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

EXDCI-2

European eXtreme Data and Computing Initiative - 2

Grant Agreement Number: 800957

D4.1

Transversal Vision Report

Final

Version: 1.1
Author(s): François Bodin, Univ. Rennes 1
Maike Gilliot, ETP4HPC
Date: 27/09/2019

Project and Deliverable Information Sheet

EXDCI Project	**Project Ref. №: 800957**	
	Project Title: European eXtreme Data and Computing Initiative - 2	
	Project Web Site: http://www.exdci.eu	
	Deliverable ID: D4.1	
	Deliverable Nature: Report	
	Dissemination Level: PU	**Contractual Date of Delivery:** 28/02/2019
		Actual Date of Delivery: 28 / 02 / 2019 (revised version uploaded 30 / 09 / 2019)
	EC Project Officer: Evangelos Floros	

* - 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: Transversal Vision Report**	
	ID: D4.1	
	Version: 1.1	**Status:** Final
	Available at: http://www.exdci.eu	
	Software Tool: Microsoft Word 2013	
	File(s): D4.1-V1.1.docx	
Authorship	**Written by:**	François Bodin, Univ. Rennes 1 Maike Gilliot, ETP4HPC
	Contributors:	Mark Asch, Univ. Jules Vernes, Picardie Olivier Beaumont, Inria Thierry Bidot, Neovia Innovation Henri Calandra, TOTAL Paul Carpenter, BSC Guillaume Colin de Verdière, CEA Jean-Francois Lavignon, TS-JFL Bernd Mohr, FZ Jülich Jean-Philippe Nominé, CEA Dirk Pleiter, FZ Jülich
	Reviewed by:	Serge Bogaerts, PRACE Catherine Lambert, CERFACS
	Approved by:	MB/TB

Document Status Sheet

Version	Date	Status	Comments
0.1	15/01/2019	Draft	Outline, structure, Into
0.2	23/01/2019	Draft	Fully integrated document, with comments and open questions
0.3	25/01/2019	Draft	Updated draft version, taking comments into account
0.35	31/01/2019	Draft	Update section 3
0.4	11/02/2019	Draft	Version for internal review
0.5	19/02/19	Draft	Integration of reviewer's comments
1.0	23/02/19	Final	After proofreading and language check, final version for submission to MB/TB
1.1	27/09/19	Final	Revised by Maike Gilliot following comments in Interim Review Report.

Document Keywords

Keywords:	PRACE, , Research Infrastructure, Transversal Vision, data driven applications, Workflow, Service-oriented architecture (SOA), Accelerators

Table of Contents

References and Applicable Documents

[1] https://irds.ieee.org/images/files/pdf/2017/2017IRDS_ES.pdf

[2] https://www.fortissimo-project.eu/about/fortissimo-2

[3] http://www.prace-ri.eu/hpc-access/shape-programme/

[4] http://www.exascale.org/bdec/sites/www.exascale.org.bdec/files/whitepapers/bdec_pathways.pdf

[5] https://technode.com/2018/12/20/taiwanese-chip-maker-tsmc-to-build-the-worlds-first-3nm-chip-factory

[6] http://mondodigitale.aicanet.net/2016-5/articoli/03_MD66_Conte-Debenedictis-Gargini-Kadin-Track.pdf

[7] https://science.energy.gov/~/media/ascr/ascac/pdf/meetings/201712/ASCAC-Future-HPC-report.pdf)

[8] http://www.softmachines.org/wordpress/?p=2097

[9] https://en.wikipedia.org/wiki/List_of_open-source_hardware_projects

[10] http://raw.necst.it

[11] https://fpl2019.bsc.es/

[12] https://blog.google/products/google-cloud/google-cloud-offer-tpus-machine-learning

[13] https://aws.amazon.com/ec2/instance-types/f1/?nc1=h_ls

[14] https://www.accelize.com

[15] https://www.youtube.com/watch?time_continue=39&v=89bYjLyffUE

[16] https://www.ibm.com/blogs/research/2017/10/ibm-scientists-demonstrate-memory-computing-1-million-devices-applications-ai

[17] https://www.ibm.com/blogs/research/2018/07/phase-change-memory

[18] https://aws.amazon.com/en/ec2/instance-types/f1/

[19] https://docs.microsoft.com/fr-fr/azure/machine-learning/service/how-to-deploy-fpga-web-service

[20] http://apt.cs.manchester.ac.uk/projects/SpiNNaker/

[21] http://apt.cs.manchester.ac.uk/projects/SpiNNaker/SpiNNchip/

[22] http://lightwave.ee.columbia.edu/files/Rumley2015a.pdf

[23] https://ieeexplore.ieee.org/abstract/document/6588351

[24] https://drive.google.com/file/d/1fVYDv8kVGUn6Ugtiru6wR2szjWZqYthV/view

[25] https://www.osapublishing.org/abstract.cfm?uri=FiO-2016-FTu1D.6

[26] http://www.hpc-escape.eu/

[27] https://ieeexplore.ieee.org/document/8310328

[28] http://www.cs.columbia.edu/~simha/preprint_micro17.pdf

[29] https://phys.org/news/2018-07-closer-optical-artificial-neural-network.html

[30] https://link.springer.com/protocol/10.1007/978-1-4939-7795-6_23
[31] https://math.nist.gov/quantum/zoo/
[32] https://www.research.ibm.com/ibm-q/
[33] https://www.dwavesys.com/tutorials/background-reading-series/introduction-d-wave-quantum-hardware
[34] https://insidehpc.com/2018/05/quantum-computing-timing-everything/

List of Acronyms and Abbreviations

AAI	Authentication and Authorisation Infrastructure
AI	Artificial Intelligence
AIOTI	The Alliance of Internet of Things Innovation
BDEC	Big Data and Extreme-scale Computing
BDVA	Big Data Value Association
CDN	Content Delivery Network
CoE	Centres of Excellence for Computing Applications
CMOS	Complementary Metal Oxide Semiconductor
CSA	Coordination and Support Action
CPU	Central Processing Unit
D	Deliverable
DNN	Deep Neural Networks
EC	European Commission
ECMWF	European Centre for Medium-range Weather Forecasts
EOSC	European Open Science Cloud
FET	Future and Emerging Technologies
FPGA	Field Programmable Gate Arrays
H2020	Horizon 2020: EC Research and Innovation Programme in Europe
HPC	High Performance Computing
ISV	Independent Software Vendor
KPI	Key-Performance Indicator
ML	Machine Learning
PCM	Phase Change Memory
R&D	Research and Development
R&I	Research and Innovation
SME	Small and Medium Enterprise
SRA	Strategic Research Agenda
TPU	Tensor Processing Unit
TRL	Technology Readiness Level
US	United States
WG	Working Group
WP	Work Package

Executive Summary

This report highlights a set of challenges that are becoming more acute as the HPC technology and use is profoundly changing. The introduction of big data in the discovery process is driving application to be implemented as large-scale distributed workflows that need to be deployed on a large set of systems, each one having its own idiosyncrasies (e.g. quantum accelerator) that questions the way application are implemented.

This is the first step on the way to a shared Transversal Vision developed in the frame of EXDCI-2. The goal is to identify key challenges for the workflow-based applications: understand the requirements and understand where research and innovation is needed to preparing Europe's HPC machines, its users, its software developers and vendors. This report presents the driving factors (cf. Section 2) for this transformation and suggests ways to cope with these challenges, presented in Section 3 of this report).

As driving factors we have identified the considerable impact of data for the design of applications, secondly an opening of HPC to other user communities and for other usages, and thirdly, the advent of new, disruptive hardware.

The data challenge can be addressed by workflows, and service-oriented architectures allow to address usage evolution.

In the upcoming months, this Transversal Vision will be developed further with support from the EXDCI community and with the support of European experts in these domains.

1 Overall process for the EXDCI2 Transversal Vision

As of today, the majority of the HPC applications could roughly be described as “high precision simulation of large and complex (multi-)physics systems”. Although this type of application will remain a predominant part of the HPC workload in the future, advances in Big Data and Artificial Intelligence (AI) have extended the scope of HPC applications, with the HPC centre being just one brick in a workflow of different, heterogenous compute and storage entities for scientific computing, moving form a “HPC Compute Centre based” perspective towards a “Workflow based” perspective.

These new, workflow based HPC applications having emerged over the past years, are expected to increase drastically in the upcoming years [4]. To prepare in Europe for this workflow based workload requires to look beyond the traditional requirements such as more flops, lower latency, or higher bandwidth.

The goal of the Transversal Vision developed in the context of EXDCI-2 is to identify key challenges arising from the workflow based applications: to understand the requirements and to understand where research and innovation is needed to preparing Europe’s HPC machines, its users, its software developers and vendors. This report is the first step towards this Transversal Vision. With EXDCI contributors, we have identified the driving factors for this transformation (cf. Section 2):

a) A drastic impact of data in the design of applications
b) An opening of HPC to other users (and thus other usages)
c) Potential highly disruptive hardware

Based on this, ways to cope with these challenges are presented in Section 3:

a) Addressing the data challenge via workflows
b) Service oriented architectures to address usage evolution
c) Coping with hardware evolution

In the upcoming months, this Transversal Vision will be developed further with support from the EXDCI community at large (including the developers of applications) and with the support of (European) experts in these domains.

2 Transversal Vision: Driving forces for tomorrow's HPC landscape

In the past High Performance Computing (HPC) was synonymous with "high precision simulation of large and complex (multi-)physics systems". This vision needs now to be enlarged, as the HPC landscape is changing. This section sketches the three factors driving this change in the HPC landscape: First, the applications we face are more and more driven by "data" rather than by "compute" (presented in Section 2.1). This leads also to broader usage of HPC, as detailed in Section 2.2. And moreover, these tendencies are supported by considerable evolution of the underlying hardware, as presented in Section 2.3.

2.1 Impact on landscape from the data-driven applications

In the future, a significant part of HPC applications will be based on the convergence of simulations (i.e., "traditional HPC") and the compute-intensive processing and analysis of data (stemming from Artificial Intelligence or Big Data applications). Overall, this will lead to a large variety of applications with different characteristics, which can very roughly be categorised as follows[1]:

- **Simulation:** relatively low amount of input data, large computation requirements (mostly in high precision floating point representation) and large amount of generated data (simulation results)
- **Big Data:** large amount of external input data, medium computational requirement and low amount of output data (information extracted from the input data)
- **AI (Deep Learning in training phase):** large input (local) database with very high access rate, large computation requirements (in low precision floating point representation) and low amount of output data (weight of the Neural Networks)
- **AI (Deep Learning in inference phase):** input of medium size (depends on the application), low processing amount (reduced precision floating point or integer) and low amount of generated data.

As part of the efforts towards a shared Transversal Vision, lead uses cases will be identified, illustrating different aspects of the applications types. This will be part of the joint efforts with WP2 and WP3 within EXDCI-2.

2.2 Beyond "traditional" HPC users and usages

For many years, High-Performance Computing in the sense of "developing and using highly scalable simulation applications on specialised hardware and software infrastructure" has been a niche area of computing science for some selected few, limited to government funded institutions, and to some large companies in some specific domains, such as in engineering (e.g., automotive/aero/spatial for simulation of combustion engines, aerodynamic properties), and in the domain of natural resources (oil & gas industry in particular).

This is currently changing: over the last decade, HPC established itself more widely in science, in particular through the interest in data-driven HPC applications (AI and Data Analytics). HPC and data analytics are today used in almost all areas of science ("computational science").

And over the past years, industry – and in particular SMEs - picked up in some limited way some of the techniques and methods developed within the HPC community. EU funded projects, such

[1] Cf. "FP-9 Vision Paper" of ETP4HPC, to be published in March 2019

as Fortissimo [2] and PRACE's SHAPE program [3] supported this evolution by accompanying SMEs in this up take.

Especially due to the wider range of data driven HPC applications, HPC has also gained in visibility by a broader public and at higher management levels in industry. Until now, HPC was more considered as part of the overall-IT cost, whereas now HPC is gaining overall awareness and becomes a topic of its own.

Nevertheless, the change we see in adopting HPC tools, methods and applications in industry is rather slow (counting in "decades" rather than in "years"). This is a serious problem for European industry today, as with the rise of the even more data-driven HPC applications, these technologies are known to be key for being and staying competitive in many domains.

The main reason for this slow adoption is that the way HPC systems are conceived and operated today is not suitable for a larger, non-expert public:

- Use of custom and non-standard hardware platforms:
  Today's HPC systems are not only complex due to their size and their parallelism, they also rely on specific hardware and specific software, making it almost impossible for non-experts to use those systems. For example, it is almost impossible to use standard applications out-of-the-box in an efficient way, as the hardware and the software configurations of every HPC system are slightly different, so that applications need to be adapted and tuned for every new system.
- Lack of a standardized software stack:
  As of today, there are only a few "de facto" standard components for the typical HPC software stack. It mostly boils down to C/C++ and Fortran compilers with OpenMP support for multi-threading, MPI libraries for communication and synchronisation, and some basic linear algebra libraries. But even they they differ in version and feature support on different systems! Most other components needed are either vendor-specific or developed by research groups, leading to a scattered landscape of different software bricks.
- Obsession with (raw) performance rather than productivity:
  Most typical applications today use only 7-10% of the full parallel capacities of the HPC system because the development of HPC architectures in the past decades concentrated on providing ever more double-precision calculation power but neglected to increase data access and movement speed accordingly. These architectures are no longer a good fit to today's more data-oriented workflow applications. Nevertheless, raw calculation performance has been for many years the main criteria for defining a successful HPC system. Today, the focus must be on more productivity and on delivering results[2].

We must find ways to speed up the adoption of HPC technologies by the European industry and we must broaden the use of HPC technologies and applications in academia. Therefore, HPC systems and setups need to be more accessible to non-experts, for example by provisioning HPC resources in a service-oriented manner. Moreover, a better integration of data pre-/post-processing - including data analytics - with visualization, alongside "traditional" simulation(s) is needed. This could be achieved by integrating HPC facilities into large-scale distributed workflows for scientific applications and by seeking more interactive usage models for HPC systems.

[2] We should build on the experience of the weather- and climate community, who is used in coping with such criteria.

2.3 Hardware evolutions disrupt the current landscape

Moore's law, which states that the number of transistors in an integrated circuit doubles every two years, is coming to an end [1]. The main reason is that the current size of transistors[3] is now close to the size of silicon atoms used in silicon chips (0.2 nm).

In this context, a first opportunity is vertical scaling, i.e. going from 2D to 2.5D or 3D. Flash memory systems have exploited efficiently this opportunity in the last few years and the use of 3D logic technologies is expected in the next few years. Nevertheless, Moore's Law is expected to end around 2030.

This event creates many opportunities and challenges, which should be reviewed in the coming years [6] [7]. The influence of the end of Moore's Law on the economics of devices and technologies is an important issue [8]. Indeed, Moore's law induced a significant concentration (with only Intel, Samsung, Globalfoundries and TSMC remaining) because of the huge investments needed and the short life expectancy of chips. The end of Moore's may induce longer lifespan for devices (10-15 years rather than 3 years), with a risk of disrupting the current economic model of silicon vendors.

Due to this transformation, radically different hardware devices are currently explored. In what follows, we survey the main future hardware architecture developments that are expected to impact HPC in the next years. Some of these developments have already started and are expected to have a huge impact in the next 5-8 years, such as reconfigurable architectures, PCM (Phase Change Memory) and Memory Centric Processing, as well as some neuromorphic architectures.

Some other technologies (analog and quantum computing in particular) are not expected to be available along this timeline. We include them here as they may induce deep changes for application developers, that require anticipation.

2.3.1 Reconfigurable architectures

To overcome the end of the growth of the number of transistors in a processor, one direction is to better use the transistors that are available. In a current processor architecture, a lot of the transistors of the chip are useless for most applications. Moreover, the Von-Neumann architecture is most of the time not efficient for specific data processing. Reconfigurable computing is a path for this optimized use of the available transistors.

Reconfigurable computing with FPGAs is a long standing research domain (with the 26th edition of the Workshop on Reconfigurable Architectures in 2019, and the 29th edition of the Conference on Field-Programmable Logic and Applications [10] [11]). FPGA architectures have evolved and diversified to more complex systems, such as embedded processors, on-chip reconfigurable memory, network interfaces, digital signal processors (DSP) arithmetic blocks. The advent of GPUs in recent years, that offer much better performance than CPUs at a lower price in terms of performance per watt slowed down research on FPGAs. They nevertheless offer an improvement of 2 orders of magnitude for some applications. Recently, the advent of Google's Tensor Processing Unit (TPU) renewed the interest in FPGAs [12].

These devices can be of us in arithmetic, signal processing, or network packet processing and in applications for bioinformatics.

[3] TSMC plans to build its first 3 nm manufacturing plant, whose production should be available in 2022; and 5 nm chips manufacturing plants are expected to start in early 2020 [4]

Among the difficulties that traditionally negatively impact the development of FPGAs is the development time to produce high quality codes and the compile time. This, however, is currently changing: AWS recently launched the opportunity to develop efficient FPGA accelerators for Amazon F1 in ½ day by using F1 Instances and additional frameworks such as Accelize [13] [14].

2.3.2 Non Von Neumann architectures

The general idea in **memory centric processing** is to put computations close to memory in order to save energy, to avoid latencies, and to enable a higher degree of parallelism when performed low in the memory hierarchy. There are many places where these computations may take place [7], with different scales and speeds. In general, the amount of memory per core is expected to decrease following the current trends, which will lead to more storage at intermediate levels and thus to additional layers in the storage hierarchy.

In [16] it is shown that it is possible to run a rather high-level computations based on the automatic detection of correlations between event-based data streams. Another example of technology is developed in the context of a collaboration between IBM Zurich and RWTH Aachen University [17]. Several systems have been produced by IBM in the context of PCM. CEA/Leti has developed a concept of PCM that can be stacked over a network of processors allowing connection of the 2 paradigms.

Neuromorphic computing is based on the idea of emulating the behaviour of the brain. There are two main approaches in this context.

The first approach is the use of deep neural networks (DNNs) that have shown their remarkable performance in specific contexts such as vision, speech recognition, translation. Having dedicated optimized devices to process part of a complex scientific application is promising, provided that the applications can be rewritten so has to benefit from such devices.

This approach has already been discussed in the context of FPGAs while discussing Google's TPU. Similar FPGA based devices can be found at AWS and Azure [18] [19].

The second approach is to build and simulate a huge distributed system by analogy with the brain, with spiking neurons and synapses, in order to understand what can be computed. In such a system, memory and processing scale together, thus creating a potentially promising model. The SpiNNaker project, supported by the Human Brain Project, built a chip based on these ideas [20] [21].

Neuromorphic architectures can be seen as specific **dataflow architectures**. The investigation of dataflow architectures is tightly linked to the use of Field Programmable Gate Array (FPGA), as most of the ideas have not led to the tape out of specific circuits but have been tested and implemented with FPGA.

Data flow architectures can provide an alternative to deliver more performance than standard processor for specific algorithms. The fact that the architecture is organized around the flows of data can save a lot of transistors and thus deliver much more energy efficient circuit. The development of reconfigurable architectures (like Intel's Configurable Spatial Accelerator (CSA)) and progress toward flexible reconfigurable FPGA will be an asset for implementing data flow architectures.

2.3.3 Optical Computing and Silicon Photonics

Photonics and optical technologies appear as potential solutions to increase network bandwidth with a lower requirement in energy, and a better robustness to noise and degradation. Due to recent advances, it is now possible to create silicon photonics switches compatible in size and design with Complementary Metal Oxide Semiconductor technology (CMOS). Silicon photonics can be used to implement neuromorphic architectures advances [22] [23].

Optical computing has been tested by the ESCAPE project [26] for the computation of spectral transform. Today, precision is not at the right level for numerical simulations but this approach can be interesting for applications looking for binary answers, like in pattern matching [24].

2.3.4 Analog Computing

Analog computing is based on the idea of using a physical process that is able to perform a computation at a much lower energetic cost. This is a relatively old idea, with recent developments based on analogy (i) with electrical devices (to solve PDEs in [28]), (ii) with thermodynamics (annealing) or (iii) with optics (optical neural networks presented in [29]). For instance, for optical neural networks, the matrix of weight can be encoded into a photonics circuit and all operations can be performed at the speed of light.

There are many developments, in optical neural networks in particular. Nevertheless, these are still experimental devices and it is unlikely to see them integrated in HPC machine in a near future (5-10 years).

2.3.5 Quantum Computing

Quantum computing is based on the idea of making use of quantum-mechanical phenomena and to apply operations to qubits (that are superposition of states), so that N qubits can represent up to 2^N states simultaneously on which operations can be applied. The computational model is called universal quantum computer.

The set of operations that could benefit from quantum computing has been the object of many studies in recent years and goes well beyond Shor's algorithm for integer factorization. In [31], a zoo of kernels that could benefit from quantum computing is presented (with many problems in graphs, verifications of matrix computations). Thee date of advent for small devices or even larger systems is still unknown. IBM [32] and DWAVE [33] have started building quantum systems, but still strongly limited in size and in functionality.

It is still unknown if physical systems implementing the universal quantum computer model can be built. This is mainly due to the problem of keeping a large set of qubits coherent over a large number of quantum gates. Even though systems have already been built, it is unlikely that actual devices could be integrated in HPC systems in the next 5-10 years. The integration of quantum devices for HPC applications is discussed in [34].

3 Tackling the transformation

The highly volatile landscape drawn in the previous section requires revisiting many of the ways of dealing with HPC applications and resources. The challenges are not only related to technology but also to uses and abilities to federate resources. Indeed, for acquiring (via scientific instruments or sensors in the IoT) and analysing the data, new elements are needed in the computation environment. This also affects the governance of compute and storage resources: it will inevitable become more complex.

In the remainder of the section, we address the transformation from three points of view.

The first one is questioning how the large-scale distributed workflows perspective with integration of data pre-/post-processing (including data analytics and visualization) and simulation(s) can be performed in an efficient manner. Examples of applications falling under this perspective are digital twins and applications using data assimilation based on edge computing to acquire and perform data analysis.

Urgent computing (e.g. seismic array processing for earthquake and tsunami prevention/warning) is another kind of applications that required revisiting the way applications interacts with resources (i.e. more interactive usage models for HPC systems). This is illustrated in Figure 1 where the capabilities are intertwined to form a complete processing chain. This figure also shows the communities which competencies need to be put together to deal with all required skills.

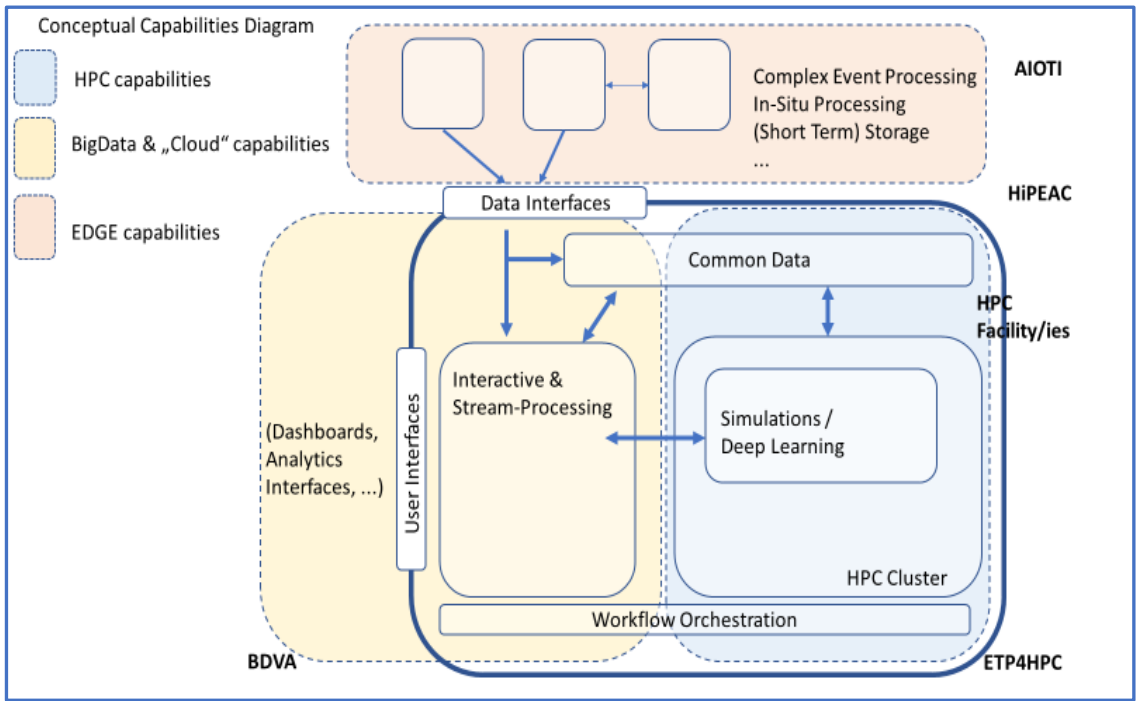

Figure 1: Intertwined Capabilities to form a complete processing chain[4]

The second point of view, developed in section 3.2, links the democratization of HPC to the capability to deliver "Service Oriented Architectures ". On one hand, making progress in the way to deliver computing and data resources is a condition to expand the use of HPC to new

[4] Source: Jens Krueger, ITWM Faunhofer, as part of collaboration BDVA-ETP4HPC-EXDCI2

communities. On the other hand, it is also the basis to construct application workflows that span over multiple infrastructures, owned by different entities.

Finally, the third point of view we consider is how to "cope with hardware evolution" (in Section 3.3). With the end of the Moore's Law, looking for extreme performance leaves few options others than looking into (algorithm-) specific hardware. In exchange for a high performance reward the programming of application is highly disrupted and the adherence of applications to a given system may be significantly increased, reducing the agility of its deployment as well as a more expensive development process.

3.1 Addressing the data challenge via workflows

Scientific discovery relies more and more on implementing workflows that bring together heterogeneous components. These components deal with experimental and observational data management and processing, as well as with computing and data analytics combined with (meta) data logistics. The workflows may include dynamicity and closed-loops where rare or unpredictable events trigger a flow of activities. They include a large ensemble of heterogeneous resources (from HPC centers to cloud facilities as well as edge/fog components). Uncertainty in data content adds uncertainty into workflow reaction/evolution. In this vision, HPC has to transition from short to long running persistency of services; from static to dynamic allocation of resources, data, and tasks; from homogeneous to heterogeneous software stack and services. The deployment of such complex workflow not only questions the HPC infrastructure (hardware and software) access and use modes, but also the networks infrastructure, and will require new approaches to security. It should be noted that when dealing with large-scale distributed workflows the technologies span over three distinct communities represented by ETP4HPC, BDVA and AIOTI (cf. Figure 1).

The first challenge for implementing application based on data/compute workflows, is to address the security issue. It is particularly difficult because this is assuming a global identification of users/application to simplify the deployment as well as some form of global governance/agreement between owners of the set of infrastructures to be used.

The allocation of resources is another challenge where different resource allocation models and orchestration must be considered (batch, stream, closed-loop, ...). For instance, the use of containers could be an option, but it requires that all resource providers agree. A related topic is the management of transversal metadata. To cope with this at application layer would lead system-specific applications (in contrast to transparent applications), which may hinder the application development.

Data logistic is another matter to be seriously considered. In many cases, moving the data between the systems must be performed (or not) in a control manner (contrary to Internet that is oblivious to the data path) in order to provide the necessary storage, compute, bandwidth and/or latency. The CDN (Content Delivery Network) approach may be necessary in some cases. Furthermore, it may be necessary to provide global addressing/naming scheme in order to help the workflow deployment and system adherence.

Last but not least of the issues is the ability of programming, debugging and monitoring the workflows across the systems in order to be able to perform global end-to-end performance analysis. A related issue is the how to address reproducibility of the application results.

3.2 Service oriented architectures to address usage evolution

With the service-oriented architecture framework we can nicely express applications as "complex workflows". This architecture framework also highlights many of the issues that need to be addressed: the linking of HPC compute and other compute services (at the edge or in the fog), with storage- and other data related services into one architectural framework.

This approach requires changing the way of providing HPC compute and extreme-scale storage resources to benefit from concepts and technologies developed for cloud computing with the following goals:

- Provide support for new types of workloads, e.g. data analytics workflows;
- Enable interactive access to HPC compute and extreme-scale storage systems in combination with elastic access to scalable compute resources. For instance, private facilities have more restrictions but they may want to get access to streams of data;
- Facilitate collaborative research;
- Improve support of domain specific platform services which allow for seamless scaling from local servers to remote supercomputers;
- Increase permeability of supercomputing systems to facilitate data (e.g. sensor and other experimental data) injection into the data centre. For instance, genomics requires processing genomic in one facility and use sequencer data on supercomputers.

Supercomputers today are typically deployed in a silo-style, with limited external connectivity, proprietary access processes[5], relatively rigid operational models that foresees users to submit batch jobs[6], and limited flexibility in terms of software stack provisioning. Cloud computing architectures on the other side are designed for being openly and easily accessible through uniform interfaces, flexible deployment of software stacks for enabling platform services, elastic provisioning of resources, and many more advantages.

Many of the outlined limitations of today's model for supercomputing systems could be overcome by defining a set of services with the following features:

- Standardised and provisioned by different supercomputing centres to facilitate more uniform access to HPC compute and extreme-scale storage resources
- Facilitate federation of services (requires support of a single authentication and authorisation infrastructure (AAI))
- Include services that are (typically) not provided today, e.g. interactive computing services, standardised interfaces for exchanging data.

These constitute key challenges[7] where R&D efforts are needed in order to adapt community practices (economical, governance, technical, security, …). Addressing these community challenges must not only be addressed by the scientists but also by the HPC centre operators and architects, the experts from organisations and projects related to distributed IT infrastructures (e.g. PRACE, the European Open Science Cloud (EOSC)).

[5] Processes for obtaining accounts, requesting resources, connecting to the systems, transfer data to or from the system largely differs between the different supercomputing systems.
[6] Job execution is scheduled primarily with the goal of maximising system utilisation and little options to take users requirements into account.
[7] Note that the challenges are not only technical but, e.g., also political.

3.3 Coping with hardware evolution

As sketched in Section 2.3, we expect within the upcoming years the advent of a multitude of different type of accelerators, each with potential benefits. This will lead overall to even more heterogenous systems, increasing the overall system complexity. Moreover, these devices will be suitable only for some specific applications (or in some specific cases), so that for each application the best system architecture needs to be understood and at runtime the corresponding resources need to be allotted accordingly.

The challenge will be to integrate those new devices wisely in order to benefit from their advantages, by coping with the induced complexity at system level and at application level (during the application development and at runtime).

Promises
Energy consumption is one of today's main bottlenecks, and ways to provide HPC resources at lower energy cost are urgently needed. FPGAs, PCM and memory centric approaches, as well as neuromorphic and optical compute units are promising to be more energy efficient.

Moreover, analog computing and optical computing are promising considerable gains in computing time. Lower latency is to be expected in particular in memory centric approaches and through optical devices.

On top, these technologies offer ways to cope with data driven architectures: FPGAs allow for a seamless integration into in-situ analysis chain. Technologies such as PCM and memory centric computing will allow the adaptation of distributed IoT algorithms for our purposes and facilitate "streaming operations", for example for encryption/decryption, search, algorithms on very large graphs and for deep learning.

Impact
FPGAs for example raise the problem of how to decompose a complex application into blocks that are well suited to existing devices, beyond having compilers to optimize existing kernels. For instance, how to rewrite a scientific application to benefit from DNN capabilities is a non-trivial task (in the case of CPU/GPU platforms, dynamic runtime schedulers have proven to be a potential solution). Another important open question is related to the variety of FPGAs that will be found in future HPC systems. Will we have to cope with five, ten, 50 or even 100s of different accelerators? This question has clear impacts on resource allocation problems, on languages to express parallelism (the recent example of CPU/GPU shows that the solution is not obvious), and on portability.

Memory centric approaches will be able to benefit form a large variety of ideas from IoT , such as streaming/sketching algorithms – assuming a conventional programming interface at all levels. Another important question raised by memory centric computing is related to coherency and consistency. Having operations achieved at potentially many levels of the memory hierarchy makes these problems extremely difficult.

Regarding **Quantum Computing**, we need to understand how quantum devices may be useful in the context of HPC applications and what level of changes are needed with respect to application conception at a high level and language and compilation at a lower level.

Although it is clear that linear algebra kernels can be implemented using **optical neural networks**, a classification saying which devices can be used to solve which specific problem would be useful to understand whether there are new kernels that can be solved efficiently with analog computing and with other types of devices, or if it is only a question of speed.

Major R&D Topics
Beside those specific technical specificities, which arise with the new hardware, we see some underlying, more general topics which need to be addressed. The extreme heterogeneity of future systems will come along with some more (parallel) programming paradigms. Ways to cope with those different paradigms will be key for a successful integration of the upcoming hardware. These reflections must include also the integration of legacy code.

4 Conclusion and next steps

This document highlights a set of challenges that are becoming more acute as the HPC technology and use is profoundly changing. The introduction of big data in the discovery process is driving application to be implemented as large-scale distributed workflows that need to be deployed on a large set of systems, each one having its own idiosyncrasies (e.g. quantum accelerator) that questions the way application are implemented.

Furthermore, it is not a matter only of technologies but also of being able to federate infrastructure. Indeed, it will be more frequent to mix the infrastructure for data collection, storage, data analytics and numerical model computation. The idea where a system fits all is becoming more and more obsolete and revisiting the HPC infrastructure as the centre of the process leads to many new challenges and potential re-organisation of the ecosystem.

Based on these elements, we will develop this Transversal Vision further in the upcoming months, with support from the EXDCI community at large, the application developers, and with the support of other European experts in this domain.