

EUROPEAN HPC TECHNOLOGY RESEARCH PROJECTS

The objective of this white paper is to present the results or the on-going work of the first HPC technology projects initiated within the Horizon 2020 programme, namely within calls FET-HPC-2014, 2016 and 2017. We have also analysed the projects' impact on the application domain

The results of these projects have been compiled (which has led to a 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).

We have established that the FET-HPC projects have delivered a large set of IPs (intellectual property elements), which represents an important asset for the entire European HPC community. The impact of the work of the projects extends beyond the domain of technology onto applications and thus the work of the Centres of Excellence.

This high-level, comprehensive analysis can be used to:

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

Table of contents

The FET-HPC projects	3
FET-HPC-2014	4
Global view	4
Analysis by project.....	8
FET-HPC 2016	15
Deepest.....	15
EuroEXA	17
FET-HPC 2017	19
Global vision	19
Analysis by group of projects	20
Impact on applications	26
Main findings and conclusion.....	29
ANNEX	31

List of tables

Table 1: List of FET-HPC calls	3
Table 2: Typology of FET-HPC-2014 projects	4
Table 3: Users of the FET-HPC-2014 project results.....	5
Table 4: DEEP-EST expected results	16
Table 5: EuroEXA expected results.....	18
Table 6: Typology of FET-HPC-2017	19
Table 7: Programming environment projects	21
Table 8: Common codes of the FET-HPC-2016-2017 projects	27
Table 9: Common codes FET-HPC-2016-2017 CoEs	27

The 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-2017¹.

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 of their activities and generate a database² 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.

¹ The complete list of the projects is given in Annex 1

² <https://exdci.eu/activities/fethpc-results>

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>.

FET-HPC-2014

Global view

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	
	ARM based HPC	
	Reconfigurable systems	
	IO	
HPC stack and application-oriented projects		
	Energy efficiency	
	Programming model	
	Multiscale	
	Generic applications: Hyperbolic PDE, Machine learning, Fluid dynamics, Numerical linear algebra, Weather models	

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 ESDs³, 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.

³ 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.

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 programming, 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.

Technology 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 Exanode/ExaNest Ecoscale MontBlanc3	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

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 be 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 Marvell) 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 MPI⁴ library
- Task aware GASPI⁵ library
- Runtime managing different programming model (based on StarPU⁶)

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).

⁴ Message Passing Interface: one of the main programming models in HPC

⁵ Global Address Space Programming Interface see <http://www.gaspi.de/gaspi/>

⁶ A Unified Runtime System for Heterogeneous Multicore Architectures <http://starpu.gforge.inria.fr/>

ESCAPE

Main results

The most noticeable IPs of the project are:

- Software representing weather forecast model components
- Benchmark suite
- Different implementation of weather dwarfs⁷

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

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.

⁷ 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.

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

Similar 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.

FET-HPC 2016

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.

Deepest

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): network-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 efficiently support the use of GPU and GPU direct communication	software
Resiliency support by software stack	software
Slurm enhancement to efficiently support the MSA	software
Complete software stack (ParaStation Modulo) to efficiently 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.

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.

FET-HPC 2017

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	
	
Heterogeneous system management	
	
IO and storage	
	
Visualization, interactive HPC, urgent computing	
	
Verification, Validation, Uncertainty Quantification	
	
Algorithm: weather and climate models	
	

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.

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 be 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 BarbequeRTRM⁸ and some enhanced features in Slurm⁹. 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.

⁸ Barbeque Run-Time Resource Manager see <https://bosp.deib.polimi.it/doku.php>

⁹ Slurm is the workload manager used by many computing centres.

Some of the developments of the MAESTRO project can be applied to generic object storage systems such as CEPH¹⁰. 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.

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 if 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.

¹⁰ CEPH is a distributed object, block, and file storage platform see <https://github.com/ceph/ceph>

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 ECMWF¹¹ the coordinator and the European climate and weather organizations that are well aware of the activities of both ESCAPE-2 and ESIWACE-2.

¹¹ European Centre for Medium-Range Weather Forecasts

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 Excellence¹²) 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 Espresso, 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.

¹² 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	D	COMBIOBIE	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				RepastHP C		
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.

Main findings and conclusion

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. In summary:

- The FET-HPC projects have delivered a large set of IPs, which represents an important asset for the entire European HPC community;
- They have contributed to most of the milestones selected in the SRAs issued to define the European HPC research 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.

Some actions are 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 ECP¹³, 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 timeframe,
- 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

¹³Exascale Computing Project see <https://www.exascaleproject.org/>

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 Agenda¹⁴ (see Chapter 10 of the SRA).

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

¹⁴ This document has been supported by EXDCI-2 and is available at [https://www.etp4hpc.eu/pujades/files/ETP4HPC_SRA4_2020_web\(1\).pdf](https://www.etp4hpc.eu/pujades/files/ETP4HPC_SRA4_2020_web(1).pdf)

ANNEX

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 platform based 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 NUMERICIS EN ENGINYERIA	http://exaquete.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/