput and energy efficiency). It also offers valuable flexibility to compute providers, allowing the set of modules and their respective size to be tailored to actual usage.

The DEEP-EST prototype plans to include three modules: general purpose Cluster Module and Extreme Scale Booster supporting the full range of HPC applications, and Data Analytics Module specifically designed for high-performance data analytics (HPDA) workloads. Proven programming models and APIs from HPC (combining MPI and OmpSs) and HPDA will be extended and combined with a significantly enhanced resource management and scheduling system to enable straightforward use of the new architecture and achieve highest system utilisation and performance. The DEEP-EST prototype is defined in close co-design between applications, system software and system component architects. Six ambitious and highly relevant European applications from the HPC and HPDA domains drive the co-design, in order to evaluate the DEEP EST prototype and demonstrate the benefits of its innovative Modular Supercomputer Architecture.

During the dialogue with the project, we were able to identify the main expected results of the DEEP-ESP project. They are presented in the table below.

Short description of the result	**Type of result**
Modular Supercomputer Architecture (MSA) that support heterogeneous computing resource and the use of accelerators	architecture
Co-design methodology with questionnaire for application developers and hardware designers, process for interaction	methodology
Interconnect with new approach to dis-aggregate the network from the servers and connect each server via PCIe cables to the network called Fabri[3]	hardware
Network attached memory (NAM) to accelerate application execution	hardware
global collective engine (GCE): networkt-attached device to speed-up collective operations	hardware

Gateway to federate networks and scale up the architecture	hardware
Software environment to distribute the application on the MSA	software
ParaStation MPI and OmpSs implementation to efficently support the use of GPU and GPU direct communication	software
Resiliency support by software stack	software
Slurm enhancement to efficently support the MSA	software
Complete software stack (ParaStation Modulo) to efficently support the MSA	software
Application optimization on the MSA for the codes: Gromacs, NEST (neuro community), High energy physics CERN, SKA correlation, Machine Learning, space weather Leuven	application optimization
Heterogeneous demonstrator with Xeon processors, FPGAs and NVIDIA GPU	demonstrator

Table 4: DEEP-EST expected results

Exploitation paths of some of the results are already in place or they have been determined by the project. For example, the co-design methodology has been reused in the EPI project where FZJ is also very active. Most of the software stack elements are planned to be industrialised and put in production on the HPC systems run by the Jülich Computing Centre. The demonstrator will be used by the project consortium but is also open to external organisations which have interesting applications to be tested on this heterogeneous platform.

Globally, the project has a lot of connections with other projects through its consortium and is able to aggregate part of the European HPC technology research effort.

2.2.2 *EuroEXA*

The EuroEXA project is related to the group of three FET-HPC-2014 projects: ExaNode, ExaNeSt and EcoScale. It started in September 2017 and it is expected to end in February 2021.

As described in Cordis, the EuroEXA project aims to achieve the demands of extreme scale and the delivery of exascale capable architecture. EuroEXA works on the co-design of a ground-breaking platform capable of scaling peak performance to 400 PFLOP in a peak system power envelope of 30MW with the use of innovative technologies coming from SMEs (Maxeler for FPGA data-flow; Iceotope for infrastructure; Zeropoint Technologies for memory bottleneck).

The objective of the project is to co-design a balanced architecture for both compute- and data-intensive applications using a cost-efficient, modular-integration approach enabled by novel inter-die links and the tape-out of a resulting EuroEXA processing unit with integration of FPGA for data-flow acceleration. The developments include a homogenised software platform offering heterogeneous acceleration with scalable shared memory access and create a unique hybrid geographically-addressed, switching and topology interconnect within the rack while enabling the adoption of low-cost Ethernet switches offering low-Latency and high-switching bandwidth.

The co-design approach uses HPC applications from across climate/weather, physics/energy and life-science/bioinformatics domains. The results of the project will be integrated through the deployment of an operational peta-flop level prototype hosted at STFC. Supported by run-to-completion platform-wide resilience mechanisms, components will manage local failures, while communicating with higher levels of the stack. Monitored and controlled by advanced runtime capabilities, EuroEXA plans to demonstrate its co-design solution supporting both existing pre-exascale and project-developed exascale applications.

During the dialogue with the project, we were able to identify the main expected results of the EuroEXA project. They are presented in the table below.

Short description of the result	**Type of result**
Euroexa architecture with ARM CPU and FPGA accelerator linked by an interconnect supporting the UNIMEM architecture	architecture
Demonstrator testbed 2 with more than 200 nodes (each with ARM A73 and FPGA resources)	demonstrator
Demonstrator testbed 3 compatible with EPI technologies	demonstrator
Test chip provided by Manchester U that implements ARM ISA and is integrated into testbed 3 ; support memory compression technology	hardware
Computer board designed for the project CRDB including interconnect and processor and FPGA accelerator	hardware
Blade designed with open compute standard for 16 daughter boards	hardware
Rack with top level interconnect (infiniband)	hardware
Cooling technology for blade and rack up to 200W	hardware

Container technology that allow up to 2MW of computing power per rack and can scale by interconnecting container in a compact way up to exascale	hardware
Interconnect at node level integrated in a FPGA	hardware
Interconnect switch at intermediate level integrated in a FPGA	hardware
Interconnect switch at blade level integrated in a FPGA ; interface between inter-rack Infiniband network and internal protocol	hardware
System level software to enable UNIMEM architecture	software
MPI communication library supporting the architecture	software
Gaspi communication library supporting the architecture	software
Slurm adaptation to support FPGA architecture	software
Programming environment including support for FPGA	software
Application optimization on the Euroexa architecture and demonstrators : Quantum Espresso, Nemo, NEST/DPSNN, image classification (astronomy), FRTM, InfOli, SMURFF, AVU-GSR, IFS, LBM, Alya, GADGET, LFRic	application optimisation
Methodology to port applications on FPGA	methodology
Methodology to assess performance acceleration of FPGA	methodology

Table 5: EuroEXA expected results

The project has consolidated multiple results coming from the ExaNeSt and EcosScale projects. The demonstrator that will be put in place will offer the largest FPGA platform available to the European researchers. Even if some of the ambitions will be difficult to achieve due to industrial problems (mainly related to integration of ARM core IP and foundry of chips), EuroEXA's experience with FPGA systems will be very valuable. Some exploitation paths in this direction should be consolidated.

Globally, EuroEXA has seen the emergence of an industrial solution ported by SMEs and it is Europe's leader in the domains of FPGA based HPC systems.

2.3 FET-HPC-2017 projects

2.3.1 *Global vision*

September 2017 was the deadline for the submission of the last set of FET-HPC projects. The call was focused on transition to exascale with 5 sub-topics:

- High productivity programming environments for exascale
- Exascale system software and management
- Exascale I/O and storage in the presence of multiple tiers of data storage
- Supercomputing for Extreme Data and emerging HPC use modes
- Mathematics and algorithms for extreme scale HPC systems and applications working with extreme data

Eleven projects were selected and they can be grouped in the following categories:

Programming environment	
	ASPIDE, EPEEC, EPiGRAMHS, exa2pro
Heterogeneous system management	
	RECIPE
IO and storage	
	Sage2, MAESTRO DATA ORCHESTRATION
Visualization, interactive HPC, urgent computing	
	VESTEC
Verification, Validation, Uncertainty Quantification	
	ExaQUte, ECMA
Algorithm: weather and climate models	
	ESCAPE

Table 6: Typology of FET-HPC-2017

Almost all projects started in September 2018 and are expected to finish during the second half of 2021 (with the exception of ASPIDE which plans to finish end of 2020). They are rather small in terms of budget (from 2.5 M€ to 4M€).

Due to the timing of the projects, the EXDCI-2 analysis has been more focused on the objectives of the projects rather than their results. Nevertheless, this analysis can lead to interesting findings in terms of project synergies and potential cooperation.

2.3.2 *Analysis by group of projects*

Programming environment

The four projects ASPIDE, EPEEC, EPIGRAM-HS and EXA2PRO address the problem of increasing the efficiency for programming the current and future HPC systems. ASPIDE has is centred on data intensive applications and works with different technologies.

Each project has a strong consortium with mainly academic partners. There are only two partners that are present in more than one project: Fraunhofer (EPEEC and EPIGRAM-HS) and INRIA (EPEEC and EXA2PRO).

All projects have adopted a co-design approach but they work on different applications (see section 4). Even the fields of the applications are very diverse and only computational fluid dynamics and image processing are covered by two projects.

The following table summarises the main technical contributions and approaches of the four projects:

Topic	ASPIDE	EPEEC	EPIGRAM-HS	EXA2PRO
DSL	Data centric and eHealth New developments		Deep Learning	
Language	DSP + existing MPI+X	C, Fortran, PGAS, MPI	C, Fortran, MPI	C, Fortran, MPI
Internal model	Mapreduce Data centric Graph of tasks	OpenACC, GASPI, ArgoDSM Use of Mercurium	GASPI, OpenCL	Composition of skeletons Use of Mercurium
Run time	New development	OmpSs	Based on GPI implementation	StarPU
Data optimization	New developments	New layer for heterogeneous memory management	New memory abstraction	Optimisation at composition level
File system	New development with underlying Lustre, GPFS			
Autotuning	Yes	Yes	Yes at communication level	Yes
Targeted architecture	CPU, GPU	CPU,GPU, FPGA	CPU,GPU,FPGA	CPU,GPU,FPGA

Table 7: Programming environment projects

There are several technical commonalities in the projects EPEEC, EPIGRAM-HS and EXA2PRO. They all plan to address the use of FPGA. Nevertheless, the level of integration seems different, with a finer grain approach in EXA2PRO.

All projects build on existing pieces of software that have been developed in previous European projects, the more important ones being: GASPI (and its implementation GPI), OmpSs, StarPU, SkePU and Mercurium.

All projects plan to introduce autotuning features for the system to optimise the execution of an application.

In conclusion, even if the approaches are different, there are obvious synergies between these four projects. They can together:

- Propose a common vision of the challenges for efficient programming of heterogeneous HPC systems;
- Provide information on the different approaches and guidelines to choose between them depending on the application features
- Develop their relationship with the application developer communities.

Indeed, the projects share a big challenge which is to build a critical mass of users in these application developer communities. Even if there is no silver bullet (i.e. a suit-all solution), some hints could be:

- Clearly defined the API that can be pushed
- Support of an industrial company
- Continuity in the funding of a potential de facto standard solution
- Programmatic research programme as the US Exascale Computing Project (ECP)

Some of the challenges clearly exceed the scope of an individual project and we will come back to this discussion in the conclusion of this Section and in the recommendations of this report.

For the time being, the four projects could strengthen their links with the CoEs. Most of the CoEs work on large applications that would be interesting targets for the methodology of these projects. Most of the applications would benefit from heterogeneous HPC systems and we see a win-win situation for the projects and the CoEs.

Interaction with the EPI could also been valuable. As the EPI has the objective to provide chips and systems with heterogeneous resources (CPU, GPU like accelerator, FPGA, specific accelerators) the programming environment developed by the four projects could be relevant. Within the EPI, the heterogeneous resources can be tightly coupled on the same chip whereas the projects work more on loosely coupled resources at the level of the board. Nevertheless, interactions could drive the EPI to a more complete software environment and the projects could avail of new opportunities such as exploitation and new research problems.

Heterogeneous system management

The RECIPE (REliable power and time-ConstraInts-aware Predictive management of heterogeneous Exascale systems) project is the only one related to the second topic of the call. The main objectives of the project are to provide:

- a hierarchical runtime resource management infrastructure optimising energy efficiency and ensuring reliability for both time-critical and throughput-oriented computation;
- a predictive reliability methodology to support the enforcing of QoS guarantees in face of both transient and long-term hardware failures, including thermal, timing and reliability models;

- a set of integration layers allowing the resource manager to interact with both the application and the underlying deeply heterogeneous architecture, addressing them in a disaggregate way.

This project has some common partners with the FET-HPC-2014 MANGO project and reuses some of the results even if the ambition is different.

The contributions of RECIPE will include a new version of BarbequeRTRM[14] and some enhanced features in Slurm[15]. The platform will be able to manage heterogeneous resources including CPU, GPU and FPGA. It will also manage NVRAM resource and of course the interconnect (based on Infiniband).

The exploitation of RECIPE results will require the adoption of the developed framework by computing centres or by HPC system providers. This is a challenging condition but that can be achieved with a relevant plan for industrialisation and some aggregation of resources and efforts.

IO and storage

SAGE2 and MAESTRO are working in a complementary way to develop a new generation of object storage middleware. The premise of both projects is that the approximation to storage in future computing facilities needs to be made in a holistic way for data access.

SAGE2, a continuation of SAGE, concentrates on the further development of Mero as object storage underlying system, and the Clovis API in order to facilitate application integration. MAESTRO will exploit those capabilities to provide the glue between the system software data I/O features developed by SAGE2, and the application level in terms of data placement optimizations at runtime. A number of use cases will be integrated in order to test the capabilities in the experimental installation in Jülich.

There is a general agreement in the importance of developing system software and supporting runtime middleware able to treat all the memory layers in a holistic way. This is particularly so in order to fully exploit the capabilities of new memory technologies such as NVRAM.

Extreme I/O and data placement in Exascale systems will become a source of energy consumption which needs to be optimised. In this context, Mero is being tested at this moment using an installation in Jülich.

Some of the developments of the MAESTRO project can be applied to generic object storage systems such as CEPH[16]. However, the components that tackle the computing-to-data features, which would be of interest, e.g. if implemented in schedulers, are implementable only in Mero, as CEPH does not tackle the problems of placing data close to the computing. The developments of Mero at the scheduler level will be implemented and tested in SLURM.

The adoption of both projects results will be facilitated by the demonstration on applications of the benefit of the approach. The co-design aspect of the projects is also of high importance. The projects face a problem similar to the programming environment ones in building a community of users. This is very challenging with the means of only one project and a global, broader approach could help. We will come back to this question in the recommendations.

[14] Barbeque Run-Time Resource Manager see https://bosp.deib.polimi.it/doku.php
[15] Slurm is the workload manager used by many computing centres.
16 CEPH is a distributed object, block, and file storage platform see https://github.com/ceph/ceph

Visualization, interactive HPC, urgent computing

VESTEC is an original project that addresses the use of extreme computing in real-time applications with high velocity data and live analytics. VESTEC intends to create the software solutions needed for urgent decision making in various fields as wildfire monitoring and forecasting, analysis of risk associated with mosquito-borne diseases and the effects of space weather on technical supply chains.

VESTEC objectives are to:

- build a flexible toolchain to combine multiple data sources,
- efficiently extract essential features,
- enable flexible scheduling and interactive supercomputing,
- realise 3D visualisation environments for interactive explorations by stakeholders and decision makers.
- develop and evaluate methods and interfaces to integrate high-performance data analytics processes into running simulations and real-time data environments.

The project addresses a field of emerging use of HPC systems. It is very close to the ideas proposed by EXDCI-2 of HPC in the loop or HPC in the transcontinuum. As being the first project to deal with this topic, it has fewer synergies with the other FET-HPC projects. Some interactions with the four projects supported by the call on HPC and Big Data (call ICT-18-2018) could be relevant for VESTEC.

Regarding the exploitation, it is obvious that the results will be highly valuable for the computing centres that would have to offer new software framework for dealing with urgent HPC. Even if VESTEC is only at half-way through its life, it would be interesting to set a communication channel with computing centres to start to build the awareness and see it some knowledge transfer can be organized.

Verification, Validation, Uncertainty Quantification

There are two projects which tackle the topic of Verification, Validation and Uncertainty Quantification: VECMA and ExaQUte. However, their approaches are quite different.

VECMA plans to develop an open source toolkit and has adopted a multiscale approach. It aims to develop:

- a collection of Uncertainty Quantification and sensitivity analysis Primitives (UQPs), tailored to efficiently use current HPC infrastructures, and to incorporate expected requirements for use on exascale architectures. The UQPs will capture, in modular form, specific sub-activities required for the uncertainty quantification and sensitivity analysis in multiscale applications.
- a constrained set of multiscale Verification and Validation Primitives (VVPs) to specifically capture and formalize activities which support the extension of single scale verification and validation procedures to multiscale settings, again tailored to the exascale.

ExaQUte deals with Uncertainty Quantification with a Multi-Level MonteCarlo (MLMC) approach that allows an analysis of a high number of stochastic variables. The domain of application addressed is also limited to the optimisation of structures under wind loads. The project plans:

- New theoretical developments to enable MLMC combination with adaptive mesh refinement, considering both, octree-based and anisotropic mesh adaptation.

- Gradient-based optimisation techniques to consider uncertainties by developing methods to compute stochastic sensitivities, this requires new theoretical and computational developments. With a proper definition of risk measures and constraints, these methods allow high-performance robust designs, also maximising the solution reliability.
- The use of complex geometries to guarantee a high robustness in the mesh generation and adaptation steps, while allowing preserving the exact geometry representation.

Even if the approaches are different, there are some common challenges for both projects, e.g. how to generate several jobs to assess the uncertainty and optimise the run of this multiplicity of jobs. Both projects have started to work together and their synergies are being analysed.

The projects can also join forces to build awareness within the user communities about the technologies developed. Even if the projects have use cases that help them to achieve a good understanding of the user needs and to disseminate in some domains, a broader user base will be beneficial to guarantee a good level of support of the way user plan to integrate VVUQ.

Weather and climate models

The ESCAPE-2 project is a continuation of the FET-HPC-2014 ESCAPE project. It aims to develop world-class, extreme-scale computing capabilities for European operational numerical weather and climate prediction, and to provide the key components for weather and climate domain benchmarks to be deployed on European pre-exascale systems.

The objectives of the projects are to:

- Combine frontier research on mathematics and algorithm development and extreme-scale, high-performance computing applications with novel hardware technology:
- Develop and apply a domain-specific language (DSL) concept for the weather and climate community:
- Establish weather and climate model benchmarks based on world class European prediction models.
- Develop a cross-disciplinary Verification, Validation, Uncertainty Quantification (VVUQ) framework.
- Produce an open-source software framework.

This project is in close contact with the ESiWACE-2 Centre of Excellence and analyses potential synergies on the VVUQ topic with the VECMA and ExaQUte projects.

The exploitation of the project results is dealt with through the ECMWF[17] the coordinator and the European climate and weather organizations that are well aware of the activities of both ESCAPE-2 and ESiWACE-2.

2.4 Main findings of the FET-HPC project analysis

Research in HPC technologies has greatly benefited from the FET-HPC effort, which amounts to over 30 projects and an investment of around 165 M€. The projects have delivered or are in the process of delivering novel, world-class results. These results will have an impact on the entire European HPC value chain, which has been explained in a document prepared on request by EuroHPC (see Annex 3 for the complete document).

17 European Centre for Medium-Range Weather Forecasts

A strategy is needed to make the most out of this effort. A potential solution could be to act like the US, which have set up the large project ECP[18], which encompasses:

- The creation of HPC libraries with Centres of Excellence,
- The creation of an HPC software stack that will be used on the exascale DoE systems coming in the 2021-2023 time frame,
- The connection with the US HPC vendors that will provides these systems.

This is an integrated project with a strong leadership of the DoE and a complete management team to deal with the different activities. The same approach of a very coordinated effort can be seen in Japan, in the form of the Fugaku project which includes not only the development of a hardware systems but also applications and software environments.

If Europe does not want to follow this tightly coordinated path, an alternative could be to put in place three types of instruments which will foster the exploitation of the results of the FET-HPC projects:

- Integration projects with high TRL level. These projects should be structured around the potential HPC system providers and help to push new technologies toward the market; the H2020-JTI-EuroHPC-2020-01 call is heading in this direction.
- Horizontal projects with high TRL level and with an objective to produce software of pre-production quality level. The projects should aggregate different technology pieces to provide software with sufficient coverage and usability level. Among the candidate topics are heterogeneous HPC system programming environment, FPGA programming environment, energy management framework, profiling and tuning tools or data management framework. The projects should result in a software that can be installed by the large European HPC centres and deliver new levels of performance at the level of one layer of the HPC software stack.
- Actions targeting the potential "users" of the technologies. There are at least three kinds of "users" to be considered: application communities, application developer communities and computing centre operational teams. They are the potential users of most of the FET-HPC project results. As it is very difficult at the level of one project to efficiently reach these communities, we suggest that specific actions will be taken to increase the interactions of technology providers (i.e. the FET-HPC projects) and these "users" communities. The actions could take the form of dedicated workshops with discussions combining user challenges and potential solutions analysis. Other activities could include trainings or demonstrations. In short, a push is needed at the European HPC ecosystem level to complement the existing (but limited in resources) projects initiatives.

These three instruments have emerged from the FET-HPC project analysis as means to leverage the ongoing effort. We do not pretend that they might be exclusive of other ideas but we are confident that they can help to achieve a better European HPC technology ecosystem. They are well aligned with the "operational recommendations" presented in the last ETP4HPC Strategic Research Agenda[19] (see Chapter 10 of the SRA).

18Exascale Computing Project see https://www.exascaleproject.org/

19 This document has been supported by EXDCI-2 and is available at https://www.etp4hpc.eu/pujades/files/ETP4HPC_SRA4_2020_web(1).pdf

3 Gap analysis: FET projects versus SRA milestones

3.1 Methodology and analysis

The HPC community has undertaken a considerable effort to write the ETP4HPC Strategic Research Agendas[20], which are recognised as high quality documents defining the European HPC research challenges. The SRAs (SRA1, SRA2 and SRA3) have been used as the basis of the text of the calls issued by the EC. These documents present the driving forces in society, industry, sciences and Information Technologies, the HPC landscape. For each technological domain of HPC, an SRA defines a set of milestones required to facilitate research projects.

How the FET projects are aligned with the SRA milestones? Such an analysis will help understand the current landscape of research actions and answer the following questions, for example:

- What are the key contributions of the FET projects to the SRA roadmap?
- Are there milestones that are not covered by any project?
- Are the milestones a good reference to analyse the project activities?

This analysis will help to better organise the future research effort (e.g. cover the gaps detected) and to improve the methodology currently used with SRAs, calls and project selections.

We decided to use the SRA3 milestones even if at the time of the FET-2014 call only the SRA1 was available and at the time of the FET-2016-2017 calls the reference was the SRA2. There are several reasons for that:

- Being the most recent, the SRA3 milestones are the most relevant in the analysis of the HPC research effort;
- One could map the SRA1 and SRA2 milestones on the SRA3 ones; with a few exceptions;
- We should only use one reference to assess the combine effort of the three calls.

This gap analysis has been based on the milestones defined in SRA3. We have compared this list of milestones with the achievements of the FET-HPC-2014 projects which are finished and the activities of the FET-HPC-2016-2017 projects which are still running.

The results are presented in 2 tables for each of the 8 domains of the SRA:

- HPC system architecture and components
- System software and management
- Programming environment
- Energy and resiliency
- Balance compute, IO ad storage performance
- Big data and HPC usage models
- Mathematics and algorithms for extreme scale HPC systems

The first table represents the contributions of the finished projects clustered by their thematic areas (19 projects) and the second table shows the topics which are addressed by the FET-2016-2017 projects (2 plus 11 projects). We have considered two levels of contribution of a project to a milestone:

- The core activity of the project contributes to the milestone (green)
- The project only has a minor contribution to the milestone (grey), i.e. even provided a full success of the project, only part of the milestone is addressed.

20 https://www.etp4hpc.eu/sra-2017.html

The tables presenting FET-HPC project relationship with the SRA3 milestones are in Annex 4.

3.2 Analysis by SRA domain

HPC system architecture and components

Number of projects addressing the milestone			
	total FET14	FET16-17	total
M-ARCH-1: New standard interfaces available for integrating CPUs and accelerators on nodes and to accommodate innovative unified memory and storage architectures on networks	6	1	7
M-ARCH-2 Having well balanced systems taking benefit from high bandwidth memories and NV memories	2	3	5
M-ARCH-3: Faster end-to-end communication networks (2x and 4x bandwidth in 2018 and 2021 compared to 2015 and lower latency) with energy and power used being proportional to bandwidth	2	2	4
M-ARCH-4: End to end optical communication chain including photonic switching in order to compensate network complexity growth on the larger fabric	1	0	1
M-ARCH-5: Optimised network and storage architectures available with dynamic features, QOS and virtualization capabilities	1	1	2
M-ARCH-6 System and hardware to support performant direct remote memory access that would enable new and easier ways to program parallel applications	2	2	4
M-ARCH-7: Exascale system power envelope in the 5-15 MW power envelop range	1	0	1
M-ARCH-8: Exascale system available, at 100x more performance for relevant applications compared to today's state-of-the-art PRACE Tier-0 systems.	0	0	0

The domain “system architecture and component” is globally covered by the projects. As the FET-HPC-2014 focused on system projects, there is more important contributions from these oldest projects.

The new architectures topics as integration of accelerators and of NVM[21] are tackled by the projects and some progress in these domains will be made.

In terms of interconnect, we have detected the contribution of several project but only a very partial activity targeting optical interconnect which is not strong enough to position the European technology in this domain. The progress on network QoS and virtualisation is also only partial and the projects have not been able to propose complete solutions.

The contribution of the projects on the global performance of the system (M-ARCH-7 and M-ARCH-8) are not significant as the main component here is either the CPU or the GPU, which have not been addressed by the projects. It is clear that such milestones are not very relevant to the projects which do not work on the computing components

21 Non Volatile Memory

System software and management

number of projects addressing the milestone			
	total FET14	FET16-17	total
M-SYS-OS-1 Memory Hierarchy-management policies and libraries for NVRAM	0	2	2
M-SYS-OS-2 OS decomposition and specialized containerisation	1	0	1
M-SYS-OS-3 HW Embedded Security integration and cross layer security support	0	0	0
M-SYS-IC-1 Efficient peer to peer and storage over fabrics support	1	2	3
M-SYS-CL-1 Initial support of mixing HPDA, AI and HPC environment	0	1	1
M-SYS-RM-1 Resource management and orchestration support for complex workflow	2	1	3
M-SYS-RM-2: data aware and Multi-criteria resource allocation integration for adaptive scheduling	2	1	3
M-SYS-RM-3 Dynamic reconfiguration scheduling support (for flexibility and resiliency) purpose	1	4	5
M-SYS-RM-4 Data aware and power efficient scheduling	2	2	4
M-SYS-VIS-1 software support for In-situ computation and visualisation	1	1	2

Again, we see some contributions of the projects to the milestones of this domain but obviously there are some objectives that will not be met with the FET-HPC effort.

The topic of NRAM support is addressed with, we expect, significant contributions.

The topics of container support or of security which are very important if we want to open more the HPC system are almost not tackled by the projects. There is here clearly a worrying gap that will need action. We cannot leave the integration of security in HPC systems a void field in Europe. EPI is working on the subject at chip level but additional efforts at system level would be welcomed.

The support of data, resilience and visualization in the stack is addressed by some projects. The AI support seems weaker if we consider the quick swift of applications toward this technology.

In this domain of system software, it must be mentioned that most of the contributing projects are limited efforts with several objectives. So, the real progress might not be significant enough to achieve the different milestones. Horizontal project has suggested at the end of the previous section might be more relevant if we want to make a real progress toward these milestones.

Programming environment

number of projects addressing the milestone			
	total FET14	FET16-17	total
M-PROG- 1: APIs and corresponding libraries, run-time and compiler support for auto-tuning of application performance (incl. energy use) and supporting legacy codes.	3	5	8
M-PROG- 2: High-level programming and domain specific language frameworks	1	6	7
M-PROG- 3: Non-conventional parallel programming approaches (i.e. not MPI, not OpenMP / pthread / PGAS - but targeting asynchronous models, data flow, functional programming, model based).	1	6	7
M-PROG-4: Enhanced programming model and run-time system support for dynamic environments (management & monitoring), optimisation of communication and data management, interaction with OS or VM - within application workflows.	2	5	7
M-PROG-6: Performance Analytics and Debugging tools at extreme scale, including data race condition detection tools and user-support for problem resolution.	3	1	4
M-PROG-7: Performance analytics and debugging tools co-designed to link to the application developer's original code and high-level programming environments	1	0	1

This is a well-covered domain with many projects working toward the milestones. The first four milestones are central to some of the FET-HPC-2017 projects. We expect strong contributions to progress toward these milestones.

The performance analytics and debugging tools topic has been addressed mainly by FET-HPC-2014 projects. To fully fulfil this milestone a horizontal project that will aggregate the different European tools could be a good idea.

The last milestone is very specific. Only one project has a weak contribution to it but this gap is not as worrying as for other more strategic milestones.

Energy and resiliency

number of projects addressing the milestone			
	total FET14	FET16-17	total
M-ENR-MS-1: Characterisation of computational advance as function of energy/power metric and standardisation of this approach with automatic and semi-automatic tools	1	2	3
M-ENR-MS-2: Methods to manage computational advance based on the pre-set energy/power metric and achieve Proportionality Computing with respect to the selected metric	0	1	1
M-ENR-MS-3: Throughput efficiency increase by scheduling instructions to the cores and functional units in the processor within its power envelope and taking the time criticality of the instructions into account	0	0	0
M-ENR-HR-4: Optimisation of the energy spend by the facility by controlling the coolant temperature down to the device level and taking the infrastructure energy cost into account	1	0	1
M-ENR-FT-5: Collection and Analysis of data from sensor networks - the Big Data challenge for measurements around the facility	0	0	0
M-ENR-FT-6: Prediction of failures and fault prediction algorithms	2	3	5
M-ENR-FT-7: Application recovery from fault conditions in the system	0	3	3
M-ENR-AR-8: Energy/Power efficient numerical libraries	4	3	7
M-ENR-MS-9: Highly efficient HPC installation	0	0	0

The gap analysis for this domain is much more difficult because of the nature of the milestones. Some are very precise and specific to one approach of the energy/resiliency issues as M-ENR-MS-3 or M-ENR-HR-4. Some projects may have addressed the underlying issue but not with the milestone proposed approach.

The main finding is that some projects have tackled the energy and resiliency topics. However, the impression is that the technologies developed are only part of the solution. To address energy efficiency, a horizontal project could be a way to progress further. To address resiliency, large integration projects could be the best option, ensuring their global scope and integration of component level resiliency features in a coherence manner.

Balance compute, IO and storage performance

number of projects addressing the milestone			
	total FET14	FET16-17	total
M-BIO-1: One or more storage class memory technology usages demonstrated as part of the persistent storage hierarchy	2	3	5
M-BIO-2: Extreme scale storage and I/O system simulation framework established.	2	0	2
M-BIO-3: Standardised Extreme scale I/O middleware API available: incorporating advanced features such as data layouts on NVRAM/Flash/Disk, in-storage computing, Object stores, etc, and also portability concerns raised by the CoEs.	1	2	3
M-BIO-4: Big Data analytics tools developed and optimised for Storage and I/O.	1	4	5
M-BIO-5: In-storage compute capability across all tiers/layers of the storage system as indicated by the data requirements within the CoEs.	0	0	0
M-BIO-6: I/O Quality-of-Service capability available for extreme scale storage systems	0	1	1
M-BIO-7: Extreme scale multi-tier data management tools available	1	2	3
M-BIO-8: Completion of co-Design with new use cases identified by the CoEs (AI/Deep learning, etc.)	2	3	5

There are multiple contributions from the FET-HPC projects. The technologies developed by the projects will help to achieve the milestones even if some of them will need additional efforts.

The new storage hierarchy introduced by NVM is a topic which is correctly addressed with several strong contributions. We have also seen effort toward the integration of the big data oriented application needs in terms of IO.

The less covered aspects are the simulation framework and the M-BIO-5 milestone. The latter does not seem to be mandatory to deliver good IO performance in exascale system and the absence of research effort is not a major concern.

Big data and HPC usage models

number of projects addressing the milestone			
	total FET14	FET16-17	total
M-BDUM-METRICS-1: Data movement aware performance metrics available	0	0	0
M-BDUM-METRICS-2: HPC-like performance metrics for Big Data systems available.	0	0	0
M-BDUM-METRICS-3: HPC-Big Data combined performance metrics available	0	0	0
M-BDUM-MEM-1: Holistic HPC-Big Data memory models available	0	3	3
M-BDUM-MEM-2: NVM-HPC memory and Big Data coherence protocols and APIs available.	0	3	3
M-BDUM-ALGS-1: Berkeley Dwarfs determination for Big Data applications available	0	0	0
M-BDUM-ALGS-2: Dwarfs in Big Data platforms implemented	1	0	1
M-BDUM-PROG-1: Heterogeneous programming paradigms for HPC-Big Data available	1	5	6
M-BDUM-PROG-2: Heterogeneous programming paradigm with coherent memory and compute unified with Big Data programming environments available	0	0	0
M-BDUM-PROG-3: Single programming paradigm across a hybrid HPC-Big Data system available	1	0	1
M-BDUM-VIRT-1: Elastic HPC deployment implemented	0	0	0
M-BDUM-VIRT-2: Full virtualisation of HPC usage implemented	0	0	0
M-BDUM-DIFFUSIVE-1: Big Data - HPC hybrid prototype available	0	2	2
M-BDUM-DIFFUSIVE-2: Big Data - HPC large-scale demonstrator integrated	0	0	0

It is obvious that the FET-HPC projects will not achieve most of the domain milestones. These milestones are ambitious and may need time and experience in dealing with big data/AI/HPC combined applications to identify research paths that would be relevant to deliver the milestones.

The need of new metrics and dwarfs is not covered by the current effort but again the lack of experience in dealing with the new type of applications could be the reason why more time is needed before having the right performance assessment framework.

An issue that is addressed is the heterogeneous nature of the new application domains. However, the current level of research is unable to reach the M-BDUM-PROG-2 milestones.

The level of effort on virtualization is more worrying. The current research activity will not be enough to move toward the milestones defined in the SRA. This is a concern if we want HPC systems to be used by a large number of user communities.

The access to prototypes and demonstrators also seems rather weak compared with the stakes and the objectives of moving fast toward new mixed BD/AI/HPC applications.

Mathematics and algorithms for extreme scale HPC systems

number of projects addressing the milestone			
	total FET14	FET16-17	total
M-ALG-1: Scalability of algorithms demonstrated for forward in time computing and 3-dimensional FFT for current architectures	0	1	1
M-ALG-2: Multiple relevant use cases demonstrated for improving performance by means of robust, inexact algorithms with reduced communication costs	2	3	5
M-ALG-3: Scalable algorithms demonstrated for relevant data analytics and artificial intelligence methods.	3	4	7
M-ALG-4: Processes established for co-design of mathematical methods for data analytics and of HPC technologies/architectures	0	1	1
M-ALG-5: Classes of data, partitioning and scheduling problems categorised and their complexity ascertained	0	0	0
M-ALG-6: Mathematical and algorithmic approaches established for the scheduling of tasks on abstract resources and exploitation of multiple memory levels	1	1	2
M-ALG-7: Research on mathematical methods and algorithms exploited for compiler technologies, runtime environments, resource schedulers and related tools.	1	0	1
M-ALG-8: Reduction of energy-to-solution demonstrated by means of appropriately optimized algorithms demonstrated for a set of relevant use cases.	0	1	1
M-ALG-9: Process for vertical integration of algorithms established together with the validation of scalability, ease of implementation, tuning and optimisation	2	0	2
M-ALG-10: Tuning of algorithmic parameters at exascale completed for a relevant set of algorithms.	0	1	1

The FET-HPC effort contributed to this domain is mainly due to the five FET-HPC-2014 projects and the four FET-HPC-2017 projects which are focused on application domains. There are also some modest contributions from the technology oriented projects which due to their co-design approach have tested new algorithms or implementation methods.

The contributions are concentrated on two topics: scalable algorithms for BD/AI and robust algorithms. The former subject is tackled by several projects that have worked on the improvement of BD/AI algorithm implementation. The latter topic is addressed partly by the projects working on Uncertainty Quantification and on some generic application domains.

The other milestones are less covered. Globally the FET-HPC projects will not have the resources needed to make a significant progress toward all the milestones of this domain. Other efforts such as the Centres of Excellence are also very relevant to tackle the domain challenges.

3.3 Main findings of the gap analysis

This analysis of the contributions of the FET-HPC projects to the SRA3 milestones provides an insight on how the technology research strategy is currently implemented. The main findings are:

- All the SRA3 domains are covered by the projects
- For the majority of the milestones, there are contributions which will facilitate their achievement;
- Some gaps exist, the most worrying being:
    - Integration of security in HPC systems
    - Support of virtualisation to open HPC system usage
    - Emergence of performance metrics for the new applications
- Most of the milestones are too ambitious to be achieved within a single FET-HPC project.

This gap analysis reinforces our conviction that, in order to implement a complete technology strategy, some additional actions besides the research done by FET-HPC-like projects are needed. The recommendations presented in Section 2.4 are a way to complement these excellent contributions of the FET-HPC projects.

This gap analysis is also valuable from a methodology point of view. Issuing a Strategic Research Agenda represents an important effort. One of the SRA objectives is to steer research projects toward achievements that are important to the progress of the whole field. The current analysis shows how to define these research objectives to make them relevant to prepare research projects. The main lessons are:

- The milestones should be stated independently from any potential solution
- They should be considered mandatory steps toward efficient exascale systems (and now post-exascale systems)
- They should highlight new trends where more innovative research is needed

In general, the quality of the SRA3 milestones is good but it would be interesting to analyse how improvements can be obtained. We hope that these observations can be helpful for the upcoming Multiannual Strategic Research and Innovation Agenda which will be issued by the Research and Innovation Advisory Group of the EuroHPC Joint Undertaking.

4 Impact on applications

The FET-HPC projects have also an impact on application domains. Most of them have a research methodology based on a "use cases", including several steps:

- Selection of relevant use cases for their research field
- Analysis of the use case requirements
- Design of new technologies with sufficient features to address the requirements
- Assessment of the benefits of the new technologies for the use cases.

Such a methodology leads to work on enhanced versions of the use cases which are often applications that have a large user base. The FET-HPC projects contribute to the effort on applications in the European HPC ecosystem.

A database of the codes that are used by the FET-HPC projects (and also by the nine domain oriented Centres of Excellence[22]) has been compiled from diverse information sources. The main fields of the database are presented in Annex 5. The database is not exhaustive but it is a solid basis for a preliminary analysis.

The first finding is that indeed the FET-HPC projects have a significant activity dedicated to applications. In total, there are 153 actions of FET-HPC projects related to work on an application.

The domains which are the most represented are (with the list of the codes that are used by more than one project or CoE):

- Computational Fluid Dynamics with the main codes being Alya, AVBP, NEK5000, Ludwig, openFOAM
- Material science and molecular dynamics with the main codes being BAC, GROMACS, LAMMPS, MiniMD, Quantum Expresso, CPMD, CP2K, QMCPACK, SIESTA, KKRnano;
- Weather, climate and space weather with the main codes being IFS, ICON, NEMO, iPIC3D
- Human body simulation with the codes NEST, HemeLB, HemoCell
- Equation solvers or libraries with contributions to ExaHyPE, ESPRESSO, FEniCS and SMURFF
- Cosmology and astronomy with contribution to GADGET and SKA data processing.

There are around 30 codes that have more than one contribution from the FET-HPC projects and the CoEs (in the range of 15 if we consider only the effort of the FET-HPC projects).

We have looked in more detail at the codes that are used by more than one FET-HPC project. Through the activities on application, there might be relevant interactions or synergies between the projects. There are more than 15 cases of applications used by more than one project. We have focused on the projects that are still running because it could be difficult to try to initiate cooperation with the FET-HPC-2014 projects that have ended. Table 8 shows the result of this analysis (as this is a symmetrical matrix, we have only filled half of it).

We can see seven potential interactions, two of them on two applications. The applications are among the one that are the most used in the research currently done in Europe.

22 The POP CoE which is horizontal has not been considered in this analysis.

	DEEPEST	EUROEXA	RECIPE	SAGE2	MAESTRO	ASPIDE	EPIGRAM-HS	EPEEC	EXA2PRO	EXAQUTE	VECMA	VESTEC	ESCAPE2
DEEPEST		NEST		SKA data processing				AVBP					
EUROEXA								Quantum ESPRESSO SMURFF					IFS NEMO
RECIPE													
SAGE2												iPIC3D	
MAESTRO													
ASPIDE													
EPIGRAM-HS													IFS
EPEEC													
EXA2PRO													
EXAQUTE													
VECMA													
VESTEC													
ESCAPE2													

Table 8: Common codes of the FET-HPC-2016-2017 projects

The same kind of analysis have been performed for the codes that are common to the FET-HPC projects and the CoEs. This could lead to interesting interactions between the projects and the CoEs. This analysis has been limited to the FET-HPC-2016-2017 projects. The results are shown in table 9.

There are 18 couples CoEs-FET-HPC projects. In two cases, the projects work on three common codes and in one case on two common applications. The most represented application domains are CFD, molecular dynamics and weather/climate.

	MAX	E-CAM	BIOEXCEL	COMPBIOMED	ESIWACE	ECoE	HIDALGO	CHEESE	EXCELLERAT
DEEPEST		GROMACS	GROMACS						AVBP
EUROEXA	Quantum ESPRESSO	Quantum ESPRESSO		Alya	IFS NEMO	Alya			Alya
RECIPE									
SAGE2									
MAESTRO									
ASPIDE									
EPIGRAM-HS					IFS				NEK500
EPEEC	Quantum ESPRESSO	Quantum ESPRESSO							AVBP
EXA2PRO									
EXAQUTE									
VECMA		LAMMPS		BAC HemeLB HemoCell			RepastHPC		
VESTEC									
ESCAPE2					ICON IFS NEMO				

Table 9: Common codes FET-HPC-2016-2017 CoEs

To summarise, the FET-HPC projects have also significant contributions to applications because most of them refer to use cases to assess their new technologies. This application angle could be a way to find potential interesting cooperation or synergies among the FET-HPC projects themselves and/or with the CoEs.

5 Potential links between FET-HPC projects and the European Processor Initiative (EPI)

The EPI (https://www.european-processor-initiative.eu/) project is central to the European HPC strategy. It aims at developing the processor and accelerator chip(s) that could be integrated in European flagship HPC systems. The initiative is managed by a consortium of 27 partners linked to the European Commission by a FPA (Framework Partnership Agreement). The actions within this FPA are defined in SGAs (Specific Grant Agreements). For the time being only one SGA has been signed by the EC and the EPI consortium.

The EPI consortium is strongly connected with the FET-HPC projects as EPI members are involved in the majority of these projects. This intersection would be the easiest option to achieve a transfer of the results and know-hows generated by the research projects. We will analyse these interactions first.

We have also approached the EPI project team to better understand what are the technologies that could complement their current actions. The outcome of these discussions is presented in Section 5.2.

For each applicable FET-HPC project, we have looked at the results (for finished projects) or the current activities (for on-going projects) that could be helpful for the EPI initiative.

This Section is concluded by a summary of the actions that can be recommended to facilitate the reuse of FET-HPC project results in EPI.

5.1 EPI partners involvement in FET-HPC projects

The partners of the EPI SGA1 (Specific Grant Agreement) have been strongly involved in the FET-HPC projects that have been launched before the start of the EPI (in December 2018).

The two tables (see below) present the intersection between EPI partners and members of the FET-HPC-2014 projects and of FET-HPC-2016-2017 projects.

		Atos	BMW	BSC	CEA	Chalmers	Cineca	E4	Elektrobit	ETHZ	Extoll	Forth	Fraunhofer	GENCI	IST	Infineon	FZJ	Kalray	KIT	Menta	Prove & Run	SiPearl	SMD	SurfSara	STM	UNIBO	UNIPISA	UNIZG	Total
ExaNoDe	hardware; interposer; unimem architecture; ML accelerator; OmpSs for unimem; openstream data flow;set of mini-applications	1		1	1					1		1	1				1	1											8
ExaNeST	FGA system; DMA engine for unimem architecture; multi-path routing interconnect; FPGA for interconnect; cooling technology;											1	1																2
ECOSCALE	FPGA system; unimem architecture; fpga bit stream tool; automatic tuning kernel tool					1						1																	2
montblanc 3	Marvell ThunderX2 platform; ARM HPC software stack	1		1						1																			3
NEXTGenIO	slurm extension for NVRAM support; workload simulator; file system for NVRAM			1																									1
SAGE	MERO object file system with Clovis API; HSM software; runtime for pre/post processing; IOR benchark; Flink connector	1			1												1												3
MANGO	FPGA system; RTL for many core acceerator; RTL for interconnect; two phase cooling system; runtime BarbequeRTRM																												0
EXTRA	FPGA system; CAOS environment for FPGA system; performance model; polymorphic register file design tool																												0
Green flash	FPGA system to control telescope mirrors; library to run RT application on GPU or FPGA; MOAO benchmark to test runtime as ompSs, StarPU;																												0
READEX	energy efficiency; extension of Score-P to support eergy measurement; connector to retrieve enrgy ad power values; MERIC library for instrumentation, profiling and optimization; API for several profiling libraries																												0
Antarex	energy efficiency; DSL and source to source compiler; cahing mecanism to avoid recomputation; management of code fragment versions; power management runtime for MPI codes; application demonstrators						1			1																			2
INTERTWINE	programming model; GASPI library; OmpSs enhanced version; StarPU supporting resource manager API; API to abstractcommunication library; benchmark			1									1																2
ALLScale	distributed system; API to manage memory in distributed system; source to source compiler and associated ; runtime; failure detection and checkpoint restart; two layer API to describe data structure s and parallel operations																												0
ComPat	distributed systems; QCJ job manager; simulator of large computing system																												0
ExaHyPE	application hyperbolic systems; library for solving hypebolic systems; seismic and astronomy applications																												0
ESCAPE	application weather forecast; dwarfs and benchmark suite;	1																											1
NLAFLET	application linear algebra; library for sparse and dense linear matrix																												0
Exaflow	application CFD; Nek5000 enhancement; waveet based compression; ExaGS gather scatter library optimized for different architectures																												0
ExCAPE	application machine learning; library ofr ML																												0
	total	4	0	4	2	1	1	0	0	3	0	3	3	0	0	0	2	1	0	0	0	0	0	0	0	0	0	0	24

Table 10: Participation of EPI partners in FET-HPC-2014 projects

		Atos	BMW	BSC	CEA	Chalmers	Cineca	E4	Elektrobit	ETHZ	Extoll	Forth	Fraunhofer	GENCI	IST	Infineon	FZJ	Kalray	KIT	Menta	Prove & Run	SiPearl	SMD	SurfSara	STM	UNIBO	UNIPISA	UNIZG	Total
DEEP-EST	system project; modular supercomputer architecture (MSA); co-design methodology; interconnet fro dis-aggregated resources; network attached memory NAM; gateway for multinetwork communication; software distribution management on MSA; parastation MPI and ompSs implementation; resiliency support; slurm enhancement			1							1		1				1												4
EuroEXA	system with ARM and FPGA; unimem architecture support; container technology; cooling technology; FPGA interconnect implementation; MPI and Gaspi communication library; slurm for FPGA; methodology for FPGA application developent and performance assessment			1								1	1																3
SAGE2	data management; global memory model that use he underlying memory hierarchy; Mero object storage software	1			1												1												3
MAESTRO	memory aware and data aware middleware; optimizaton of data movement at all level of the stack (from Os to application)				1					1							1												3
RECIPE	heterogeneous resource management environment; FPGA focus; energy efficiency and reliability; hardware abstraction layer that is use by BarbequeRTM and at a higher level slurm			1																									1
EPEEC	programming model for heterogeneous systems; directive based approach; profiling tools; automatic code generation targeting accelerator; using OmpSs, Gaspi, parallware and ArgoDSM			1			1						1																3
Epigram-HS	programming model for heterogeneous systems; ML focus; MPI (GPI) for FPGA support; automatic data placement									1			1																2
EXA2PRO	programming model for heterogeneous systems; skeleton based approach SkePU; composition layer; sorce to source compiler using Mercurium; runtime based on StarPU; multi criteria scheduling, foaul tolerance and plaftform modelling features																1												1
ASPIDE	programming model for data applications; DSL over mapreduce model; autotuning features; focus on IOPS;	1																											1
VECMA	verification, validation and uncertainty quantification approach; focus on multiscale	1																											1
ExaQUte	verification, validation and uncertainty quantification approach; focus on stochastic methods			1																									1
VESTEC	urgent decision making; visualization of data analytics in real time; coupling of data analytics with simulation																												0
ESCAPE-2	application for weather forecast; new model and associated DSL; uncertainty quantification; benchmark	1		1	1																								3
	total	4	0	6	3	0	1	0	0	2	1	1	4	0	0	0	4	0	0	0	0	0	0	0	0	0	0	0	26

Table 11: Participation of EPI partners in FET-HPC-2016-2017 projects

The tables show that four EPI partners, namely Atos (8 participations), BSC (Barcelona Supercomputing Centre, 10 participations), Fraunhofer (7 participations) and FZJ (Forschungszentrum Jülich, 6 participations) have strong connections with the FET-HPC projects and can play a major role in the exploitation of these project achievements. As they have some key roles in the EPI, the transfer of the IP will be even easier. In addition, there are three research organisations, namely CEA, ETHZ and Forth, which have also an important role both in EPI and some of the FET-HPC projects.

There are 10 projects from the FET-HPC-2014 call and one project from the FET-HPC-2016-2017 call that are not related to the EPI consortium. These projects will be carefully looked at to analyse whether it is appropriate to set-up their relationships with the EPI.

5.2 EPI developments and potential complementary actions

EPI is a stand-alone project that plans to develop all the technology bricks required to deliver a product to the market, at the end of the project. It should not depend on other projects to reach this objective. However, there could be areas in which complementary actions could accelerate the achievement of this goal and could increase its impact on the market. These complementary actions could be either oriented toward hardware or software technologies.

5.2.1 *Hardware oriented complementary actions*

The EPI has introduced the "EPI Common Platform" (https://www.european-processor-initiative.eu/european-processor-initiative-first-year-of-activities/) which is composed of Arm architecture general purpose cores and prototypes of high energy-efficient accelerator tiles: RISC-V based (EPAC), Multi-Purpose Processing Array (MPPA), embedded FPGA (eFPGA) and cryptography HW engine (see Figure 2).

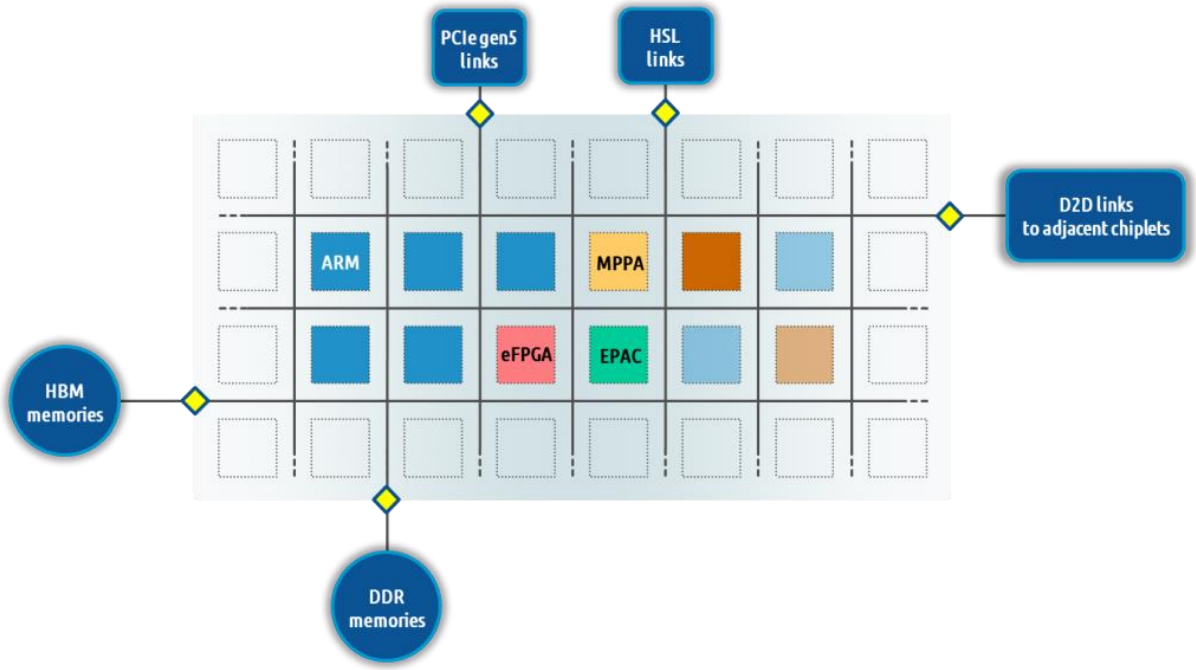

Figure 2 - EPI common platform architecture concept

This architecture provides two levels of cooperation with hardware-oriented R&D intended by entities outside of the EPI consortium:

- Integration at the network-on-chip (NOC) level with components designed outside the EPI project
- Integration at interposer level of chips founded outside the EPI project completing the EPI chip (e.g. an AI accelerator on the same interposer as the EPI chip and memory modules)

To enable cooperation, the interfaces at these two levels need to be made available by the EPI, allowing interested projects to develop interoperable hardware solutions.

5.2.2 *Software oriented complementary actions*

EPI platform software activities encompass the following components:

- BIOS
- Firmware
- Software platform
- Compilation chain
- Associated runtime

As the general purpose core uses the ARM architecture, the results of the set of Montblanc projects will be reused to provide some of the basic elements of the software stack. Nevertheless, there are at least two important aspects of the EPI architecture that will require additional features in the software stack:

- Presence of accelerator(s) leading to heterogeneous computing resources;
- Co-existence of high bandwidth memory HBM and large capacity memory DDR5 (Dual Data Rate 5)

To be able to efficiently use these two elements, some new software developments are needed. The presence of accelerators of potentially different types will require to develop specific compilation chains and libraries and to have either at compile time or at runtime the ability to decide which computing resources to use. The EPI developments will propose a solution but some other approaches could also be interesting and could provide improvement especially if they are coupled with DSLs (Domain Specific Language) which help to take efficient decisions by encapsulating some domain semantics about the data and the operations.

For the HBM and DDR management, the fact that the EPI does not intend to propose a cache mode (HBM being automatically used by the system to provide a new level of cache memory) means that there is an important effort to achieve a good placement of the data at application level. Very few architectures have this double level of memory. It was the case for the Intel Knights Landing (KNL) that was discontinued in 2018. Some efforts have been made to exploit the flat mode of KNL that could be reused for EPI but the announcement of the end of the Xeon Phi family has put an end to most of the efforts. The EPI will develop a solution but other options could be proposed and help to handle this difficult problem of memory allocation optimisation.

Besides these two important subjects for EPI, there could also be complementary actions on the following topics:

- Support for additional programming models
- Profiling and tuning tools
- Energy efficiency management
- Specific developments for new market targets

The EPI has plans to support different programming models including MPI+OpenMP[23]. For the time being, the support of PGAS (Partitioned Global Address Space) approaches is not a priority. As Europe has developed with GASPI a flavour of PGAS that has some user base, it could be interesting to support this programming model on EPI platform[24].

European teams have also developed a large number of profiling and tuning tools (extrae, paraver, dimenas, scalasca, cube, Score-P, Extra-P, Vampir, TAU, MAQAO, SimGrid…). These pieces of software are useful to optimise the applications especially when the architecture is heterogeneous and complex as in the EPI case. The availability of the profiling and tuning tools will be an advantage for the take-off of EPI based platform.

Regarding energy efficiency, EPI develops a low level framework for the power management based on a software platform. This action could be complemented by higher level environments that will deliver information on the power utilisation of applications. Such software is important both for the system administrators to optimise the energy and for the application developers to analyse the impact on power consumption of different versions.

The EPI targets the HPC and the automotive markets at first. A lot of effort is made to deliver the software stack for the HPC domains. Most of it can be reused to address automotive use. In addition, specific software is developed to give automotive domain users the whole environment they are used to. For example, the support of AUTOSAR will be provided by the EPI developments to foster the use of EPI chip in the automotive market. Beside these two markets, additional specific developments could open new markets to the EPI chips. Even if it is not the current focus of the EPI, some complementary actions could be planned in parallel in order to prepare the penetration of new markets by the EPI which might increase the volume of chips and in consequence the return on investment of the initiative.

5.3 Relevant technologies developed by FET-HPC projects

In this paragraph, we identify the technologies developed by the FET-HPC projects that could be used by the EPI. We present our findings either by project or a group of projects and we show the technologies that could be useful for the EPI, including ideas on how the interaction needed could be organised.

MontBlanc 3 has developed a HPC software stack for platform based on ARM architecture processor which is the EPI choice. The results of MontBlanc 3 are already integrated through the participation of Atos and BSC two main contributors of both projects.

The co-design methodology developed by **DEEP-EST** is also already used by the EPI. This transfer is easy due to the central contribution of FZJ in both projects. We will come back below on other technologies coming from DEEP-EST that could be useful for EPI.

The **ExaNode** project has developed and active interposer technology for the integration of chiplet. This specific technology may not be appeal to the EPI for industrial reasons, but the experience gained in ExaNode is valuable for the EPI. A key partner of both projects, i.e. CEA will ensure that the transfer of the relevant knowhow and experience will happen.

[23] Message Passing Interface + Open Multi-Processing
[24] This is planned for EPAC but a more general support could be interesting

The group of projects **ExaNeST**, **ECOSCALE** and **EuroEXA** has also worked on ARM based environment and on heterogeneous resources with CPU and FPGA. These two domains have synergies with the EPI approach. However as there are different technical choices for the implementation, it remains to be confirmed that some of the technologies be reused in the context of the EPI. This could be analysed in the dialogue framework that is suggested at the end of this Section.

The **NEXTGenIO** project has developed some software components to manage a complex memory hierarchy. One of the challenges of the EPI platform will be to manage the different levels of memory especially the DDR and HBM levels. Whether NEXTgenIO technologies can be applied to the EPI case remains to be seen. However, a technical analysis may be worthwhile.

The three data and storage projects **SAGE**, **SAGE2** and **MAESTRO** have developed technologies that are not specific to any given platform. Nevertheless, it could be good for the systems based on EPI to have a rich set of software to increase their attractiveness for users. The environments developed by these three projects could be good candidates for an innovative data management stack.

The results of **ANTHAREX**, **READEX** and **RECIPE** could also be integrated and they could offer to the EPI platforms an efficient energy management framework. Of course, some adaptations to the specificities of the EPI chip energy counters and energy management controls will be needed but the exploitation of the results of these projects could provide the EPI platform with a competitive advantage, with some energy savings due to the this energy management framework.

The FPGA oriented projects **MANGO**, **EXTRA** and **RECIPE** have developed FPGA programming tools. As the EPI plans to produce chips based on its common platform architecture that will include FPGAs, a FPGA programming framework will be needed. As the level of FPGA integration is different (in the previous project, FPGAs can be viewed as external accelerators whereas in the EPI they can access the cash) there may be some adaptation to be done. Again, a dialogue between the projects technical experts and EPI needs to be organised.

DEEP-EST, InterTwine, EpiGRAM-HS, ASPIDE, EPEEC and **EXA2PRO** have all worked on programming environment for heterogeneous architectures including CPU, accelerator and FPGAs. Even if these technologies have been developed for heterogeneous resources integrated at the server or systems levels whereas the EPI will offer a package integration, there might be interesting results to be adapted to this EPI approach within the work of these project. The programming environment of the EIP will be key in a fast adoption of this technology. It is worth to investigate how the current effort of the FET-HPC projects on programming environment could help to propose a more complete software for the EPI platform.

The same objective to propose a rich environment for the EPI platforms could be supported by the results of the **EXTRA, NLAFET, ExaFlow** and **ExaHyPE** projects. These projects have developed libraries that are important for some application communities in Europe. The work have been in some case to redesign in order to increase the scalability and to be able to integrate accelerators. Even if additional adaptations to the instruction set and architecture of the EPI platform are needed, the three project partners have the experience to tune the libraries and to provide to the EPI platform a richer environment which will correctly exploit the resources of EPI chip. As the memory hierarchy management of the EPI platform may be too complex for a basis application developer, it is important to invest in libraries for which highly skilled people can precisely optimised the use of the DDR and HBM and deliver easy to use high performance codes.

5.4 Summary

The EPI project is a critical research and development action in the European HPC strategy. As the consortium members have been very active in the FET-HPC projects, they can exploit the results of some of these projects provided they are in line with the needs of the EPI. Apart from this transfer taking place through the EPI consortium members, it would be beneficial to set-up a dialogue between the EPI and other research organisations of the European HPC ecosystem. This would be a win-win situation because of:

- Additional opportunities for the ecosystem to perform research that will have an impact on one of the strategic axes of the European HPC strategy
- Richer environment for the EPI products and a more attractive offering in a competitive landscape.

In order to organise this dialogue, we see several actions that could be put in place:

1. To facilitate interaction at the hardware level, if the EPI consortium wishes to promote an open ecosystem, the release of some interface specifications could be carried out. There are three potential levels of integration and the EPI could publish either all of them or only part of them, again based on their strategy:
    a. Specifications for integration at the die level of accelerators or specialized ASIC IPs; the objective of this action would be to have additional IPs that can be integrated at NOC[25] level.
    b. Specifications for integration at the package level; this could lead to additional accelerators founded independently to the EPI chips; they would be tightly integrated in a single package and could share some memory access.
    c. Specifications of a memory coherent protocol (if supported by EPI and if different from the previous level); this would enable outside package accelerators that will share some memory access with the EPI chips.

    The release of the specifications could be either public or on a private basis after an agreement between the EPI consortium and the teams that would work on providing complementary components.

2. To facilitate the development of cooperation at software level, some thematic workshops (software stack, programming environment, libraries…) could be organised between the EPI team and other organisations willing to complement the current effort. These workshops could lead to cooperation proposals from the external EPI teams that could be accepted as a one to one relationship by the EPI consortium. Several funding scheme for the related research could be proposed. This could be the form of self-investment that would generate return on investment through IPs valorisation. This funding could come from a cascade funding budget given by EuroHPC to the EPI that would use it for the applicable proposals.

3. To increase the level of investment on EPI complementary software, a specific call can be organised by EuroHPC with an objective to select projects that will develop software components helping to have a rich environment for the EPI future products. This call could be either completely open or if we prefer to give a priority to the exploitation of the FET-HPC effort it could be restricted to the teams that have already been funded (as it is today the case for the "FET Innovation Launchpad" calls). In any case, the selected

[25] Network on Chip

consortium would have to sign an agreement with the EPI consortium before being funded.

The timing of these potential actions needs to be organised around the EPI timeline. It is obvious that some of the researches and developments that will be the result of this dialogue will require the existence of simulation environment, part of the compilation stack, availability of prototype servers, software development kit, etc. A time dependency analysis will be needed to plan the complementary effort.

In conclusion, the EPI is a strategic axis of the European HPC plan. It can be complemented by additional actions that have to be tightly coordinated with the EPI consortium to provide a rich and appealing environment that will help this new technology to reach market acceptance. Europe has the potential to further develop existing assets especially in the software area that could deliver, for example:

- An adhoc programming environment for the heterogeneous architecture of the EPI;
- Specific FPGA programming environment;
- Profiling and tuning tools;
- Energy management framework;
- Application domain libraries.

The existence of such a software ecosystem would be a big asset in facilitating the market uptake of EPI chips.

6 Recommendations for a structured technology European effort

The following are the conclusions of this EXDCI-2 work aimed at understanding the performance of the FET-HPC projects:

- The FET-HPC projects have delivered a large set of IPs, which represents an important asset for the entire European HPC community;
- Relating this asset to the SRA provides a valuable insight, which could be used in the preparation of a more detailed future HPC development strategy;
- The FET-HPC projects have an impact not only on the technology but also on the applications leading to interactions with the application pillar implemented through the Centres of Excellence initiative;
- The results developed by the FET-HPC projects can serve as the basis of a rich software ecosystem that could help the EPI project penetrate the market.

As the next step, Europe could choose a very integrated approach like the US Exascale Computing Project or the Japanese Fugaku project. This path will require strong leadership in the area of both the technical and managerial aspects (i.e. in order to select a team of recognised people in charge of the technical choices and the management of the efforts and resources). The experience of the US and Japan shows that this approach can deliver good results but it also takes time to be implemented (ECP is a 7 years project started in 2016 up to 2023; post-K (Fugaku) started in 2014 and will be completed next year).

Another option could be to continue with the framework of calls used in Horizon2020 which leads to a set of research efforts. To increase the impacts of this approach at least four orientations could be proposed:

- A more programmatic vision for the preparation of the calls;
- High TRL objectives for the calls;
- A selection of the project based on their contribution to the overall objectives;
- A mandatory cooperation between the selected projects.

HPC has a very structured technology value chain. This is true for hardware (processor, memory, interconnect, storage, cooling, energy, etc.) and for software (bios, firmware, server management, system management, programming models, compiler, runtime, libraries, etc.). It is possible to choose the components which should constitute the European research priorities and which we require in order to build our exascale and post-exascale HPC infrastructure. This programmatic vision should be the basis of writing the calls with objectives related to this global technology strategy.

The past projects have been efficient to generate results but in the case of most of them the level of development is too low to have a “pre-production” quality needed to build a user community and to enter in a virtuous circle of improvements made by the users themselves. In this second phase of the HPC strategy, a large part of the effort should be dedicated to research and development that will have a correspondingly high TRL level. This is the only way for some of the project results to be used in the exascale and post-exascale infrastructure.

The selection process of the projects has to take into account how the projects will contribute to the overall strategy and should not be based on the own merit of the research only. The selection criteria need to be closely related to the goal selected in a tight manner. The set of the selected projects should also be complementary (unless the resources are sufficient to fund parallel researches on one objective). This complementarity should also be a criteria at the level of the call.

In the case of the past calls, the cooperation between the projects has been hindered by the lack of resources. Each project had its DoA[26] with all the resources allocated to actions not aimed at a cooperation with other projects which were unknown when the proposal was submitted. In result, such cooperation is limited to small reallocation of resources without clear commitments. There are several implementation options to allow for more structured cooperation between projects of the same call (i) The projects could reserve a share of their budget for actions to be defined when the selection of the projects has been made public (ii) An additional budget can be distributed to the projects that propose to add cooperation tasks after a discussion among themselves and before the signature of the grant agreement (iii) An additional budget could be distributed on a change request including additional cooperation tasks decided during the lifetime of the projects. In any case, we should allow for cooperation tasks to be added to the projects activities to increase the global impact of the call.

These four directions would make the European technology research effort more efficient and enable the delivery of results that will be used by the exascale and post-exascale HPC infrastructure.

To summarise, the analysis that has been conducted by EXDCI-2 on the European HPC technology research effort leads to the following recommendations:

1. Maintain a global survey and data base of the results of the HPC technology projects to build a global vision which can be used to update the HPC strategy; EuroHPC team could do this analysis or delegate it to the relevant organisation(s)/project in the HPC ecosystem;
2. Set up calls with higher TRL objectives for the technologies to enter the virtuous circle of continuous improvements by their user bases; These calls could be:
    a. Integration projects with the objective to deliver a complete HPC solution;
    b. Horizontal projects with the objective to develop a layer that could be used by several European computing centres or application developer communities
3. Develop a programmatic approach of the research programme with a strong focus on the strategic axes as the EPI;
4. Implement a framework to facilitate strong cooperation between the projects selected.

We are confident that these recommendations will improve the impact of the European research technology effort and strengthen European leadership in HPC.

26 Description of Action

7 Acknowledgements

We would like to thank the FET-HPC project coordinators and participants for their participation to this activity, in particular those of the FET-HPC-2014 who have continued to help us although after the completion of their projects.

We would like also to thank Michael Malms for his comments on the gap analysis and Romain Dolbeau from the EPI project who has helped us to identify the potential synergies between the EPI and the FET-HPC projects.

8 Annexes

The annexes are:

- Annex 1: FET-HPC project lists
- Annex 2: FET-HPC-2014 project questionnaire
- Annex 3: Analysis of the impact of Horizon2020 on the HPC value chain
- Annex 4: FET-HPC project relationship with the SRA3 milestones
- Annex 5: Codes used by FET-HPC projects and CoEs

8.1 Annex 1: FET-HPC project lists

FET HPC 2014 projects

ACRONYM	Title	Leader	access to content
ExaNoDe	European Exascale Processor Memory Node Design	COMMISSARIAT A L ENERGIE ATOMIQUE ET AUX ENERGIES ALTERNATIVES	http://exanode.eu/
ExaNeSt	European Exascale System Interconnect and Storage	FOUNDATION FOR RESEARCH AND TECHNOLOGY HELLAS	http://www.exanest.eu/
NEXTGenIO	Next Generation I/O for Exascale	THE UNIVERSITY OF EDINBURGH	http://www.nextgenio.eu
Mont-Blanc 3	Mont-Blanc 3, European scalable and power efficient HPC platformbased on low-power embedded technology	Bull SAS	http://montblanc-project.eu/
SAGE	SAGE	XYRATEX TECHNOLOGY LIMITED	http://www.sagestorage.eu/
MANGO	MANGO: exploring Manycore Architectures for Next-GeneratiOn HPC systems	UNIVERSITAT POLITECNICA DE VALENCIA	http://www.mango-project.eu
ECOSCALE	Energy-efficient Heterogeneous COmputing at exaSCALE	TELECOMMUNICATION SYSTEMS INSTITUTE	http://www.ecoscale.eu/
EXTRA	Exploiting eXascale Technology with Reconfigurable Architectures	UNIVERSITEIT GENT	https://www.extrahpc.eu/
ESCAPE	Energy-efficient SCalable Algorithms for weather Prediction at Exascale	EUROPEAN CENTRE FOR MEDIUM-RANGE WEATHER FORECASTS	http://www.hpc-escape.eu/
ComPat	Computing Patterns for High Performance Multiscale Computing	UNIVERSITEIT VAN AMSTERDAM	https://www.compatproject.eu/
ExCAPE	Exascale Compound Activity Prediction Engine	INTERUNIVERSITAIR MICRO-ELECTRONICA CENTRUM VZW	http://www.excape-h2020.eu/
NLAFET	Parallel Numerical Linear Algebra for Future Extreme-Scale Systems	UMEA UNIVERSITET	https://www.nlafet.eu/
INTERTWINE	Programming Model INTERoperability ToWards Exascale (INTERTWinE)	THE UNIVERSITY OF EDINBURGH	http://www.intertwine-project.eu/
greenFLASH	Green Flash, energy efficient high performance computing for real-time science	OBSERVATOIRE DE PARIS	https://lesia.obspm.fr/
READEX	Runtime Exploitation of Application Dynamism for Energy-efficient eXascale computing	TECHNISCHE UNIVERSITAET DRESDEN	https://www.readex.eu/
ALLScale	An Exascale Programming, Multi-objective Optimisation and Resilience Management Environment Based on Nested Recursive Parallelism	UNIVERSITAET INNSBRUCK	http://www.allscale.eu/
ExaFLOW	Enabling Exascale Fluid Dynamics Simulations	KUNGLIGA TEKNISKA HOEGSKOLAN	http://www.exaflow-project.eu/
ANTAREX	AutoTuning and Adaptivity appRoach for Energy efficient eXascale HPC systems	POLITECNICO DI MILANO	http://www.antarex-project.eu/
ExaHyPE	An Exascale Hyperbolic PDE Engine	TECHNISCHE UNIVERSITAET MUENCHEN	https://exahype.eu/

FET HPC 2016 projects

ACRONYM	Title	Leader	access to content
Deepest	Dynamical Exascale Entry Platform - Extreme Scale Technologies	Juelich	https://www.deep-projects.eu/
EuroEXA	ExaScale supercomputers	ICCS	https://euroexa.eu/

FET-2017 HPC projects

ACRONYM	Title	Leader	access to content
ASPIDE	exAScale Programing models for extreme Data procEssing	U Madrid	https://www.aspide-project.eu/
EPEEC	European joint Effort toward a Highly Productive Programming Environment for Heterogeneous Exascale Computing (EPEEC)	BSC	https://epeec-project.eu/
EPiGRAM HS	Exascale programming model for heterogeneous systems	KTH	https://epigram-hs.eu/
ESCAPE-2	Energy-efficient SCalable Algorithms for weather and climate Prediction at Exascale	ECMWF	http://www.hpc-escape2.eu/
EXA2PRO	Enhancing Programmability and boosting Performance Portability for Exascale Computing Systems	ICCS	https://exa2pro.eu/
ExaQUte	EXAscale Quantification of Uncertainties for Technology and Science Simulation	UPC CENTRE INTERNACIONAL DE METODES NUMERICS EN ENGINYERIA	http://exaqute.eu/
MAESTRO	Middleware for memory and data-awareness in workflows	Juelich	https://www.maestro-data.eu/
Recipe	European Exascale Processor Memory Node Design	Politecnico di Milano – Dipartimento di Elettronica, Informazione e Bioingegneria	http://www.recipe-project.eu/
SAGE2	Percipient Storage for Exascale Data Centric Computing2	Seagate	http://www.sagestorage.eu/
VECMA	Verified Exascale Computing for Multiscale Applications	UCL	https://www.vecma.eu/
VESTEC	Visual Exploration and Sampling Toolkit for Extreme Computing	DEUTSCHES ZENTRUM FUER LUFT - UND RAUMFAHRT EV	https://vestec-project.eu/

8.2 Annex 2: FET-HPC-2014 project questionnaire

FET-HPC 2014 Project Questionnaire

Information on this Questionnaire

Why do we need it?

We are aware of the fact that we have requested your participation in a number of questionnaires that require a lot of time – a scare resource in the case of most of us. Nevertheless, we believe that your investment in this effort is a win-win opportunity for at least three reasons:

1. **The funding for a call such as FET-HPC can continue only if we are able to report some tangible results** and show the progress enabled by this program. Your contribution is a way to create new research opportunities in the next 'Horizon Europe' framework program.
2. **There might be opportunities in the European HPC ecosystem for the exploitation of your project's results.** This questionnaire will help find the right connections and advance your research.
3. **The EXDCI2 project has several activities that can benefit the FET-HPC projects**: support for start-ups, liaison with the ESDs, international collaboration, standardisation or communication actions – all of these can help you in exploiting what has been achieved during your project. The EXDCI2 project plans to use your inputs to help you through these activities.

Your answers will not be public and will only be used to concentrate the EXDCI2 resources on activities that can help you to increase the impact of your project.

This questionnaire is not:

1. An evaluation of your project
2. A statistical analysis that will only serve 'bureaucracy'

How to answer this questionnaire?

We offer various ways of answering the questionnaire:

- You fill the document and send it back to us by mail
- A phone interview: please propose some time slots and we will set up a conference call to fill in the questionnaire with you
- A face-to-face meeting at the EXDCI 2 Booth at ISC 2018: we set up an appointment and complete the questionnaire together.

If you want to distribute the work to other people, we can receive several partial contributions sent back to us or several online form completed (we will forward you all the contributions submitted by your project participants).

The timing:

We expect to receive your feedback in the coming weeks (as an option, we could discuss this material together during the ISC) but you can also answer during the summer.

If you believe that this summer is not the right time for your project to answer this questionnaire, please let us know and tell us when it suits you.

----- Questionnaire starts here -----

Please just write your answers below the questions in the tables below:

This is a question?
Please write your answer here…

If there are multiple choices, please put an x next to your answer:

I am a human	x

Some questions require a yes/no answer.

Have you been to the Moon (yes/no)?
no

Then, save the document and send back to us by replying to the email the document was sent with.

1. Information on the Project
Name of the project:
When do you plan to finish the project?
Do you plan to organise an open conference or a workshop at the end of the project?
If yes, when and where?
Who is the main contact?

2. Results and follow-up
This section aims to describe the outputs of the projects that you consider the main results. *If you want to describe more than one result in each category and if there is no room provided, please duplicate the paragraph.*

2.1 Hardware/hardware IP (Intellectual Property)	
*The questions below refer to the hardware/hardware IP produced by your project – there is one question per each piece of IP. If you have more than one piece of IP to provide an answer on, we provided room for **five** answers (if you have less pieces of IP than the room provided, please leave the rest of the space empty).*	
Have you produced hardware/hardware IP (yes/no)?	
For each hardware/hardware IP:	
HARDWARE 1	
What is the domain of this hardware?	
Processor	
Storage	
Interconnect	
Other (please specify below)	
Short description of the hardware:	
What is the maturity level of the hardware?	
What do you plan to do in the future with this hardware? (please choose at least one)	
Develop it further within another R&D project (e.g., other FETHPC project)	
Develop it further in a project for maturing innovation (e.g.; FET Launchpad Call, Tetracom calls, etc.)	
Integrate it in Extreme-Scale Demonstrators	
Integrate in participant product roadmap	
Creating a start-up with this IP	
No idea yet	
Any other ideas? (free text below)	
What potential do you see for your hardware? Do not censor yourself – express your "wildest dreams", such as for examples: "My new hardware IP is licensed by X HPC system provider and installed in the first European exascale system"	

Who is the main contact if someone is interested in this hardware? (please provide a name)	
Main contact	
Another person	
NEXT - HARDWARE 2	
What is the domain of this hardware?	
Processor	
Storage	
Interconnect	
Other (please specify below)	
Short description of the hardware:	
What is the maturity level of the hardware?	
What do you plan to do in the future with this hardware? (please choose at least one)	
Develop it further within another R&D project (e.g., other FETHPC project)	
Develop it further in a project for maturing innovation (e.g.; FET Launchpad Call, Tetracom calls, etc.)	
Integrate it in Extreme-Scale Demonstrators	
Integrate in participant product roadmap	
Creating a start-up with this IP	
No idea yet	
Any other ideas? (free text below)	
What potential do you see for your hardware? Do not censor yourself – express your "wildest dreams", such as for examples: "My new hardware IP is licensed by X HPC system provider and installed in the first European exascale system"	
Who is the main contact if someone is interested in this hardware? (please provide a name)	

Main contact	
Another person	
NEXT - HARDWARE 3	
What is the domain of this hardware?	
Processor	
Storage	
Interconnect	
Other (please specify below)	
Short description of the hardware:	
What is the maturity level of the hardware?	
What do you plan to do in the future with this hardware? (please choose at least one)	
Develop it further within another R&D project (e.g., other FETHPC project)	
Develop it further in a project for maturing innovation (e.g.; FET Launchpad Call, Tetracom calls, etc.)	
Integrate it in Extreme-Scale Demonstrators	
Integrate in participant product roadmap	
Creating a start-up with this IP	
No idea yet	
Any other ideas? (free text below)	
What potential do you see for your hardware? Do not censor yourself – express your "wildest dreams", such as for examples: "My new hardware IP is licensed by X HPC system provider and installed in the first European exascale system"	
Who is the main contact if someone is interested in this hardware? (please provide a name)	
Main contact	
Another person	
NEXT - HARDWARE 4	
What is the domain of this hardware?	
Processor	
Storage	

Interconnect	
Other (please specify below)	
Short description of the hardware:	
What is the maturity level of the hardware?	
What do you plan to do in the future with this hardware? (please choose at least one)	
Develop it further within another R&D project (e.g., other FETHPC project)	
Develop it further in a project for maturing innovation (e.g.; FET Launchpad Call, Tetracom calls, etc.)	
Integrate it in Extreme-Scale Demonstrators	
Integrate in participant product roadmap	
Creating a start-up with this IP	
No idea yet	
Any other ideas? (free text below)	
What potential do you see for your hardware? Do not censor yourself – express your "wildest dreams", such as for examples: "My new hardware IP is licensed by X HPC system provider and installed in the first European exascale system"	
Who is the main contact if someone is interested in this hardware? (please provide a name)	
Main contact	
Another person	
NEXT - HARDWARE 5	
What is the domain of this hardware?	
Processor	
Storage	
Interconnect	
Other (please specify below)	
Short description of the hardware:	

What is the maturity level of the hardware?	
What do you plan to do in the future with this hardware? (please choose at least one)	
Develop it further within another R&D project (e.g., other FETHPC project)	
Develop it further in a project for maturing innovation (e.g.; FET Launchpad Call, Tetracom calls, etc.)	
Integrate it in Extreme-Scale Demonstrators	
Integrate in participant product roadmap	
Creating a start-up with this IP	
No idea yet	
Any other ideas? (free text below)	
What potential do you see for your hardware? Do not censor yourself – express your "wildest dreams", such as for examples: "My new hardware IP is licensed by X HPC system provider and installed in the first European exascale system"	
Who is the main contact if someone is interested in this hardware? (please provide a name)	
Main contact	
Another person	

2.2 Software/software IP (Intellectual Property)
*The questions below refer to the software IP produced by your project – there is one question per each piece of IP. If you have more than one piece of IP to provide an answer on, we provided room for **five** answers (if you have less pieces of IP than the room provided, please leave the rest of the space empty).*
Have you produced software/software IP (yes/no)?
For each software/software IP:
SOFTWARE IP 1
Is the software open source (yes/no)? If so, what license does it use?
What is the domain of this software?

Firmware	
OS	
Runtime	
System software	
System middleware	
Compiler	
Application middleware	
Application	
Other (please specify below)	
Short description of the software (e.g. core functionality, target usage scenario):	
What is the maturity level of the software (e.g. work in progress, stable, release)?	
Is the software available on a public repository? Please provide a link.	
If the answer to the previous question is “no”, who is the main contact if someone is interested in exploring the use of this software? (please provide a name)	
Main contact	
Another person	
Is there a user base for the software? If so, is this an emerging (e.g. originating from the project) or an established user base?	
What do you plan to do in the future with this software?	
Develop it further within another R&D project (e.g., other FETHPC project)	
Develop it further in a project for maturing innovation (e.g.; FET Launchpad Call, Tetracom calls, etc.)	
Integrate it in Extreme Scale Demonstrators	
Integrate in participant product roadmap	
Creating a start-up with this IP	
No idea yet	
Any other ideas? (free text below)	
What potential do you see for your IP? Do not censor yourself – express your “wildest dreams”, such as for example: ”My software is licensed by Y to increase the efficiency of its HPC cloud offering”, or “My solver is becoming the basis for exascale application in the Z field”	

NEXT - SOFTWARE IP 2
Is the software open source (yes/no)? If so, what license does it use?
What is the domain of this software?
Firmware
OS
Runtime
System software
System middleware
Compiler
Application middleware
Application
Other (please specify below)
Short description of the software (e.g. core functionality, target usage scenario):
What is the maturity level of the software (e.g. work in progress, stable, release)?
Is the software available on a public repository? Please provide a link.
If the answer to the previous question is “no”, who is the main contact if someone is interested in exploring the use of this software? (please provide a name)
Main contact
Another person
Is there a user base for the software? If so, is this an emerging (e.g. originating from the project) or an established user base?
What do you plan to do in the future with this software?
Develop it further within another R&D project (e.g., other FETHPC project)
Develop it further in a project for maturing innovation (e.g.; FET Launchpad Call, Tetracom calls, etc.)
Integrate it in Extreme Scale Demonstrators
Integrate in participant product roadmap
Creating a start-up with this IP

No idea yet	
Any other ideas? (free text below)	
What potential do you see for your IP? Do not censor yourself – express your "wildest dreams", such as for example: "My software is licensed by Y to increase the efficiency of its HPC cloud offering", or "My solver is becoming the basis for exascale application in the Z field"	
NEXT - SOFTWARE IP 3	
Is the software open source (yes/no)? If so, what license does it use?	
What is the domain of this software?	
Firmware	
OS	
Runtime	
System software	
System middleware	
Compiler	
Application middleware	
Application	
Other (please specify below)	
Short description of the software (e.g. core functionality, target usage scenario):	
What is the maturity level of the software (e.g. work in progress, stable, release)?	
Is the software available on a public repository? Please provide a link.	
If the answer to the previous question is "no", who is the main contact if someone is interested in exploring the use of this software? (please provide a name)	
Main contact	
Another person	
Is there a user base for the software? If so, is this an emerging (e.g. originating from the project) or an established user base?	

What do you plan to do in the future with this software?	
Develop it further within another R&D project (e.g., other FETHPC project)	
Develop it further in a project for maturing innovation (e.g.; FET Launchpad Call, Tetracom calls, etc.)	
Integrate it in Extreme Scale Demonstrators	
Integrate in participant product roadmap	
Creating a start-up with this IP	
No idea yet	
Any other ideas? (free text below)	
What potential do you see for your IP? Do not censor yourself – express your "wildest dreams", such as for example: "My software is licensed by Y to increase the efficiency of its HPC cloud offering", or "My solver is becoming the basis for exascale application in the Z field"	
NEXT - SOFTWARE IP 4	
Is the software open source (yes/no)? If so, what license does it use?	
What is the domain of this software?	
Firmware	
OS	
Runtime	
System software	
System middleware	
Compiler	
Application middleware	
Application	
Other (please specify below)	
Short description of the software (e.g. core functionality, target usage scenario):	
What is the maturity level of the software (e.g. work in progress, stable, release)?	
Is the software available on a public repository? Please provide a link.	

If the answer to the previous question is “no”, who is the main contact if someone is interested in exploring the use of this software? (please provide a name)	
Main contact	
Another person	
Is there a user base for the software? If so, is this an emerging (e.g. originating from the project) or an established user base?	
What do you plan to do in the future with this software?	
Develop it further within another R&D project (e.g., other FETHPC project)	
Develop it further in a project for maturing innovation (e.g.; FET Launchpad Call, Tetracom calls, etc.)	
Integrate it in Extreme Scale Demonstrators	
Integrate in participant product roadmap	
Creating a start-up with this IP	
No idea yet	
Any other ideas? (free text below)	
What potential do you see for your IP? Do not censor yourself – express your “wildest dreams”, such as for example: ”My software is licensed by Y to increase the efficiency of its HPC cloud offering”, or “My solver is becoming the basis for exascale application in the Z field”	
NEXT - SOFTWARE IP 5	
Is the software open source (yes/no)? If so, what license does it use?	
What is the domain of this software?	
Firmware	
OS	
Runtime	
System software	
System middleware	
Compiler	
Application middleware	
Application	
Other (please specify below)	

Short description of the software (e.g. core functionality, target usage scenario):	
What is the maturity level of the software (e.g. work in progress, stable, release)?	
Is the software available on a public repository? Please provide a link.	
If the answer to the previous question is "no", who is the main contact if someone is interested in exploring the use of this software? (please provide a name)	
Main contact	
Another person	
Is there a user base for the software? If so, is this an emerging (e.g. originating from the project) or an established user base?	
What do you plan to do in the future with this software?	
Develop it further within another R&D project (e.g., other FETHPC project)	
Develop it further in a project for maturing innovation (e.g.; FET Launchpad Call, Tetracom calls, etc.)	
Integrate it in Extreme Scale Demonstrators	
Integrate in participant product roadmap	
Creating a start-up with this IP	
No idea yet	
Any other ideas? (free text below)	
What potential do you see for your IP? Do not censor yourself – express your "wildest dreams", such as for example: "My software is licensed by Y to increase the efficiency of its HPC cloud offering", or "My solver is becoming the basis for exascale application in the Z field"	

2.3 Other outcomes	
a. API	
b. Benchmark suites	
c. Application optimization	
d. Demonstrator	
e. Training	

f. Report	
a. API:	
Short description of the API (e.g. core functionality, target usage scenario)	
What is the level of maturity of the API (e.g. work in progress, stable, release)?	
Has the API description been published?	
Who is the main contact if someone is interested in this API? (please provide a name)	
Main contact	
Another person	
b. Benchmark suite:	
What is the level of maturity of the benchmark(s) (e.g. work in progress, stable, release)?	
Are there sample benchmarks results that are publically available?	
Are the benchmarks available on a public repository? Please provide a link.	
Who is the main contact if someone is interested in this benchmark? (please provide a name)	
Main contact	
Another person	
c. Application optimisations:	
What applications are they?	
Please describe briefly any improvements that were achieved.	
Have you developed a methodology from this experience?	
If the answer to the previous question is “yes”, have you published this methodology?	
Who is the main contact if someone is interested in this methodology? (please provide a name)	

Main contact	
Another person	
d. Demonstrator:	
Short description of the demonstrator:	
Is this demonstrator open to external access (i.e. people not in the project)?	
Who is the main contact if someone is interested in this demonstrator? (please provide a name)	
Main contact	
Another person	
e. Training:	
Short description of the training:	
Is the training material reusable?	
Who is the main contact if someone is interested in this training material? (please provide a name)	
Main contact	
Another person	
f. Report:	
Short description of the report:	
Is the report in open access?	
Cooperation:	
Have you established cooperation that will last after the project?	
Who is this collaboration with?	
Project partners	
Other FET HPC projects	
Other European projects	
Industry	
International outside Europe	
Other (please specify)	
Short description of the cooperation (organisations or people involved, topics)	
Other results:	

If you there are results that have not been mentioned in the previous categories, please describe them:

3. Best way to learn about project results					
Have you a report, a presentation or other means that you would recommend to people interested in your project (in the case of open access and web access please provide the url)?					
If you plan to issue some additional material aiming to explain the outcomes of the project, what is the best way for EXDCI2 to be informed of its release, when it will be available?					
Check your web site		Ask you directly by Send you an email (if yes when)		You will inform us via e-mail	
If so, please details how to have access to reports or presentations or who are the contact persons?					

4. Links with nanoelectronics and photonics
In your project, have you integrated innovation coming from nanoelectronics of photonics fields?
Have you been in contact during your project with organisations researching in the fields of nanoelectronics or photonics?
Have you identified a need for future advanced nanoelectronic or photonic research?
Short description of the technologies and/or relationship/need?

5. Actions related to standard (standard organisation or de facto standard)
Do you believe some of the project results could contribute to HPC standard (most of the HPC standard are de facto standard)?
Short description of the potential standard

Are any of the participants of the project experts in standard organisations? If yes, who and which standard organisations?
Have standards been a barrier to innovation coming out of your project?
Do you see anything related to standard (as presence of European experts in standard organisation, funding of expert travel to attend standard organisation, etc.) that could have helped your project?

6. Communication
EXDCI2 will implement collective actions to promote FET projects and results at SC, ISC and others events.
Are you interested in EXDCI2's promoting your results and projects?
Have you any suggestions for promotion material and/or actions that you would like EXDCI2 to implement?
Would you be interested in participating directly in the EXDCI2 booths (at SC and ISC) even after the end of your project?
Would you like to participate in a common presentation of the FET HPC 2015 project results?

7. Other comments
Is there any action that you would like EXDCI2 to undertake to help you?
How would you want to be informed about the results of this questionnaire?

Thank you!

8.3 Annex 3: Analysis of the impact of Horizon2020 on the HPC value chain

Introduction

This document is issued by the EXDCI-2 project to answer a request from the EuroHPC Governing Board to perform an analysis of the value chain in the HPC ecosystem that we have been supporting in Horizon 2020 and to issue a short document with the findings of this analysis.

The method adopted is to consider four main domains

- HPC system hardware,
- HPC system and application development software,
- HPC applications
- Transversal aspects

of the complete HPC value chain. For each of them, a high level view of the different components of the value chain is presented and then an analysis of how Horizon 2020 has supported the development of this value chain in Europe.

Due to the complexity of the HPC value chain, to the time given to perform the analysis and to the shortness of the report, this document does not pretend to capture the complete picture but should be seen as a first contribution[27].

HPC system hardware (including operation of HPC systems)

Presentation of the value chain

The production and operation of HPC systems involve a complex value chain that can be decomposed as follows:

- Foundry of chips: HPC systems mainly use chips manufactured with the most advanced CMOS technology, because of energy efficiency and due to the cost to assemble too many medium/low performance chips. The main players are non-European (TSMC, Intel, Samsung, SK hynix, Micron).
- Design of HPC chips:
    - Compute:
        - CPU: the dominant players are Intel and AMD with smaller players such as Fujitsu, IBM, Sugon or Marvell. In Europe, the EPI plans to provide a processor based on ARM and Tachyum is also designing a chip aiming to this market.
        - GPU, generic accelerator: the main actors are NVIDIA, AMD and Intel. EPI plans to develop an accelerator based on the RISC-V instruction set.
        - Specific accelerators and FPGA (Field Programmable Gate Arrays): Intel and Xilinx are the dominant FPGA players; the specific accelerator field is a booming area with ca. 50 companies proposing or preparing new chips. Some are European such as Kalray or Graphcore.
    - Interconnect: the main actors are Mellanox (NVIDIA) and Cray (now HPE). Some European companies have developed their own solutions such as Atos with BXI and Extoll.

[27] Additional information can be found in the HPC cPPP monitoring reports issued by ETP4HPC.
https://www.etp4hpc.eu/cppp.html
https://www.etp4hpc.eu/pujades/files/2017_PMR_HPC_cPPP%20Final_web.pdf
https://www.etp4hpc.eu/pujades/files/2018_PMR_HPC_cPPP-final_web.pdf

    - Memory: the domain has grown beside DDR memory to propose either very high bandwidth memory or non-volatile memory. The main players are Samsung, SK hynix, Micron and Intel.
- Integration of chips: this new layer in the value chain integrates different chips (eventually from different technologies) into a single package. The leaders are TSMC, Samsung and Intel. Europe has competences in this domain at research institute level (IMEC, CEA, Fraunhofer).
- Integration at board level and server level: even if the manufacturing is done mainly in Asia (India, China and Taiwan), Europe has several companies mastering this step either for HPC systems (Atos, E4, 2CRSI) or for other markets (Kontron, Topic…). There are also in Europe some companies proposing FPGA based boards (Pro Design, PLDA, Reflex…)
- Integration of HPC systems: integrating several thousand nodes in an efficient HPC system is an industrial know-how that today includes the cooling of the system. For very large systems, Atos is the only European provider in competition with HPE (including now former SGI and Cray companies), IBM, Dell, Lenovo, Fujitsu, Sugon, Inspur, Huawei. Europe has also other smaller players such as Megware or E4, and more recently C2RSI, which are active for midsize HPC systems (sometimes placing a system in the Top500 – whereas Atos has had ca. 20 for several years). Some cooling technology providers have also emerged in Europe such as Asetek, Iceotope or Submer. Optical cables, most of the time, are supplied by Mellanox (NVIDIA) or Finisar.
- Storage system: beside the computing part, the storage system is an important part of the HPC installation. The leaders are HPE, IBM and DDN. Some research is done in Europe at Seagate which is also one of the main disk suppliers with Hitachi. The long-term storage (tape system) providers are US-based- such as StorageTek (Oracle) and IBM.
- Data centre installation and containers: data centres are another component of the value chain with main technologies being power systems and cooling systems. Europe has some providers in this value chain as Saiver or Schneider Electric. Installation of containers is an alternative to more conventional data centres. In Europe Atos and Iceotope have developed container solutions.
- The operation of HPC systems is most of the case performed by the computing centre but can also be proposed by the system providers. Europe has built a very strong ecosystem of computing centres with top level expertise (EPCC, STFC, CEA, Cines, IDRIS, BSC, Cineca, FZJ, LRZ, HLRS, CSC, CSCS, SurfSara, PSNC…). HPC on demand is an alternative to centralised on premise operation. The three main cloud providers AWS, Microsoft and Google have all developed solutions trying to attract HPC users with cluster as a service or server instances with GPU or FPGA. In Europe, OVHcloud is present as well as specific providers such as Arctur, Gompute, T-Systems or Atos Extreme Factory. European computing centres have also developed “as a service” solutions (EPCC, STFC, Cesga, HLRS…).

The economic weight of some parts of this value chain can be assessed using Hyperion market analysis. The global market size in 2018 was 12.5 B€ for HPC systems, 5 B€ for storage systems and 2 B€ for services (mainly maintenance). From the size of the HPC systems market we can deduce that the market for compute chips is at least 5B€, and the interconnect and memory chips for HPC are over 1B€ each. Using standard industry ratio, we can make the hypothesis that the foundry part is in the range of 2-3 B€.

This value chain is not static but evolves over time. The layers that will either emerge or take more importance in the mid-term are accelerator chips, package level integration, photonics interconnects and HPC as a service. In a more long-term perspective, we can expect some new materials to be used

for compute, communication or storage. These new layers might rather complement the silicon centric value chain than replacing it.

Impact of Horizon 2020 actions on this value chain

The Horizon 2020 program has supported several elements of this value chain:

- The EPI project aims to develop in Europe a CPU and accelerator solution. The effort could reposition Europe in a layer that was and still is dominated by US technologies.
- At the basic technology level, the Exanode project has demonstrated a chip integration solution which could be further industrialized.
- For the interconnect layer, several projects have addressed the topic (Exanest, EuroExa, Mango) and demonstrated some solutions at prototype level. The most mature solutions remain Atos and Extoll ones which have been developed mainly outside Horizon 2020 support.
- Some cooling technologies have been developed under Horizon 2020 leading to industrialization through SMEs such as Iceotope or Submer.
- The set of Montblanc projects have contributed to the emergence of ARM processor HPC systems.
- In the field of FPGAs, several interesting developments can be mentioned. Some have been brought to the market by SMEs such as Maxeler.
- The integration of a more efficient storage hierarchy in HPC systems is also a domain where we can see a contribution of Horizon 2020 projects.

Besides the HPC oriented calls, the electronics and photonics call have a potential to impact the HPC value chain. EXDCI-2 works on bridging the technologies developed in these fields with the HPC requirements for more efficient solutions in compute, communication or storage.

As mentioned before, the HPC system value chain is a dynamic one. Developing a European based value chain could be easier in some emerging domains.

HPC system and application development software

Presentation of the value chain

The value chain can be explained by the description of two complementary software stacks that interact together to enable the run of applications on HPC systems.

The first one is the HPC system management stacks which is usually best described bottom-up:

- Operating system: HPC is a Linux only domain but with different variants being used. Some HPC system providers have their own optimized kernel for very large systems. But for most HPC systems, RedHat (IBM) and SuSE are the dominant options. Virtualization software is not used on large HPC systems but can be present in HPC as a service offers.
- Server management: this layer is most of the time proprietary software supplied by the HPC system vendor.
- Cluster management: the management of the nodes and the monitoring of the operation are done by this piece of software. It can be based on open source software, provided by HPC system vendors or by independent software companies such as Bright Computing or ParTec.
- Interconnect management: the software managing the interconnect can be provided by Mellanox (NVIDIA), the HPC system vendors (e.g. Atos, HPE) or independent software companies such as FabriScale (Norway)
- Resource management and scheduling: the most used resource manager is an open source one: SLURM. Other options can be supplied by software companies such as Altair.

- Workflow management: this piece of software aims to efficiently support applications that have complex workflow connected to external resources (sensors, instruments, data streams…). No leader has emerged yet for this field.
- File system: the most common solutions are either Lustre (community effort) or Spectrum Scale (GPFS from IBM). The European BeeGFS is another option.
- Data management: for managing the hot tiers (near the CPU) of the storage hierarchy, new software solutions are emerging; for the management of cold tiers the landscape is more established with solutions from either storage providers, computing centres or software companies as Grau Data.
- Maintenance tools: this layer is most of the time provided by the system providers to optimize the maintenance and detect potential failure in advance.

The complementary application development stack is usually described top-down:

- Domain specific language (DSL) or application packages: some domains develop high level abstractions relevant to express the application problems and that have attached semantic allowing the automatic generation of codes. In other domains like data analytics or deep neural networks, software frameworks provide an interface to specify an application and to link to data and then transform these specifications into a program using the underlying layers.
- Language and compilers: most common HPC languages are still Fortran and C++. For programming complex workflows some more recent languages such as Python have been used. Julia has also emerged as an alternative for some application fields. For the compilers, community efforts like GNU or LLVM do exist. Nevertheless, as performance is the key element for HPC, some players have invested in their own solutions such as Intel, NVIDIA or Cray (HPE). In Europe, NAG provides its own Fortran compiler.
- Scientific libraries: scientific libraries are a useful tool to reduce application development time and to port on different architectures with good performance level. A lot of these libraries come from a community effort. Again, as performance is so important to HPC, some players have developed their own solutions. The most notable ones are Intel and Cray libraries. In Europe, NAG is also active in this domain.
- Parallel programming models: the dominant parallel models are MPI (Message Passing Interface), OpenMP (Open Multi-Processing) and PGAS (Partitioned Global Address Space). For these programming models, you have a community effort to propose the relevant software stack support. For PGAS, we can mention the GASPI option promoted by Fraunhofer and its GPI implementation.
- Communication libraries: to support parallel programming model, communication libraries are provided mainly by HPC system vendors to efficiently use the features of their interconnect networks.
- Run-times: this piece of software makes decision during the execution of the application to optimise the time to solution (or other criteria such as the energy consumption). European research centres have developed some solutions as OmpSs (BSC), StarPU (Inria) or MPC (CEA)
- Debugger: HPC software with high parallelism needs powerful tools to help the developers. Besides basic open source approaches like GDB (GNU), the two leaders are commercial software TotalView (RogueWave) and DDT (ARM, formerly Allinea).
- Performance monitoring tools: performance tuning is a key aspect in HPC. Most of the software come from community effort. European teams have developed several top-level tools (Vampir, Scalasca, Macao, Paraver…).
- Visualization tools: exploitation of HPC application results requires most of the time powerful visualization tools. Some of these tools can also be used to interact with the application

(parameter change, mesh refinement…) introducing a new “in situ” or even “steering” dimension in HPC applications.

As seen along the previous descriptions, the value chain actors for these two software stacks are either publicly funded research teams or software developers from HPC vendors - mainly at processor (Intel, NVIDIA, ARM…) or systems level (HPE, Fujitsu, Atos…) - or else some independent software companies (as European Bright Computing). Since part of this software is delivered bundled with HPC systems, the market annual value of 1.4 B€ reported by Hyperion is an underestimated figure of the business related to this domain. In addition to this commercial value, you have also to consider the community effort. Even if precise estimates are not available, it is clear that HPC software stacks mobilize tens of thousands of full time high-level developers.

As for the HPC system value chain, we have a dynamic value chain that changes to adapt itself to the current major evolutions:

- Fast integration of AI: this integration goes in two directions, adaptation of the stacks to support AI applications, and integration of AI techniques into the stacks to increase HPC or simulation efficiency.
- Integration of HPC into the digital continuum: HPC systems will serve applications ranging from edge to HPC centres. This evolution drives new features for the software stacks such as complex end-to-end workflow management, security, integration with cloud capability or interactive HPC.
- Adaptation to HPC system architecture: here the main current drivers are new data storage hierarchy and integration of heterogeneous computing options.
- Sustainability: developing new features to reduce the environmental impacts of HPC; one of the axes is to balance energy consumption and performance; another one is to extend the life time of hardware resources.

Impact of Horizon 2020 actions on this value chain

The HPC software value chain has been addressed mainly by the FET-HPC projects. The contribution of these projects is very diverse as the analysis of the results of the FET-HPC 2014 has showed (171 results has been listed for these 19 projects ranging from demonstrators to application optimizations). The two FET-HPC2016 projects will also deliver demonstrators with an integrated software stack. The eleven FET-HPC2017 projects selected with a bottom up approach will also contribute in several key domains and the four ICT-11 projects will bridge HPC and big data technologies.

The domains in which we have seen or are expecting to see the more significant results from the ongoing projects are:

- A software stack for ARM-based processor HPC system
- New software for efficient use of the new levels of the storage hierarchy
- New IO (input/output) API (Application Programming Interface) for object storage
- New tools for FPGA resource management and programming
- Monitoring tools or auto-tuning tools with features to optimize the energy consumption
- Scientific libraries for some specific domains
- DSLs and the supporting stack for weather applications and machine learning applications
- Programming environments for heterogeneous systems (CPU+GPU+FPGA)
- Frameworks for uncertainty quantification
- Visualization tools for interactive HPC

It must be noted that the new technologies are developed by the projects, most of the time, up to TRL 5-6 and sometimes up to TRL-7. Nevertheless, there is a gap for production use of the innovations

proposed by the projects. As the majority of projects partners are from research organizations (3/4 of the total funding of the FET projects) whose mainspring is not industrialization, the projects have difficulty to impact the day to day use of HPC. One option to solve this issue would be to support a second type of projects with the clear objective to cover the TRL 6-9 levels gap. Another option, not exclusive to the first one, would be to have more industrial partners in the projects. Most of the SMEs funded by the projects have been effective in bringing their project results to the market, such as Maxeler, Appentra, Arctur, Synelexis or Kitware.

HPC applications

Presentation of the value chain

This part of the value chain is of utmost importance especially in Europe as 1) the scientific and industrial applications are THE meaning of HPC, 2) Europe owns the development of many applications used worldwide, for example in chemistry/materials more than 50% of the applications used worldwide are coming from Europe and 3) Europe is one of the biggest generator of data in the world thanks to simulations and large-scale instruments like CERN.

In the scientific domains, each community has developed its own value chain that connect different kinds of researchers (data producers, experimental scientists, expert in mathematics and algorithmics, application code developers, HPC experts…). These communities have domain specific challenges that are linked to different HPC usage challenges but all fields of research are facing a bottleneck to access to programming and development expertise which is a major obstacle to addressing the scientific challenges. Some of the major issues are presented below:

- As stated into the last edition of the PRACE Scientific Case[28], simulations are critical in **Climate, Weather, and Earth Sciences**. Exascale resources will enable sub-kilometer resolution instead of 10km, better mathematical models, and ensembles of simulations for uncertainty quantification. This will extend the forecasting ability, it will enable researchers to include soil and ocean effects into global models, and by quantifying the accuracy, climate simulations will increasingly be able to predict the efficiency of climate actions. Next-generation weather forecasting will provide local prognoses and longer lead times, with substantial financial impact. Also, researchers will increasingly be able to understand, image, and predict where earthquakes will happen and what their intensity will be.
- In **Life Sciences & Medicine**, bioinformatics will have tremendous impact for personalised (high precision) medicine. Researchers are already able to rapidly identify genetic disease variants, and it will become possible to identify diseases that are caused by combinations of variants, with treatments tailored both to the patient and state of the disease. Structural biology will increasingly rely on computational tools, allowing researchers to predict how the flexibility and motion of molecules influence function and disease. AI and deep learning techniques will provide more specific diagnosis and treatment plans than human doctors, making medical imaging one of the largest future computing users.
- **Chemistry & Materials Science** will remain one of the largest users of computing, with industry increasingly relying on simulation to design, for example, catalysts, lubricants, polymers and liquid crystals. Traditional electronic structure and molecular mechanics methods are being complemented both with multi-scale models and data-driven approaches using deep learning to predict properties of materials. This will enable researchers to fulfil the grand challenge of designing and manufacturing all aspects of a new materials from scratch, which will usher in a

[28] http://www.prace-ri.eu/third-scientific-case/

new era of targeted manufacturing and so enable European industry to compete on know-how rather by having to reduce salary-related costs to make it more competitive.

In addition to the very strong research application value chain, there is also a value chain for the industrial HPC users. Most of the industrial sectors use HPC application software provided by independent software vendors (ISVs). Some of these software are supplied by the PLM (Product Life Management) leaders Dassault Systems and Siemens. These two European companies have acquired or developed HPC applications to complement their main PLM business. You have also big ISVs dedicated to HPC applications often coming from the consolidation of smaller companies. Examples of these companies that can propose a complete set of HPC applications for some industrial sectors are ANSYS, ESI Group, Altair or Schrödinger. In addition, Europe has also some small ISVs targeting specific sectors (AVL, GNS, MAGMA, CULVI and tens of others). To complete the value chain, you have in Europe a set of service companies that propose help for product design through HPC usage. They can offer specific studies as CFD (Computational Fluid Dynamics) or electromagnetic optimization for companies designing new products that have not in-house skill for these complex processes (Noesis, Open Ocean, Lauer Weiss, K-epsilon and tens of others). They are an important layer of the value chain to expend the use of HPC in the industry.

The commercial HPC application market is just the emerged part of the iceberg as a lot of work is done by research communities organized to develop their own codes. The estimated value of the commercial HPC software market is 4.2 B€ in 2018 according to Hyperion.

Impact of Horizon 2020 actions on this value chain

In the domain of HPC application the main Horizon 2020 activity has come from the Centre of Excellence (CoE) action. The main results from this action are:

- weather and climate simulation framework ready for future exascale systems
- new advances in codes for material science
- new developments for life science applications
- improved codes for carbon free energy applications
- methodology to assess the performances of codes and improve their efficiency and application of this methodology for tens of codes
- promotion of HPC research applications within industry.

Beside the CoEs, Horizon 2020 technology projects have also an impact on the application value chain. Almost all the project's consortia include application oriented partners that work to improve some codes by taking advantage of new HPC technologies developed inside the projects. The more important impacts are:

- new libraries in some domains (hyperbolic equations; fluid dynamics, machine learning)
- use of FPGA for some kernels of codes
- restructuring of code for better use of accelerators and of new memory hierarchy
- increased scalability of big data/AI oriented codes.

In total (CoE and FET projects), we have listed more than 100 codes that have been impacted by the Horizon 2020 programme.

Finally in the field of the PRACE implementation or inside the Fortissimo projects, Horizon 2020 funding have allowed to develop tailored and complementary solution for 1) raising awareness of HPC and now AI to European SMEs through the PRACE SHAPE (SME HPC Adoption Programme in Europe) initiative (training, expertise and PoC) and 2) once evangelised provided Cloud-HPC services to European SMEs through the Fortissimo marketplace. Such combined effort supported more than 100 European SMEs at date, and thus increasing their competitiveness.

Transversal aspects

To develop a sound and dynamic European HPC ecosystem other aspects are also to be considered.

- Availability of educated and trained people: HPC is a people intensive domain with various mandatory profiles to push forward the field;
- Ability of the different HPC value chain actors to work together to increase the efficiency of the whole chain;
- Ability to work with external ecosystems to share technologies and to deliver more value in the digital continuum.

Training plays an important role in building a European base of highly-skilled, knowledgeable HPC users and application developers. This will be key to EuroHPC. Efficient exploitation of exascale hardware will undoubtedly raise new challenges for application developers and users, who will contribute their requirements and experiences into a co-design cycle for exascale technology development under EuroHPC. Hence developing human expertise is of paramount importance in the ecosystem. Europe is in a strong position in that regard and must sustain investments in HPC education and training programmes in order to remain competitive and to realise the aims of EuroHPC. In the Horizon 2020 work programme training was provided mainly by:

- PRACE (general HPC topics, cross disciplinary and domain specific in collaboration with CoEs)
- Centres of Excellence (domain specific)

The following topics are presently covered and will need to be supported in future as well:

- HPC technologies
- Scientific codes, tools and methods
- Programming models and environments
- Operation and use of HPC systems

Since the training activity is always part of a project, it is difficult to estimate the overall funding in Horizon 2020 dedicated to training. PRACE provided within the PRACE-IP projects training to more than 13 000 people and dedicated approximately 20% of the project budget to training, which sums up to 10,1 M€ in Horizon 2020. Assuming 10% of the available budget is spent by CoEs for training, CoEs and PRACE provided jointly training for approximately 23 M€ in Horizon 2020.

Horizon 2020 has been successful in developing better links between the different levels of the **HPC value chain**. Most of the technology projects consortia include some application partners. As a result, the technology providers have a better understanding of the current status of applications and of what could help to have application efficiency on future HPC systems. Vice versa, the application players have a better view of the evolution of the programming environments and of the HPC system architecture. This capacity to avoid to work in silo has a positive impact of the efficiency of the European HPC value chain.

Another transversal topic that has been successfully supported by Horizon 2020 is the development of **relationships with other ecosystems**. The evolution of simulation centric HPC toward the integration of HPC in a continuum of IT infrastructure from edge to HPC is a major challenge. To tackle it and to also address critical mass issues, it is mandatory to develop the relationship of the HPC community with other ecosystems such as Big Data, AI or IoT. Here Europe can be a worldwide leader if the momentum created by Horizon 2020 is further pushed.

Concluding remark

Up to now Horizon 2020 has committed 430 M€ of funding for technology projects (211 M€), Centres of Excellence (118 M€), processor actions (92 M€) and Coordination and Support actions (CSA 10M€) of which 230 M€ has already been spent (the rest will finance the activities of still running projects).

Horizon 2020 has been successful in achieving significant impacts in the European HPC value chains with:

- Interesting technology results developed for most of them up to TRL6-7;
- Progresses in applications in some domains
- Stronger and more connected ecosystem.

Nevertheless, due to the limited volume of this investment compared to the complete value chain R&D efforts and its bottom up approach (except for the EPI investment) Horizon 2020 impact has not yet changed the worldwide position of the European players and has not touched some parts of the value chain (the industrial HPC application sector for example).

In the next years, the addition of projects targeting the TRL6-9 gap and a more programmatic approach should help to maximize the impact of the future investments.

An enhanced and sustained training effort is needed in order to fully exploit the next EuroHPC funded pre-exascale and exascale systems.

8.4 Annex 4: FET-HPC project relationship with the SRA3 milestones

Priorities of SRA3: HPC system architecture and components

	EXANODE	EXANEST	ECOSCAL	MONTBLANC3	SAGE	NEXTGENIO	MANGO	EXTRA	GRENNFLASH	INTERWINE	COMPAT	ALLSCALE	READEX	ANTAREX	ESCAPE	EXSCAPE	EXAFLOW	EXAHYPE	NLAFET
M-ARCH-1: New standard interfaces available for integrating CPUs and accelerators on nodes and to accommodate innovative unified memory and storage architectures on networks																			
M-ARCH-2 Having well balanced systems taking benefit from high bandwidth memories and NV memories																			
M-ARCH-3: Faster end-to-end communication networks (2x and 4x bandwidth in 2018 and 2021 compared to 2015 and lower latency) with energy and power used being proportional to bandwidth																			
M-ARCH-4: End to end optical communication chain including photonic switching in order to compensate network complexity growth on the larger fabric																			
M-ARCH-5: Optimised network and storage architectures available with dynamic features, QOS and virtualization capabilities																			
M-ARCH-6 System and hardware to support performant direct remote memory access that would enable new and easier ways to program parallel applications																			
M-ARCH-7: Exascale system power envelope in the 5-15 MW power envelop range																			
M-ARCH-8: Exascale system available, at 100x more performance for relevant applications compared to today's state-of-the-art PRACE Tier-0 systems.																			

	DEEPEST	EUROEXA	RECIPE	SAGE2	MAESTRO	ASPIDE	EPIGRAM-HS	EPEEC	EXA2PRO	EXAQUTE	VECMA	VESTEC	ESCAPE2
M-ARCH-1: New standard interfaces available for integrating CPUs and accelerators on nodes and to accommodate innovative unified memory and storage architectures on networks													
M-ARCH-2 Having well balanced systems taking benefit from high bandwidth memories and NV memories													
M-ARCH-3: Faster end-to-end communication networks (2x and 4x bandwidth in 2018 and 2021 compared to 2015 and lower latency) with energy and power used being proportional to bandwidth													
M-ARCH-4: End to end optical communication chain including photonic switching in order to compensate network complexity growth on the larger fabric													
M-ARCH-5: Optimised network and storage architectures available with dynamic features, QOS and virtualization capabilities													
M-ARCH-6 System and hardware to support performant direct remote memory access that would enable new and easier ways to program parallel applications													
M-ARCH-7: Exascale system power envelope in the 5-15 MW power envelop range													
M-ARCH-8: Exascale system available, at 100x more performance for relevant applications compared to today's state-of-the-art PRACE Tier-0 systems.													

Priorities of SRA3: System software and management

	EXANODE	EXANEST	ECOSCAL	MONTBLANC3	SAGE	NEXTGENIO	MANGO	EXTRA	GRENNFLASH	INTERWINE	COMPAT	ALLSCALE	READEX	ANTAREX	ESCAPE	EXSCAPE	EXAFLOW	EXAHYPE	NLAFET
M-SYS-OS-1 Memory Hierarchy-management policies and libraries for NVRAM																			
M-SYS-OS-2 OS decomposition and specialized containerisation																			
M-SYS-OS-3 HW Embedded Security integration and cross layer security support																			
M-SYS-IC-1 Efficient peer to peer and storage over fabrics support																			
M-SYS-CL-1 Initial support of mixing HPDA, AI and HPC environment																			
M-SYS-RM-1 Resource management and orchestration support for complex workflow																			
M-SYS-RM-2: data aware and Multi-criteria resource allocation integration for adaptive scheduling																			
M-SYS-RM-3 Dynamic reconfiguration scheduling support (for flexibility and resiliency) purpose																			
M-SYS-RM-4 Data aware and power efficient scheduling																			
M-SYS-VIS-1 software support for In-situ computation and visualisation																			

	EUROEXA	DEEPEST	RECIPE	SAGE2	MAESTRO	ASPIDE	EPIGRAM-HS	EPEEC	EXA2PRO	EXAQUTE	VECMA	VESTEC	ESCAPE2
M-SYS-OS-1 Memory Hierarchy-management policies and libraries for NVRAM													
M-SYS-OS-2 OS decomposition and specialized containerisation													
M-SYS-OS-3 HW Embedded Security integration and cross layer security support													
M-SYS-IC-1 Efficient peer to peer and storage over fabrics support													
M-SYS-CL-1 Initial support of mixing HPDA, AI and HPC environment													
M-SYS-RM-1 Resource management and orchestration support for complex workflow													
M-SYS-RM-2: data aware and Multi-criteria resource allocation integration for adaptive scheduling													
M-SYS-RM-3 Dynamic reconfiguration scheduling support (for flexibility and resiliency) purpose													
M-SYS-RM-4 Data aware and power efficient scheduling													
M-SYS-VIS-1 software support for In-situ computation and visualisation													

Priorities of SRA3: Programming environment

	EXANODE	EXANEST	ECOSCAL	MONTBLANC3	SAGE	NEXTGENIO	MANGO	EXTRA	GRENNFLASH	INTERWINE	COMPAT	ALLSCALE	READEX	ANTAREX	ESCAPE	EXSCAPE	EXAFLOW	EXAHYPE	NLAFET
M-PROG- 1: APIs and corresponding libraries, run-time and compiler support for auto-tuning of application performance (incl. energy use) and supporting legacy codes.																			
M-PROG- 2: High-level programming and domain specific language frameworks																			
M-PROG- 3: Non-conventional parallel programming approaches (i.e. not MPI, not OpenMP / pthread / PGAS - but targeting asynchronous models, data flow, functional programming, model based).																			
M-PROG-4: Enhanced programming model and run-time system support for dynamic environments (management & monitoring), optimisation of communication and data management, interaction with OS or VM - within application workflows.																			
M-PROG-6: Performance Analytics and Debugging tools at extreme scale, including data race condition detection tools and user-support for problem resolution.																			
M-PROG-7: Performance analytics and debugging tools co-designed to link to the application developer's original code and high-level programming environments																			

	EUROEXA	DEEPEST	RECIPE	SAGE2	MAESTRO	ASPIDE	EPIGRAM-HS	EPEEC	EXA2PRO	EXAQUTE	VECMA	VESTEC	ESCAPE2
M-PROG- 1: APIs and corresponding libraries, run-time and compiler support for auto-tuning of application performance (incl. energy use) and supporting legacy codes.			■			■	■	■	■				
M-PROG- 2: High-level programming and domain specific language frameworks						■	■			■	■	■	■
M-PROG- 3: Non-conventional parallel programming approaches (i.e. not MPI, not OpenMP / pthread / PGAS - but targeting asynchronous models, data flow, functional programming, model based).	■	■	■				■	■	■				
M-PROG-4: Enhanced programming model and run-time system support for dynamic environments (management & monitoring), optimisation of communication and data management, interaction with OS or VM - within application workflows.						■	■	■	■			■	
M-PROG-6: Performance Analytics and Debugging tools at extreme scale, including data race condition detection tools and user-support for problem resolution.						■							
M-PROG-7: Performance analytics and debugging tools co-designed to link to the application developer's original code and high-level programming environments													

Priorities of SRA3: Energy and resiliency

	EXANODE	EXANEST	ECOSCAL	MONTBLANC3	SAGE	NEXTGENIO	MANGO	EXTRA	GRENNFLASH	INTERWINE	COMPAT	ALLSCALE	READEX	ANTAREX	ESCAPE	EXSCAPE	EXAFLOW	EXAHYPE	NLAFET
M-ENR-MS-1: Characterisation of computational advance as function of energy/power metric and standardisation of this approach with automatic and semi-automatic tools																			
M-ENR-MS-2: Methods to manage computational advance based on the pre-set energy/power metric and achieve Proportionality Computing with respect to the selected metric																			
M-ENR-MS-3: Throughput efficiency increase by scheduling instructions to the cores and functional units in the processor within its power envelope and taking the time criticality of the instructions into account																			
M-ENR-HR-4: Optimisation of the energy spend by the facility by controlling the coolant temperature down to the device level and taking the infrastructure energy cost into account																			
M-ENR-FT-5: Collection and Analysis of data from sensor networks - the Big Data challenge for measurements around the facility																			
M-ENR-FT-6: Prediction of failures and fault prediction algorithms																			
M-ENR-FT-7: Application recovery from fault conditions in the system																			
M-ENR-AR-8: Energy/Power efficient numerical libraries																			
M-ENR-MS-9: Highly efficient HPC installation																			

	EUROEXA	DEEPEST	RECIPE	SAGE2	MAESTRO	ASPIDE	EPIGRAM-HS	EPEEC	EXA2PRO	EXAQUTE	VECMA	VESTEC	ESCAPE2
M-ENR-MS-1: Characterisation of computational advance as function of energy/power metric and standardisation of this approach with automatic and semi-automatic tools			■										■
M-ENR-MS-2: Methods to manage computational advance based on the pre-set energy/power metric and achieve Proportionality Computing with respect to the selected metric			■										
M-ENR-MS-3: Throughput efficiency increase by scheduling instructions to the cores and functional units in the processor within its power envelope and taking the time criticality of the instructions into account													
M-ENR-HR-4: Optimisation of the energy spend by the facility by controlling the coolant temperature down to the device level and taking the infrastructure energy cost into account													
M-ENR-FT-5: Collection and Analysis of data from sensor networks - the Big Data challenge for measurements around the facility													
M-ENR-FT-6: Prediction of failures and fault prediction algorithms	■		■										■
M-ENR-FT-7: Application recovery from fault conditions in the system	■		■						■				
M-ENR-AR-8: Energy/Power efficient numerical libraries							■	■	■				
M-ENR-MS-9: Highly efficient HPC installation													

Priorities of SRA3: Balance compute, IO ad storage performance

	EXANODE	EXANEST	ECOSCAL	MONTBLANC3	SAGE	NEXTGENIO	MANGO	EXTRA	GRENNFLASH	INTERWINE	COMPAT	ALLSCALE	READEX	ANTAREX	ESCAPE	EXSCAPE	EXAFLOW	EXAHYPE	NLAFET
M-BIO-1: One or more storage class memory technology usages demonstrated as part of the persistent storage hierarchy																			
M-BIO-2: Extreme scale storage and I/O system simulation framework established.																			
M-BIO-3: Standardised Extreme scale I/O middleware API available: incorporating advanced features such as data layouts on NVRAM/Flash/Disk, in-storage computing, Object stores, etc, and also portability concerns raised by the CoEs.																			
M-BIO-4: Big Data analytics tools developed and optimised for Storage and I/O.																			
M-BIO-5: In-storage compute capability across all tiers/layers of the storage system as indicated by the data requirements within the CoEs.																			
M-BIO-6: I/O Quality-of-Service capability available for extreme scale storage systems																			
M-BIO-7: Extreme scale multi-tier data management tools available																			
M-BIO-8: Completion of co-Design with new use cases identified by the CoEs (AI/Deep learning, etc.)																			

	EUROEXA	DEEPEST	RECIPE	SAGE2	MAESTRO	ASPIDE	EPIGRAM-HS	EPEEC	EXA2PRO	EXAQUTE	VECMA	VESTEC	ESCAPE2
M-BIO-1: One or more storage class memory technology usages demonstrated as part of the persistent storage hierarchy													
M-BIO-2: Extreme scale storage and I/O system simulation framework established.													
M-BIO-3: Standardised Extreme scale I/O middleware API available: incorporating advanced features such as data layouts on NVRAM/Flash/Disk, in-storage computing, Object stores, etc, and also portability concerns raised by the CoEs.													
M-BIO-4: Big Data analytics tools developed and optimised for Storage and I/O.													
M-BIO-5: In-storage compute capability across all tiers/layers of the storage system as indicated by the data requirements within the CoEs.													
M-BIO-6: I/O Quality-of-Service capability available for extreme scale storage systems													
M-BIO-7: Extreme scale multi-tier data management tools available													
M-BIO-8: Completion of co-Design with new use cases identified by the CoEs (AI/Deep learning, etc.)													

Priorities of SRA3: Big data and HPC usage models

	EXANODE	EXANEST	ECOSCAL	MONTBLANC3	SAGE	NEXTGENIO	MANGO	EXTRA	GRENNFLASH	INTERWINE	COMPAT	ALLSCALE	READEX	ANTAREX	ESCAPE	EXSCAPE	EXAFLOW	EXAHYPE	NLAFET
M-BDUM-METRICS-1: Data movement aware performance metrics available																			
M-BDUM-METRICS-2: HPC-like performance metrics for Big Data systems available.																			
M-BDUM-METRICS-3: HPC-Big Data combined performance metrics available																			
M-BDUM-MEM-1: Holistic HPC-Big Data memory models available																			
M-BDUM-MEM-2: NVM-HPC memory and Big Data coherence protocols and APIs available.																			
M-BDUM-ALGS-1: Berkeley Dwarfs determination for Big Data applications available																			
M-BDUM-ALGS-2: Dwarfs in Big Data platforms implemented																			
M-BDUM-PROG-1: Heterogeneous programming paradigms for HPC-Big Data available																			
M-BDUM-PROG-2: Heterogeneous programming paradigm with coherent memory and compute unified with Big Data programming environments available																			
M-BDUM-PROG-3: Single programming paradigm across a hybrid HPC-Big Data system available																			
M-BDUM-VIRT-1: Elastic HPC deployment implemented																			
M-BDUM-VIRT-2: Full virtualisation of HPC usage implemented																			
M-BDUM-DIFFUSIVE-1: Big Data - HPC hybrid prototype available																			
M-BDUM-DIFFUSIVE-2: Big Data - HPC large-scale demonstrator integrated																			

	EUROEXA	DEEPEST	RECIPE	SAGE2	MAESTRO	ASPIDE	EPIGRAM-HS	EPEEC	EXA2PRO	EXAQUTE	VECMA	VESTEC	ESCAPE2
M-BDUM-METRICS-1: Data movement aware performance metrics available													
M-BDUM-METRICS-2: HPC-like performance metrics for Big Data systems available.													
M-BDUM-METRICS-3: HPC-Big Data combined performance metrics available													
M-BDUM-MEM-1: Holistic HPC-Big Data memory models available													
M-BDUM-MEM-2: NVM-HPC memory and Big Data coherence protocols and APIs available.													
M-BDUM-ALGS-1: Berkeley Dwarfs determination for Big Data applications available													
M-BDUM-ALGS-2: Dwarfs in Big Data platforms implemented													
M-BDUM-PROG-1: Heterogeneous programming paradigms for HPC-Big Data available													
M-BDUM-PROG-2: Heterogeneous programming paradigm with coherent memory and compute unified with Big Data programming environments available													
M-BDUM-PROG-3: Single programming paradigm across a hybrid HPC-Big Data system available													
M-BDUM-VIRT-1: Elastic HPC deployment implemented													
M-BDUM-VIRT-2: Full virtualisation of HPC usage implemented													
M-BDUM-DIFFUSIVE-1: Big Data - HPC hybrid prototype available													
M-BDUM-DIFFUSIVE-2: Big Data - HPC large-scale demonstrator integrated													

Priorities of SRA3: Mathematics and algorithms for extreme scale HPC systems

	EXANODE	EXANEST	ECOSCAL	MONTBLANC3	SAGE	NEXTGENIO	MANGO	EXTRA	GRENNFLASH	INTERWINE	COMPAT	ALLSCALE	READEX	ANTAREX	ESCAPE	EXSCAPE	EXAFLOW	EXAHYPE	NLAFET
M-ALG-1: Scalability of algorithms demonstrated for forward in time computing and 3-dimensional FFT for current architectures																			
M-ALG-2: Multiple relevant use cases demonstrated for improving performance by means of robust, inexact algorithms with reduced communication costs																			
M-ALG-3: Scalable algorithms demonstrated for relevant data analytics and artificial intelligence methods.																			
M-ALG-4: Processes established for co-design of mathematical methods for data analytics and of HPC technologies/architectures																			
M-ALG-5: Classes of data, partitioning and scheduling problems categorised and their complexity ascertained																			
M-ALG-6: Mathematical and algorithmic approaches established for the scheduling of tasks on abstract resources and exploitation of multiple memory levels																			
M-ALG-7: Research on mathematical methods and algorithms exploited for compiler technologies, runtime environments, resource schedulers and related tools.																			
M-ALG-8: Reduction of energy-to-solution demonstrated by means of appropriately optimized algorithms demonstrated for a set of relevant use cases.																			
M-ALG-9: Process for vertical integration of algorithms established together with the validation of scalability, ease of implementation, tuning and optimisation																			
M-ALG-10: Tuning of algorithmic parameters at exascale completed for a relevant set of algorithms.																			

	EUROEXA	DEEPEST	RECIPE	SAGE2	MAESTRO	ASPIDE	EPIGRAM-HS	EPEEC	EXA2PRO	EXAQUTE	VECMA	VESTEC	ESCAPE2
M-ALG-1: Scalability of algorithms demonstrated for forward in time computing and 3-dimensional FFT for current architectures													
M-ALG-2: Multiple relevant use cases demonstrated for improving performance by means of robust, inexact algorithms with reduced communication costs													
M-ALG-3: Scalable algorithms demonstrated for relevant data analytics and artificial intelligence methods.													
M-ALG-4: Processes established for co-design of mathematical methods for data analytics and of HPC technologies/architectures													
M-ALG-5: Classes of data, partitioning and scheduling problems categorised and their complexity ascertained													
M-ALG-6: Mathematical and algorithmic approaches established for the scheduling of tasks on abstract resources and exploitation of multiple memory levels													
M-ALG-7: Research on mathematical methods and algorithms exploited for compiler technologies, runtime environments, resource schedulers and related tools.													
M-ALG-8: Reduction of energy-to-solution demonstrated by means of appropriately optimized algorithms demonstrated for a set of relevant use cases.													
M-ALG-9: Process for vertical integration of algorithms established together with the validation of scalability, ease of implementation, tuning and optimisation													
M-ALG-10: Tuning of algorithmic parameters at exascale completed for a relevant set of algorithms.													

8.5 Annex 5: Codes used by FET-HPC projects and CoEs

project	type	code name
exanode	FET-HPC-2014	Abinit
E-CAM	CoE	AiiDA
MAX	CoE	AiiDA
ESCAPE	FET-HPC-2014	Aladin
ECoE	CoE	Alya
CompBioMed	CoE	Alya
Excellerat	CoE	Alya
montblanc3	FET-HPC-2014	Alya
Readex	FET-HPC-2014	Alya
EuroEXA	FET-HPC-2016	Alya
excape	FET-HPC-2014	ambitcli
allscale	FET-HPC-2014	AMDADOS
Readex	FET-HPC-2014	AMG203
Hidalgo	CoE	AMOS
COMPAT	FET-HPC-2014	Amuse
Hidalgo	CoE	Apache Flink
Hidalgo	CoE	Apache Spark
Hidalgo	CoE	Apache Storm
Cheese	CoE	ASHEE
Excellerat	CoE	AVBP
EPEEC	FET-HPC-2017	AVBP
DEEP-EST	FET-HPC-2016	AVBP
EuroEXA	FET-HPC-2016	AVU-GSR
CompBioMed	CoE	BAC
COMPAT	FET-HPC-2014	BAC
VECMA	FET-HPC-2017	BAC
intertwine	FET-HPC-2014	BAR
VERTEX	FET-HPC-2017	BATS-R-US
MAX	CoE	BigDFT
excape	FET-HPC-2014	binet
RECIPE	FET-HPC-2017	biomedecine
SAGE2	FET-HPC-2017	Bout++
exanode	FET-HPC-2014	BQCD
NLAFET	FET-HPC-2014	BSD
RECIPE	FET-HPC-2017	BSIT
DEEP-EST	FET-HPC-2016	CERN code
SAGE2	FET-HPC-2017	Cervical Cancer Screening
VECMA	FET-HPC-2017	CGMD
EXA2PRO	FET-HPC-2017	CO2 capture
NLAFET	FET-HPC-2014	Code Saturne
DEEP-EST	FET-HPC-2016	CoreNeuron
ESCAPE	FET-HPC-2014	COSMO
Hidalgo	CoE	COVISE
E-CAM	CoE	CP2K

bioexcel	CoE	CP2K
MAX	CoE	CP2K
E-CAM	CoE	CPMD
bioexcel	CoE	CPMD
CompBioMed	CoE	CT2S
VECMA	FET-HPC-2017	DALES
ASPIDE	FET-HPC-2017	Deep learning
ecoscale	FET-HPC-2014	detection of moving objects
SAGE	FET-HPC-2014	diamond data
EPEEC	FET-HPC-2017	DIOGENes
E-CAM	CoE	DL_MESO_DPD
E-CAM	CoE	DL-POLY
ExaNest	FET-HPC-2014	DPSNN
ESIWACE	CoE	Dynamico
E-CAM	CoE	Electronic Structure Library Project (ESL)
Readex	FET-HPC-2014	Elmer
E-CAM	CoE	ELSI
DEEP-EST	FET-HPC-2016	EMAC
MAESTRO	FET-HPC-2017	ERA
ECoE	CoE	ESIAS-Chem
ECoE	CoE	ESIAS-Meteo
Readex	FET-HPC-2014	ESPRESO
E-CAM	CoE	ESPResSo++ .
RECIPE	FET-HPC-2017	EULAG
ECoE	CoE	EURAD-IM
Cheese	CoE	ExaHyPE
Exahype	FET-HPC-2014	ExaHyPE
excape	FET-HPC-2014	excape-db
Cheese	CoE	FALL3D
Antarex	FET-HPC-2014	FCD engine
EXaQute	FET-HPC-2017	FEMPAR
Excellerat	CoE	FEniCS
Hidalgo	CoE	FEniCS
MAX	CoE	FFTXLib
COMPAT	FET-HPC-2014	Ficsion
allscale	FET-HPC-2014	FINE/OpenTM
Hidalgo	CoE	FLEE
VECMA	FET-HPC-2017	Flee
MAX	CoE	Fleur
Readex	FET-HPC-2014	FLEXI
Excellerat	CoE	Flucs
Excellerat	CoE	Fluidity
EuroEXA	FET-HPC-2016	FRTM
ExaNest	FET-HPC-2014	GADGET
EuroEXA	FET-HPC-2016	GADGET
E-CAM	CoE	GC-AdResS
VECMA	FET-HPC-2017	GEM
Antarex	FET-HPC-2014	GEODOCK
excape	FET-HPC-2014	GFA

intertwine	FET-HPC-2014	Graph BLAS
E-CAM	CoE	GROMACS
bioexcel	CoE	GROMACS
DEEP-EST	FET-HPC-2016	GROMACS
ECoE	CoE	Gysela
Mango	FET-HPC-2014	H265 transcoder
bioexcel	CoE	HADDOCK
CompBioMed	CoE	HemeLB
COMPAT	FET-HPC-2014	HemeLB
VECMA	FET-HPC-2017	HemeLB
CompBioMed	CoE	HemoCell
VECMA	FET-HPC-2017	HemoCell
SAGE2	FET-HPC-2017	high-resolution images of a brain
ESCAPE	FET-HPC-2014	Hirlam
DEEP-EST	FET-HPC-2016	HPDBSCAN
CompBioMed	CoE	HTBAC
CompBioMed	CoE	HTMD
ASPIDE	FET-HPC-2017	Human brain Magnetic Resonance Imaging
exanode	FET-HPC-2014	HydroC
ExaNest	FET-HPC-2014	Hy-Nbody
excape	FET-HPC-2014	hyperloom
ESIWACE	CoE	ICON
ESCAPE2	FET-HPC-2017	ICON
ESIWACE	CoE	IFS
ESCAPE	FET-HPC-2014	IFS
ESCAPE2	FET-HPC-2017	IFS
EPIGRAM-HS	FET-HPC-2017	IFS
EuroEXA	FET-HPC-2016	IFS
EuroEXA	FET-HPC-2016	image classification code
Mango	FET-HPC-2014	image rendering
extra	FET-HPC-2014	image segmentation
Readex	FET-HPC-2014	Indeed
EuroEXA	FET-HPC-2016	InfOli
CompBioMed	CoE	Insigneo Bone Tissue
CompBioMed	CoE	InSilicoMRI
SAGE	FET-HPC-2014	Integrated Modelling and Analysis Suite (IMAS)
Readex	FET-HPC-2014	Intel MKL
SAGE	FET-HPC-2014	iPIC3D
allscale	FET-HPC-2014	iPIC3D
intertwine	FET-HPC-2014	iPIC3D
SAGE2	FET-HPC-2017	iPIC3D
VERTEX	FET-HPC-2017	iPIC3D
Mango	FET-HPC-2014	IPSEC
exanode	FET-HPC-2014	KKRnano
EXA2PRO	FET-HPC-2017	KKRnano
ECoE	CoE	KMC/DMC
EXaQute	FET-HPC-2017	Kratos
Readex	FET-HPC-2014	Kripke

E-CAM	CoE	LAMMPS
ExaNest	FET-HPC-2014	LAMMPS
VECMA	FET-HPC-2017	LAMMPS
MAX	CoE	LAXLib
EuroEXA	FET-HPC-2016	LBM
EuroEXA	FET-HPC-2016	LFRic
Cheese	CoE	L-HySEA
CompBioMed	CoE	Living heart Human Model
EXA2PRO	FET-HPC-2017	LQCD (high energy physics)
E-CAM	CoE	Ludwig
intertwine	FET-HPC-2014	Ludwig
Readex	FET-HPC-2014	Lulesh
EPIGRAM-HS	FET-HPC-2017	lung cancer
Hidalgo	CoE	MASON
Readex	FET-HPC-2014	MCB
VERTEX	FET-HPC-2017	Meso-NH
EXA2PRO	FET-HPC-2017	MetaWalls (energy storage)
ExaNest	FET-HPC-2014	MiniMD
Readex	FET-HPC-2014	MiniMD
Greenflash	FET-HPC-2014	mirror control
ExaNest	FET-HPC-2014	MonetDB
VERTEX	FET-HPC-2017	mosquito disaese
E-CAM	CoE	MP2C
Hidalgo	CoE	MUSCLE2
Excellerat	CoE	Nek5000
Exaflow	FET-HPC-2014	Nek5000
EPIGRAM-HS	FET-HPC-2017	NEK5000
Exaflow	FET-HPC-2014	Nektar++
ESIWACE	CoE	NEMO
ESCAPE2	FET-HPC-2017	NEMO
EuroEXA	FET-HPC-2016	NEMO
VECMA	FET-HPC-2017	NEMORB
exanode	FET-HPC-2014	NEST
DEEP-EST	FET-HPC-2016	NEST
EuroEXA	FET-HPC-2016	NEST
SAGE	FET-HPC-2014	Next-Generation Sequencing data
Exaflow	FET-HPC-2014	NS3D
MAESTRO	FET-HPC-2017	NWP
CompBioMed	CoE	OpenBF
ExaNest	FET-HPC-2014	OpenFOAM
NextgenIO	FET-HPC-2014	OpenFOAM
Readex	FET-HPC-2014	OpenFOAM
VECMA	FET-HPC-2017	openIFS
ASPIDE	FET-HPC-2017	Opinion Mining platform
E-CAM	CoE	OPS
EPEEC	FET-HPC-2017	OSIRIS
MAESTRO	FET-HPC-2017	PaDaWAn
COMPAT	FET-HPC-2014	Palabos
CompBioMed	CoE	Palabos Flow Diverter

CompBioMed	CoE	Palabos Vertebroplasty
E-CAM	CoE	PaPIM
ECoE	CoE	ParFlow
Cheese	CoE	PARODY_PDAF
EPIGRAM-HS	FET-HPC-2017	PIC3D
ExaNest	FET-HPC-2014	Pinocchio
DEEP-EST	FET-HPC-2016	piSVM
intertwine	FET-HPC-2014	PLASMA
CompBioMed	CoE	Playmolecule
bioexcel	CoE	PMX
CompBioMed	CoE	PolNet
ECoE	CoE	PVnegf
E-CAM	CoE	Q-Chem
ECoE	CoE	QMCPACK
E-CAM	CoE	QMCPACK
E-CAM	CoE	Quantics
MAX	CoE	Quantum ESPRESSO
EPEEC	FET-HPC-2017	Quantum ESPRESSO
EuroEXA	FET-HPC-2016	Quantum ESPRESSO
E-CAM	CoE	Quantum ESPRESSO
extra	FET-HPC-2014	quantum monte carlo code
ExaNest	FET-HPC-2014	RegCM
Hidalgo	CoE	RepastHPC
VECMA	FET-HPC-2017	RepastHPC
ecoscale	FET-HPC-2014	reservoir simulation
extra	FET-HPC-2014	risk model
ExaNest	FET-HPC-2014	SailFish
Cheese	CoE	Salvus
SAGE	FET-HPC-2014	satellite data
Exaflow	FET-HPC-2014	SBLI
Cheese	CoE	SeisSol
ECoE	CoE	SHEMAT-Suite
E-CAM	CoE	SIESTA
MAX	CoE	SIESTA
MAX	CoE	SIRIUS
MAESTRO	FET-HPC-2017	SIRIUS
DEEP-EST	FET-HPC-2016	SKA data processing
SAGE2	FET-HPC-2017	SKA data processing
excape	FET-HPC-2014	SMURFF
EuroEXA	FET-HPC-2016	SMURFF
EPEEC	FET-HPC-2017	SMURFF
Cheese	CoE	SPECFEM3D
Hidalgo	CoE	SUMO
Antarex	FET-HPC-2014	Sygic
intertwine	FET-HPC-2014	TAU
Hidalgo	CoE	TensorFlow
MAESTRO	FET-HPC-2017	TerrSysMP
Cheese	CoE	T-HySEA
Hidalgo	CoE	Torch

ASPIDE	FET-HPC-2017	Urban Computing (UNICAL)
CompBioMed	CoE	Virtual Assay
Hidalgo	CoE	VISTLE
CompBioMed	CoE	VisualGec
SAGE	FET-HPC-2014	visuallisation
E-CAM	CoE	Wannier90
VERTEX	FET-HPC-2017	wildfire
DEEP-EST	FET-HPC-2016	xPic
Cheese	CoE	XSHELLS
MAX	CoE	Yambo
SAGE2	FET-HPC-2017	YouTube-8M Video Understanding Challenge