

HPC4AI, an AI-on-demand federated platform endeavour

Original

HPC4AI, an AI-on-demand federated platform endeavour / Aldinucci, Marco; Rabellino, Sergio; Pironti, Marco; Spiga, Filippo; Viviani, Paolo; Drocco, Maurizio; Guerzoni, Marco; Boella, Guido; Mellia, Marco; Margara, Paolo; Drago, Idilio; Marturano, Roberto; Marchetto, Guido; Piccolo, Elio; Bagnasco, Stefano; Lusso, Stefano; Vallero, Sara; Attardi, Giuseppe; Barchiesi, Alex; Colla, Alberto; Galeazzi, Fulvio. - ELETTRONICO. - (2018), pp. 279-286. (15th ACM International Conference on Computing Frontiers, CF 2018 ita 2018) [10.1145/3203217.3205340].

Availability:

This version is available at: 11583/2712210 since: 2018-09-13T08:03:27Z

Publisher:

Association for Computing Machinery, Inc

Published

DOI:10.1145/3203217.3205340

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

HPC4AI, an AI-on-demand federated platform endeavour

Invited Paper

Marco Aldinucci*
University of Torino, Italy
marco.aldinucci@di.unito.it

Filippo Spiga
University of Torino, Italy
filippo.spiga@unito.it

Marco Guerzoni
University of Torino, Italy
marco.guerzoni@unito.it

Paolo Margara
Politecnico di Torino, Italy
paolo.margara@polito.it

Guido Marchetto
Politecnico di Torino, Italy
guido.marchetto@polito.it

Stefano Lusso
INFN Torino, Italy
stefano.lusso@to.infn.it

Alex Barchiesi
Consortium GARR, Italy
alex.barchiesi@garr.it

Sergio Rabellino
University of Torino, Italy
sergio.rabellino@unito.it

Paolo Viviani
University of Torino, Italy
paolo.viviani@unito.it

Guido Boella
University of Torino, Italy
guido.boella@unito.it

Idilio Drago
Politecnico di Torino, Italy
idilio.drago@polito.it

Elio Piccolo
Politecnico di Torino, Italy
elio.piccolo@polito.it

Sara Vallero
INFN Torino, Italy
sara.vallero@to.infn.it

Alberto Colla
Consortium GARR, Italy
alberto.colla@garr.it

Marco Pironti
University of Torino, Italy
marco.pironti@unito.it

Maurizio Drocco
University of Torino, Italy
maurizio.drocco@unito.it

Marco Mellia[†]
Politecnico di Torino, Italy
marco.mellia@polito.it

Roberto Marturano
Politecnico di Torino, Italy
roberto.marturano@polito.it

Stefano Bagnasco[‡]
INFN Torino, Italy
stefano.bagnasco@to.infn.it

Giuseppe Attardi[§]
Consortium GARR, Italy
giuseppe.attardi@garr.it

Fulvio Galeazzi
Consortium GARR, Italy
fulvio.galeazzi@garr.it

ABSTRACT

In April 2018, under the auspices of the POR-FESR 2014-2020 program of Italian Piedmont Region, the Turin's Centre on High-Performance Computing for Artificial Intelligence (HPC4AI) was founded with a capital investment of 4.5M€ and began its deployment. HPC4AI aims to facilitate scientific and engineering research in Artificial Intelligence, Big Data Analytics and the methods to provide them as a service to the regional and national industrial community, including the large regional ecosystem of Small-Medium Enterprises (SMEs) active in many different sectors such as Automotive, Aerospace, Mechatronics, Manufacturing, Health and Agrifood.

*HPC4AI coordinator, Principal Investigator at University of Torino

[†]Principal Investigator at Politecnico di Torino

[‡]Principal Investigator at INFN Torino

[§]Principal Investigator at Consortium GARR

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

CF'18, Ischia, Italy

© 2016 Copyright held by the owner/author(s). 978-x-xxxx-xxxx-x/YY/MM...\$15.00
DOI: 10.1145/nnnnnnn.nnnnnnn

KEYWORDS

AI-on-demand, Federated Cloud, HPC, Big Data, Fog, Industry 4.0

ACM Reference format:

Marco Aldinucci, Sergio Rabellino, Marco Pironti, Filippo Spiga, Paolo Viviani, Maurizio Drocco, Marco Guerzoni, Guido Boella, Marco Mellia, Paolo Margara, Idilio Drago, Roberto Marturano, Guido Marchetto, Elio Piccolo, Stefano Bagnasco, Stefano Lusso, Sara Vallero, Giuseppe Attardi, Alex Barchiesi, Alberto Colla, and Fulvio Galeazzi. 2016. HPC4AI, an AI-on-demand federated platform endeavour. In *Proceedings of ACM Computing Frontiers conference, Ischia, Italy, May 3018 (CF'18)*, 8 pages.
DOI: 10.1145/nnnnnnn.nnnnnnn

1 INTRODUCTION

Artificial Intelligence (AI), together with data-centric research areas such as Machine Learning (ML) and Big Data Analytics (BDA), currently represents the most valuable technological breakthrough for the Italian and European manufacturing and service sectors. It is commonly viewed as an enabling technology for the ongoing industrial revolution, commonly referred as *Industry 4.0*.

From the computer science viewpoint, the Industry 4.0 vision is outlined as an increasingly connected ecosystem of devices that

produces digital data of increasing variety, volume, speed and volatility. In order to fully exploit its potential, AI applications of the next future must embrace High-Performance Computing (HPC) techniques and platforms. Secondly, computing and data management capabilities of HPC must be made readily and easily accessible *on-demand* to data scientists, who are more used to perform their work locally on interactive platforms.

The Turin's Centre on High-Performance Computing for Artificial Intelligence (HPC4AI) addresses these two synergistic aspects by leveraging the skills available in two large universities and eight research institutes, federating their laboratories in a single regional centre of national importance. This new centre, equipped with state-of-the-art components of adequate size to be effective, will achieve self-sustainability thanks collaborations and business opportunities with local companies.

Over the last several years, there has been a dramatic increase in the amount data available from multiple heterogeneous sources. Local companies struggle to tap the vast potential of the data deluge and gain competitive advantage through the ability to develop solutions at the frontier of research avoiding any technological lock-in by international technology conglomerates such as Google, Microsoft or Amazon. For this reason, it is necessary to create a flexible and agile research-oriented environment, inspired by a model of public-private partnership, where skills can be developed to the highest level and capable of self-sustaining through the synergy between universities, industries and SMEs.

This position paper describes the motivation, the objectives and the general structure of the HPC4AI centre.

2 MOTIVATIONS AND OBJECTIVES

General-purpose AI platforms and services are currently offered by commercial cloud providers as an integral part of their cloud portfolio. The vast majority of commercial cloud offerings have significant drawbacks with respect to the needs of an open AI-on-demand platform for the European research and industry:

- They are based on proprietary solutions that produce technological lock-in. After applications have been built, the platform will continue to exact a rent for the services deployed. The long-tail of this *rental economy* model increases inequalities rather than controlling them [28].
- Modern AI algorithms require specialised hardware, typically general purpose Graphical Processing Units (GPUs) or Field-Programmable Gate Array (FPGA). The cost of virtual machines that use specialized hardware is still very high in commercial public clouds; for such kind of resources, the break-even point of buying versus renting occurs quite early.
- AI and BDA algorithms work very effectively mainly on annotated datasets, which require significant human effort for curation and maintenance. Such datasets may also contain sensitive information, therefore should be stored appropriately and be compliant with European regulations if within European borders.

A cloud infrastructure dedicated to AI able to address such drawbacks should be considered a strategic asset for the regional, national and European communities.

The generation of innovative results in the field of Artificial Intelligence and Deep Learning requires performing many experiments on huge volume of data. It is also well understood that the time and the effort involved in each experiment is somehow related to the size of the training data, the characteristics of the network and the number of parameters involved. To overcome this situation, HPC4AI explicitly addresses two main technical challenges: (1) Provide AI and BDA users with an *on-demand* platform supporting AI and BDA solutions as services; (2) Review HPC methods and tools to meet the needs of AI and BDA.

2.1 Everything-as-a-Service by way of a Declarative Modelling Approach

HPC4AI embraces a declarative modelling approach that demonstrated its effectiveness at different levels of abstraction. A couple of examples in the direct experience of partners are the design of GARR cloud and BDA services in the *Toreador* project [31].

In the GARR cloud, the platform and the federation are managed and maintained through automated tools, greatly reducing the manpower required to implement a modern and efficient cloud computing platform. Automation is based on a declarative modelling approach [3] that is used to describe the structure and constraints of the cloud system components. An automation tool generates a deployment plan consisting of the steps required to achieve the requested configuration, transforming the state of the system until it satisfies all constraints. The approach offers several benefits:

- *Portability*: models are described in a declarative fashion, abstracting from the specifics of a cloud provider and hence they can be ported to different platforms
- *Consistency*: both physical and virtual infrastructures can be modelled, as well as the relationships between infrastructure, network, and application components.
- *Automation*: mapping a model onto infrastructure and cloud-specific deployment operations is a responsibility of the orchestrator rather than of a system administrator.

At a higher level of abstraction, methodologies to design BDA-as-a-Service (BDAAA) based on model transformations were developed in the context of the *Toreador* project. Abstract declarative models are used for expressing customers' goals and requirements in all aspects of BDA, including data representation, processing, analytics, security and privacy specification, deployment-strategy and visualisation tool definition. Such models are then converted into platform-independent procedural models specifying the desired procedures for data preparation and algorithms (including parallelisation), to be later compiled in provisioning and configuration of the computational resources on the execution platform [31].

2.2 HPC4AI, where HPC cross-pollinate AI

As digital society evolves, an increasingly connected ecosystem of devices produces more volumes and variety of data. They range from devices-in-the-fog to large and instruments addressing the grand challenges of science. To keep up with the pace, very large volumes of dynamically changing data ought to be processed, synthesised, and eventually turned into knowledge. High-velocity data

brings high value, especially to volatile business processes, mission-critical tasks and scientific grand challenges. Some of this data loses its operational value in a short time frame, some other is simply too much to be stored.

The support of parallel programming models and platforms allows both BDA and AI to address this challenge. The convergence between BDA and HPC is witnessed by the ecosystem of Big Data stacks, which are evolving at a fast pace but still revolve around the *MapReduce* [9] model, which eventually is a parallel computing paradigm itself.

Nowadays, without parallel computing it will be impossible to conduct any modern experiment in a reasonable amount of time. For example, training with 150 epochs of *BoNet* on 2000 images (X-ray) requires 8 days on a dual core CPU (Intel Core i7 @2.50Ghz) but only 3 hours on NVidia Titan X GPU [27]. On the same platforms, the training with 25 epochs of a deep generative network for video synthesis on over 600k videos require 420 days on the CPU and 7 days on the GPU [26].

Such disruptive advances are dramatically changing computing provisioning. Both general-purpose hardware and traditional HPC systems are not fully adequate for these new kinds of computations. To achieve system efficiency elements execution speed, memory, accuracy and energy consumption need to be taken into account during the design. For example, setting an energy constraint will enforce an upper bound on the maximum achievable accuracy and model complexity [24]. Floating point operations at mixed precision (FP16 and FP32) are acceptable in many ML and DL applications and allow reducing both memory footprint and energy consumption against single precision (FP32) and resulting in an increased computation speed. For example, the training of AlexNet, VGG-D, GoogLeNet, Inception v1, and Resnet50 on ILSVRC12 validation data result almost identical on top-1 accuracy when trained with single and mixed precisions [21]. Half precision (FP16) uses 16 bits, compared to 32 bits for single precision (FP32). Lowering the required memory enables training of larger models or training with larger minibatches. Execution time can be sensitive to memory or arithmetic bandwidth. Half-precision halves the number of bytes accessed, thus reducing the time spent in memory-limited layers. NVIDIA Volta GPUs offer up to 8× more half precision arithmetic throughput when compared to single-precision, thus speeding up math-limited layers.

For all these reasons it is imperative to start reviewing the organisation of data centres to accommodate AI and AI-related workloads, which may require specialised hardware (e.g. GPU, FPGA or ASIC) and complex memory hierarchies to feed them with large data streams.

Deep Learning algorithms can be more forgiving towards traditional approaches to parallelism. For example, due to the statistical nature of computations, it is possible to relax synchronisation constraints, by allowing access to shared memory without locking primitives or barriers, since computations on the same data will be performed repeatedly and hence one can afford to miss an update during an iteration. Many AI workloads still reserve unexplored territories for parallel programming models. In fact today there is no such thing as a *"Map Reduce for ML"*, i.e. a universal and scalable programming model to transparently scale-out next-generation distributed multi-modal training workloads.

3 A FEDERATED LABORATORY

HPC4AI will be implemented as four federated laboratories connected to each other by high capacity networks. A large part of the computing and storage resources acquired within HPC4AI will be integrated into the Italian national GARR cloud federation as the "Piedmont cloud zone" (*federated islands*) [14]. By doing so, these resources will be made available to other laboratories according to a standard rate card published on the portal of the centre. Within the federation umbrella, the management of the resources of each laboratory is under the sole responsibility of the university that supplies them. This greatly simplifies the governance and the sustainability of HPC4AI since many aspects will be delegated directly to individual laboratories. The remaining computing resources will form *Experimental Islands* and will be available for experimentation and research within the university and wide open to project partners and external consultancy engagements. The big picture of HPC4AI is sketched in Fig. 1.

3.1 Participants

As summarized by Figure 1, HPC4AI projects has two main partners: University of Turin (UNITO) and Polytechnic University of Turin (POLITO). Each of them participates in HPC4AI contributing with different specialisations via their Interdisciplinary Centers:

- **PdF@UNITO**, the *Piero della Francesca* centre of UNITO, to provide AI-on-demand its applications to innovation.
- **C3S@UNITO**, the *Competency Centre for Scientific Computing* of UNITO, to provide a containerised distributed platform for HPC applications.
- **SmartData@POLITO**, the research centre at POLITO, to provide infrastructure and experts' support to researchers and companies interested in Big Data and Data Science.
- **HPC@POLITO**, the *Academic Computing Centre* of POLITO, to provide the HPC infrastructure and technical support for academic teaching and research activities.

Several departments from the two Universities participate into the four interdisciplinary laboratories: Computer Science, Economics and Statistics, Mathematics, Law, Philosophy, Control and Computer Engineering, Electronics and Telecommunication Engineering.

To conclude, eight Technological Partners contributes into HPC4AI with their own expertise:

- The **GARR Consortium**, contributing by sharing its federated cloud software.
- The **National Institute of Nuclear Physics (INFN-TO)**, contributing by sharing its experience in designing and managing Cloud infrastructures (e.g. INDIGO-DataCloud [16]).
- The **Collegio Carlo Alberto**, contributing by sharing its experience in Economics and business models.
- The **Città della Salute e della Scienza**, which manages Turin's hospitals, contributing by sharing its domain specialism.
- The **Human Technopole**, contributing by sharing their large datasets of medical data.
- The **ISI Foundation**, contributing by sharing its experience in Complex Systems.

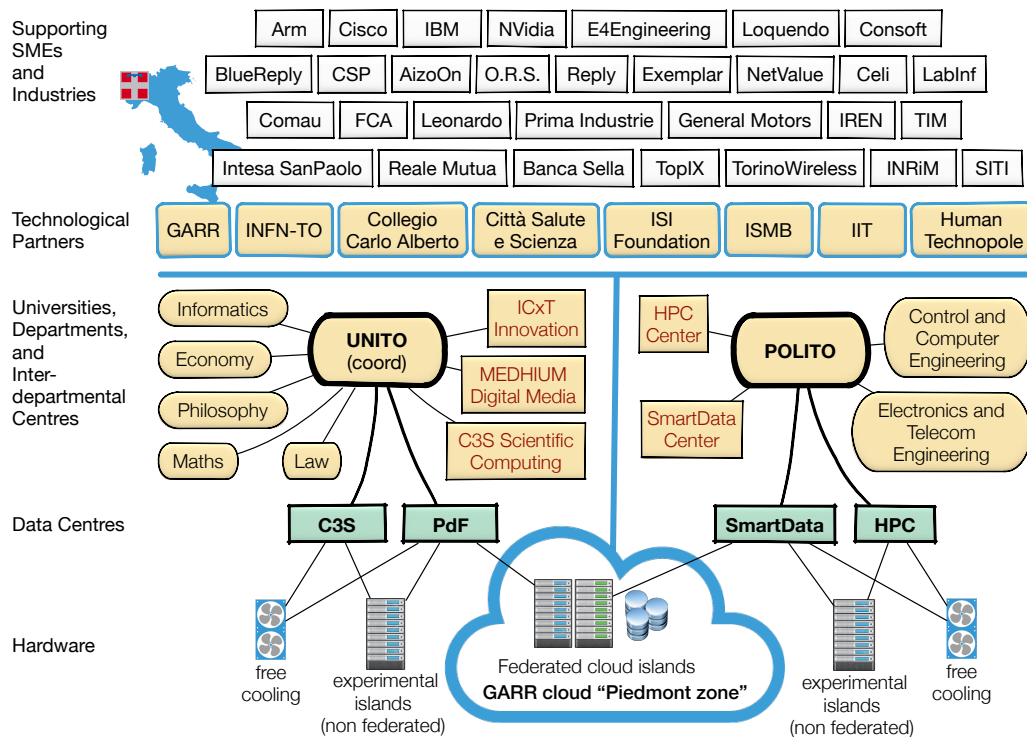


Figure 1: HPC4AI vision: shareholders and partnerships.

- The **Italian Institute of Technology (IIT)**, contributing by starting a new HPC4AI-related laboratory on medical imaging.
- The **Istituto Superiore Mario Boella**, contributing by sharing experience on managing large technological projects.

4 CLOUD FEDERATION: RELATED WORKS

Several large research infrastructures on a European scale are foreseen in the context of the ESFRI road-map [12], which aims to provide scientists adequate tools for scientific investigations. The common driver for all these initiatives is the ever-increasing demand for processing in terms of data volumes and computing power.

Seven universities in the state of Baden-Württemberg (Germany) have built federated community science cloud, through the *bw-Cloud* project [6]. The *bwCloud* architecture is based on *OpenStack* and uses a single *Keystone* service for authorization, while authentication is delegated to the *Eduroam* service among the affiliated universities. The project is focusing on a governance scheme suitable for managing a federated system that aims to support over 20,000 users. The scheme must take into account the needs and interests of the various stakeholders and different funding agencies. The coordination within *bwCloud*, therefore, includes the development of a solid business model and procedures that allow the participation and satisfaction of users.

SeaClouds aims to standardize management across multiple providers and to support a robust and scalable orchestration of services

on them [4]. *SeaClouds* will intrinsically support the evolution of the systems developed with it, addressing any necessary changes, even during the execution phase. The development, monitoring and reconfiguration through *SeaClouds* include a unified management service, where services can be implemented, replicated and managed through harmonised standard APIs such as the *CAMP* specification and the *Cloud4SOA* project.

Apache Brooklyn is an open source framework for modelling, monitoring and managing applications for cloud environments [5]. *Apache Brooklyn* is able to manage the deployment of applications and prescriptions, to monitor the health and metrics of an application, to understand dependencies between components and to apply complex criteria to manage the application defined in a project.

Acumos AI Platform [30] is a federated platform for managing Artificial Intelligence and Machine Learning applications and sharing AI models, following *simplicity* as main design principle. AT&T has been developing the project together with the Indian consulting firm *Tech Mahindra*, only recently the full open source code has been made publicly available and hosted by the Linux Foundation archives. *Acumos* includes a standard mechanism for downloading AI code from a marketplace, bringing AI functionality to where data resides instead of requiring developers to load data into a proprietary cloud service to access AI functions.

INDIGO-DataCloud developed a middleware to implement cloud services specifically relevant for scientific computing: authentication, workload and data management [16]. The *Indigo-DataCloud*

project focuses mainly on bridging the gap between cloud developers and the services offered by existing cloud providers, rather than provisioning a cloud service. The work started with the INDIGO-DataCloud project is currently being furthered by the DEEP-HybridDataCloud [10] and XDC-eXtremeDataCloud [13] projects for workload- and data-management aspects, respectively.

Nectar Cloud provides computing infrastructures, software and services that enable the Australian research community to store, access and manage data remotely, quickly and independently [22]. The Nectar Cloud self-service feature allows users to access their data at any time and collaborate with others from their desktop quickly and efficiently. Nectar Cloud uses OpenStack to automate deployment, including resizing and load balancing.

The EGI Federated Cloud e-Infrastructure implements a federated cloud [11], which was originally based on OCCI [23] and CDMI [2] as interfaces to web services for access to resources from OpenNebula and OpenStack cluster or from public providers. The approach is to provide an extra abstraction layer on the resources provided by the national network initiatives that remain separate and independent of each other.

HelixNebula [15] aims to explore the best way to exploit the commercial cloud providers in the procurement of cloud infrastructures for research and education. The approach is to create a private audience partnership for the acquisition of hybrid clouds.

The *GARR Consortium* is the Italian National Research Network organisation that provides high-speed connectivity to all Italian research institutions. GARR offers an advanced Federated Cloud Platform that provides both typical *IaaS* services as well as a convenient *DaaS* (Deployment-as-a-Service) that allows even non-experts to easily deploy cloud applications. The platform currently consists of about 9000 virtual cores and about 10 PetaBytes of storage, distributed in 5 datacentres. The cloud is fully based on OpenStack for the cloud infrastructure, Juju [18] for automating deployment and Kubernetes for containers orchestration. The platform architecture is a federation of multiple regions [3, 14], that allows other institutions to participate by creating their own region and share a quota of those on-premises resources with the rest of the trusted users. The federation can seamlessly include resources in public clouds. HPC4AI will implement the *Piedmont cloud zone* of the GARR federation.

5 METHODS AND TOOLS

5.1 Federated Cloud Platform

Building a well-supported and comprehensive Cloud infrastructure for Science is a far-reaching goal that requires significant investments coming either from commercial cloud providers or public organisations. A possible alternative is a federated approach, where the infrastructure is built bottom-up by combining medium facilities into larger ones, to achieve the adequate scale [3]. HPC4AI deploy the federated GARR cloud described in Sec. 4. Cloud federation makes possible to share resources preserving the ownership, which in the case of the GARR cloud potentially embraces many Italian research institutes [14]. Likewise, since users are also owners, they are directly involved in creating and sharing services and DevOps best practices, facilitating the transition of many on-premises research activities into an off-premises cloud computing paradigm.

5.1.1 Services. As described in Table 1, HPC4AI provides to users a stack of services with increasing level of abstraction: Metal, Infrastructure, Platform, BDA-as-a-Service (BDAaaS) and AI-as-a-Service (AIaaS).

The innovation potential of HPC4AI is built on four cornerstones:

- A *BDAaaS/AIaaS* layer specifically oriented to AI and BDA well-defined services. These services must be accessible to not tech-savvy domain experts: medical doctors, biologists, mechanical engineers, lawyers, sociologists, etc.
- A powerful and intuitive Deployment-as-a-Service (DaaS) layer from the GARR cloud layer to easily deploy complex applications in a few clicks. HPC4AI DaaS leverages on the GARR cloud that designed on top of the Juju application modelling tool and its predefined bundles of charms [18], pre-packaged recipes defining requirements and steps involved at each stage of a service lifetime. The charms to deploy Spark and Kubernetes clusters are shown in Fig. 2.
- A rich *marketplace* of Platforms, BDA, AI services (either for free or not) also including large datasets of specific interest for the local research community (see also the next section).
- The possibility to optimise and personalise services by moving them down to the next level of the stack and then moving them up again, e.g. to allow the optimisation or the re-training of a neural network (at the DaaS level) that is offered as a model for inference at the AIaaS level.

5.1.2 Datasets. Along with the computation resources, a repository for training datasets will be provided, allowing a fast access to the data, without the need to replicate or transfer it from other sources. In addition to many publicly accessible datasets, the project will store particularly interesting datasets like:

- A mirror of the Software Heritage archive that holds over 20 million software projects, with an archive of over 2.7 billion unique source files [8].
- Labelled brain Computer Tomography scans from the Interventional Neuroradiology Unit of Molinette Hospital in Turin.
- Labelled digital histopathology slides from the Department of Medical Sciences of the University of Torino.

Collaboration between scientific and technological partners will improve and extend available datasets. For example, the Città della Salute produces hundreds of thousands of medical images per year. A significant fraction of this dataset is usable for scientific research but often cannot be moved elsewhere (e.g. public cloud) because of their size or specific restrictions (e.g. permission of use, privacy concerns, etc.). Collaboration with technological partners is also expected to make data ready to direct apply deep learning techniques, for example by improving labelling practices with respect to current standards (e.g. marking regions of interest).

5.2 Containerised Cloud-like HPC Platform

The C3S@UNITO Centre provides the University of Torino research community with a multi-purpose facility called OCCAM, with fairly conventional HPC nodes, high-memory nodes, GPUs, a high-performance parallel storage based on the Lustre filesystem

Users	Kind of service	Services	Artifacts
Application domain experts with no skills on ML and/or BDA	AlaaS, BDAaaS	AI and BDA services directly designed within HPC4AI or third parties	Marketplace for ML and BDA solutions, e.g. by way of AcumosAI marketplace [30]: dashboards, ML/DL models in several domains, BDA pipelines
Application domain experts skilled on ML and/or BDA but not expert in parallel and cloud computing	PaaS by way of Deployment-as-a-Service (DaaS)	PaaS solutions for ML and BDA directly designed within HPC4AI and related projects; GARR cloud DaaS	Marketplace of Platforms (e.g. Juju bundles) to deploy software stacks for ML and BDA; solutions for data ingestion, data lake, etc.
Researchers, cloud engineering, ML and BDA framework designers, cloud engineers, stack and automation designers	Infrastructure-as-a-Service (IaaS)	GARR federated cloud	Virtual Machines and networks, object storage, file storage, etc.
Researchers, cloud engineering, ML and BDA framework designers, cloud engineers, stack and automation designers	Metal-as-a-Service (IaaS)	Job scheduler for HPC resources	Job queue, Container orchestrators, Big Data stack (Spark, etc.)
Researchers, run-time designers	Hardware	Bare Metal	Multi-core CPU, GPU, storage, network

Table 1: HPC4AI service stack.

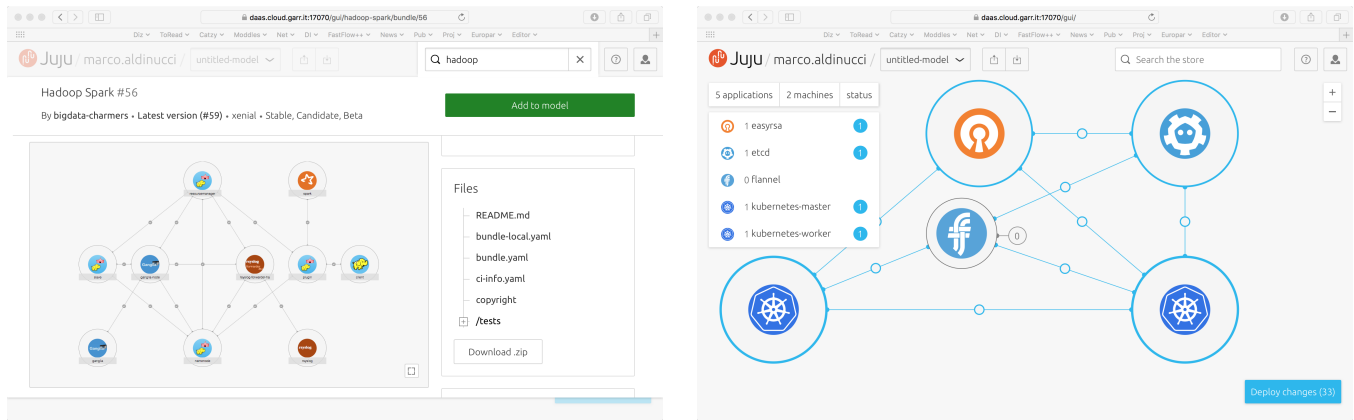


Figure 2: Two screenshots of the GARR cloud Deployment-as-a-Service (DaaS) interface: deployment of Spark (left) and Kubernetes (right) clusters.

[19] for data processing, and a high capacity archival storage area, with both 10 Gb/s Ethernet and 56 Gb/s InfiniBand interconnections.

The execution model on OCCAM [1] is fully based on containerisation technologies, giving the users maximum freedom in the choice of their runtime environment and, at the same time, decoupling the application software management from the infrastructure management.

A Computing Application is defined by a runtime environment (OS, libraries, software packages, services), its resource requirements (e.g. one large single-image node, HPC cluster, GPUs) and an execution model (e.g. batch jobs, pipelines, interactive access).

Linux containers provide an isolated execution environment for user software, and their use in HPC environments and, more generally, for scientific computing is currently a hot research topic. In most cases, special tools are used to run containers in user-space

on HPC nodes, such as Singularity [25] or Shifter [17]. In our case we decided to rely on Docker for container management and execution, Apache Mesos and Marathon for container scheduling [20], along with several tools and configuration templates produced by the INDIGO-DataCloud EU project [16].

Docker provides a popular ecosystem of tools for easy creation and distribution of computing applications in the form of container images, together with the availability of many off-the-shelf base images ready to use and customize. Thus, a user can develop her Computing Application container on her private workstation, then run it on OCCAM as-is or as a component in a more complex structure. Allowing users to run Docker containers natively on worker nodes poses obvious security concerns; some countermeasures were adopted to sanitise images before running them, cleaning up the attack vectors that a malicious user can leverage to break in the system security, without impacting on Docker performance or

features needed to run the legitimate user tasks. Furthermore, the users cannot access directly Docker's APIs on worker nodes but are provided with command line interface tools that interface with Docker in a controlled way.

The choice of a containerised platform has risen from the wide range of use cases, which in turns calls for a great flexibility in the way C3S@UNITO provisions its users with resources, moving from pre-installed software packages to consolidated applications. Even though the user interfaces and work-flow are fully defined and will not change, development work is still ongoing under the hood in order to strengthen the infrastructural back-end and allow for more advanced cross-application scheduling.

All the images are stored on a private registry integrated with the main user interface provided by a repository service based on GitLab, so the users can freely manage their images, share them with other users avoiding duplication and making simulations reproducible.

5.3 HPC Platform

Founded in 2012, the HPC@POLITO Academic Computing Centre provides the HPC infrastructure inside Politecnico di Torino to many engineering research groups.

The centre directly supports a wide range of researchers coming from nine different departments, offering computing resources to scientific topics ranging from computational fluid dynamics to bioinformatics, molecular dynamics, and machine learning. Furthermore, an application engineering support is provided to help users to take advantage of the power of modern HPC systems. Individual research groups can also contribute to the improvement of the centre's clusters by acquiring computational nodes that conform to the concurred standards.

Following the University educational mandate, HPC@POLITO also supports teaching activity such as hosting specialized programming courses of both Master and PhD students. An increasing number of students use Centre's computing facilities to perform complex and intensive numerical simulations for their thesis. Furthermore, the number of new research activities increases year by year as well as the number of papers, published on both journals and conference proceedings, whose scientific results are obtained thanks to the computational resources made available.

All the clusters hosted by HPC@POLITO are made by distributed-memory commodity nodes interconnected by InfiniBand and sharing the same high-performance Lustre storage. The centre already hosts specialised node equipped with accelerators, whose number is expected to grow thanks to HPC4AI.

Users have the ability to create customised environments by using Singularity [25] containers. Singularity is a software container technology specifically created with the HPC use case in mind. It can be used to package entire scientific work-flows, software and libraries. Unlike many other container systems designed to support trusted users running trusted containers, Singularity supports the opposite model of untrusted users running untrusted containers. Images are managed through a private Singularity registry application on which HPC's users can upload and download their images. Some pre-built images for widely used applications, for

example commonly used Deep Learning frameworks like Keras or TensorFlow, are made available by the staff.

5.4 Big Data Platform

The SmartData@POLITO centre was created to target the fundamental research challenges and practical applications of big data and data science. SmartData@POLITO developed experience with big data workloads using a small-scale infrastructure, built in 2015, which is able to store around 750 TeraBytes.

The cluster has been designed to be horizontally scalable based on commodity hardware, which means that new resources can be deployed besides the existing ones, enhancing performance without impacting end users. It includes both the Hadoop distributed file systems (e.g, HDFS) and frameworks for resource management (e.g, YARN) both needed to enable general compute engines for big data (e.g., Map-Reduce and Spark).

Such infrastructure has proven instrumental for bootstrapping the centre's activities, supporting research projects operating with dozens of TeraBytes, new collaborations with industries and Master's Courses attended by hundreds of students annually. It is however not sufficient to explore the new possibilities offered by the combination of AI and big datasets.

HPC4AI will allow to reach new frontiers on big data processing capabilities, moving the available capacity of storing and processing data from the TeraByte to the PetaByte space.

As big data applications are increasingly relying on machine learning and AI algorithms, the SmartData@POLITO will explore novel alternatives to deploy a big data infrastructure that allows advanced AI to be executed on extremely large datasets. In terms of hardware, this will include the provisioning of infrastructure that is not only focused on provisioning large storage capacity but also compute to better handle memory-intensive workloads and the presence of specialised hardware (e.g., GPUs).

6 HPC4AI BUSINESS MODEL

Implementing a successful business model is increasingly seen as a source of innovation. Such model has to be managed in line with an organisation of business processes and strategy to become a key shaping factor for long-term sustainability [7].

Regardless of how innovation in the business model occurs, its structure [29] must enable the organisation to provide value to all its stakeholders whilst protecting its competitive position. This requires a business model that is flexible enough to allow the organisation to anticipate problems, identify and exploit opportunities, correct projects which are going astray and respond to social and technological change. For example, HPC is inextricably linked to innovation by enabling breakthroughs in science, engineering. In a business perspective, HPC is viewed as a cost-effective tool for speeding up the R&D process. HPC4AI aims to adopt a AI-centric business model which can create and defend values, seizing opportunities, able to adapt according to market needs.

For HPC4AI, we considered a model in which technologically advanced players begin to play a dual role. They are not only actors within their industry, providing value propositions to their segments of customers, but also facilitators, enabling other players within the sector to innovate their value proposition. We are

planning a technology platform which provides not products or services, but an integrated business support service. However, what is less reproducible is the creation of an infrastructure that will facilitate the company's business (1) by promoting collaboration among key stakeholders such as suppliers, customers and producers, (2) through the creation of new opportunities by monitoring transactions and recording a high volume of information and data, and (3) by disseminating innovative trends and new technologies within different sectors.

7 CONCLUSIONS

HPC4AI starts on April 2018 with an investment of 4.5M€ and a seven years' work plan divided into two periods: two years of design and deployment and five years of operations. In the second period, its total cost of ownership is expected to be sustained with a revenue stream coming mainly from research projects funded by several Italian/European initiatives and industrial collaborations.

Beyond sustainability, it is expected as main outcome the strengthening of the scientific communities involved in this endeavor. This will be evaluated by looking at the scientific outputs of both resident scientists and their students, which ultimately embody the innovation potential of the regional and national industry.

HPC4AI is not designed to bring an immediate financial gain for partners but to support the long tail of science, allowing researchers of any discipline, without expertise in computing systems, to deploy applications readily available for production, and the researchers in computer science to master enabling technologies for the imminent Industry 4.0 revolution.

ACKNOWLEDGEMENTS

This paper has been partially supported by the HPC4AI project funded by Italian Piedmont Region, by the EU-H2020-RIA Rephrase project (no. 644235), the EU-H2020-RIA Toreador project (no. 688797), and the EU-H2020-IA Fortissimo-2 project (no. 680481).

REFERENCES

- [1] Marco Aldinucci, Stefano Bagnasco, Stefano Lusso, Paolo Pasteris, Sergio Rabbellino, and Sara Vallero. 2017. OCCAM: a flexible, multi-purpose and extendable HPC cluster. *Journal of Physics: Conference Series* 898, 8 (2017), 082039. DOI: <http://dx.doi.org/10.1088/1742-6596/898/8/082039>
- [2] Storage Networking Industry Association. 2018. Cloud Data Management Interface. <https://www.snia.org/cdmi/>. (Last accessed, Mar 2018).
- [3] Giuseppe Attardi, Alex Barchiesi, Alberto Colla, Fulvio Galeazzi, Giovanni Marzulli, and Mario Reale. 2017. Declarative Modeling for Building a Cloud Federation and Cloud Applications. *CoRR abs/1706.05272* (2017), 1–23. <http://arxiv.org/abs/1706.05272>
- [4] Antonio Brogi, José Carrasco, Javier Cubo, Francesco D'Andria, Elisabetta Di Nitto, Michele Guerriero, Diego Pérez, Ernesto Pimentel, and Jacopo Soldani. 2016. SeaClouds: An Open Reference Architecture for Multi-cloud Governance. In *Proc. of the 10th European Conference on Software Architecture (ECSA) (LNCS)*, Vol. 9839. Springer, Copenhagen, Denmark, 334–338. DOI: http://dx.doi.org/10.1007/978-3-319-48992-6_25
- [5] Brooklyn 2018. Apache Brooklyn. <https://brooklyn.apache.org/>. (Last accessed, Mar 2018).
- [6] bwCloud 2018. bwCloud Project. <http://www.bw-cloud.org/>. (Last accessed, Mar 2018).
- [7] Henry W. Chesbrough. 2010. Business model innovation: opportunities and barriers. *Long Range Planning* 43, 2-3 (2010), 354 – 363.
- [8] Roberto Di Cosmo and Stefano Zacchiroli. 2017. Software Heritage: Why and How to Preserve Software Source Code. In *iPRES 2017: 14th International Conference on Digital Preservation* (2017-09-25). Kyoto, Japan, 1–10. <https://hal.archives-ouvertes.fr/hal-01590958>
- [9] Jeffrey Dean and Sanjay Ghemawat. 2008. MapReduce: Simplified Data Processing on Large Clusters. *Commun. ACM* 51, 1 (Jan. 2008), 107–113. DOI: <http://dx.doi.org/10.1145/1327452.1327492>
- [10] DEEP Hybrid DataCloud 2018. DEEP Hybrid DataCloud project. <https://deep-hybrid-datacloud.eu>. (Last accessed, Mar 2018).
- [11] Enol Fernández del Castillo, Diego Scardaci, and Álvaro López García. 2015. The EGI Federated Cloud e-Infrastructure. *Procedia Computer Science* 68 (2015), 196 – 205. DOI: <http://dx.doi.org/https://doi.org/10.1016/j.procs.2015.09.235> 1st International Conference on Cloud Forward: From Distributed to Complete Computing.
- [12] ESFRI 2016. ESFRI Roadmap 2016. <http://www.esfri.eu/roadmap-2016>. (2016).
- [13] eXtreme Hybrid DataCloud 2018. eXtreme DataCloud project. <http://www.extreme-datacloud.eu>. (Last accessed, Mar 2018).
- [14] GARR. 2016. Federated Cloud Reference Architecture. (2016).
- [15] Helix 2018. Helix Nebula. <http://www.helix-nebula.eu/>. (Last accessed, Mar 2018).
- [16] Indigo-DataCloud 2018. Indigo DataCloud project. <https://www.indigo-datacloud.eu/>. (Last accessed, Mar 2018).
- [17] Douglas M Jacobsen and Richard Shane Canon. 2015. Contain This, Unleashing Docker for HPC. Cray User Group Conference (CUG'15). (2015). Cray User Group 2015.
- [18] Canonical Group Ltd. 2018. <https://jujucharms.com>. (Last accessed, Mar 2018).
- [19] Lustre 2018. Lustre file system. <http://lustre.org/>. (Last accessed, Mar 2018).
- [20] Mesos 2018. Apache Mesos. <http://mesos.apache.org>. (Last accessed, Mar 2018).
- [21] Paulius Micikevicius, Sharan Narang, Jonah Alben, Gregory F. Diamos, Erich Elsen, David Garcia, Boris Ginsburg, Michael Houston, Oleksii Kuchaiev, Ganesh Venkatesh, and Hao Wu. 2017. Mixed Precision Training. *CoRR abs/1710.03740* (2017), 1–12. <http://arxiv.org/abs/1710.03740>
- [22] Nectar 2018. Nectar Cloud. <https://nectar.org.au/>. (Last accessed, Mar 2018).
- [23] OCCI 2018. Open Cloud Computing Interface. <http://occi-wg.org/>. (Last accessed, Mar 2018).
- [24] Canada Department of Finance. 2018. Growing Canada's Advantage in Artificial Intelligence. <http://www.fin.gc.ca/n17/17-026-eng.asp>. (Last accessed, Mar 2018).
- [25] Singularity 2018. <http://singularity.lbl.gov/>. (Last accessed, Mar 2018).
- [26] Concetto Spampinato, Simone Palazzo, Pierluca D'Oro, Francesca Murabito, Daniela Giordano, and Mubarak Shah. 2018. VOS-GAN: Adversarial Learning of Visual-Temporal Dynamics for Unsupervised Dense Prediction in Videos. *CoRR abs/1803.09092* (March 2018), 1–18. <https://arxiv.org/abs/1803.09092>
- [27] Concetto Spampinato, Simone Palazzo, Daniela Giordano, Marco Aldinucci, and Rosalia Leonardi. 2017. Deep learning for automated skeletal bone age assessment in X-ray images. *Medical Image Analysis* 36 (2017), 41–51. DOI: <http://dx.doi.org/10.1016/j.media.2016.10.010>
- [28] Joseph E. Stiglitz. 2018. America Has a Monopoly Problem-and It's Huge. <https://www.thenation.com/article/america-has-a-monopoly-problem-and-its-huge/>. (Last accessed, Mar 2018).
- [29] David J. Teece. 2010. Business models, business strategy, and innovation. *Long Range Planning* 43, 2-3 (2010), 172–194.
- [30] The Linux Foundation Project 2018. AcumosAI. <https://www.acumos.org/platform/>. (Last accessed, Mar 2018).
- [31] Toreador 2018. TOREADOR project: Trustworthy model-aware Analytics Data platform. (Last accessed, Mar 2018).