

## CONFERENZA GARR 2022

18-20 maggio 2022 - Palermo

**CondiVisioni**

La rete come strumento per  
costruire il futuro

# Da HPC4AI al living lab dello spoke FutureHPC del CN HPC

**una storia dove non si è fatto altro che scopiazzare soluzioni  
intelligenti che altri hanno avuto ...**



UNIVERSITÀ  
DEGLI STUDI  
DI TORINO

Marco Aldinucci, Dip. Informatica, Università degli Studi di Torino  
Direttore del lab nazionale "HPC Key Technologies and Tools" del CINI



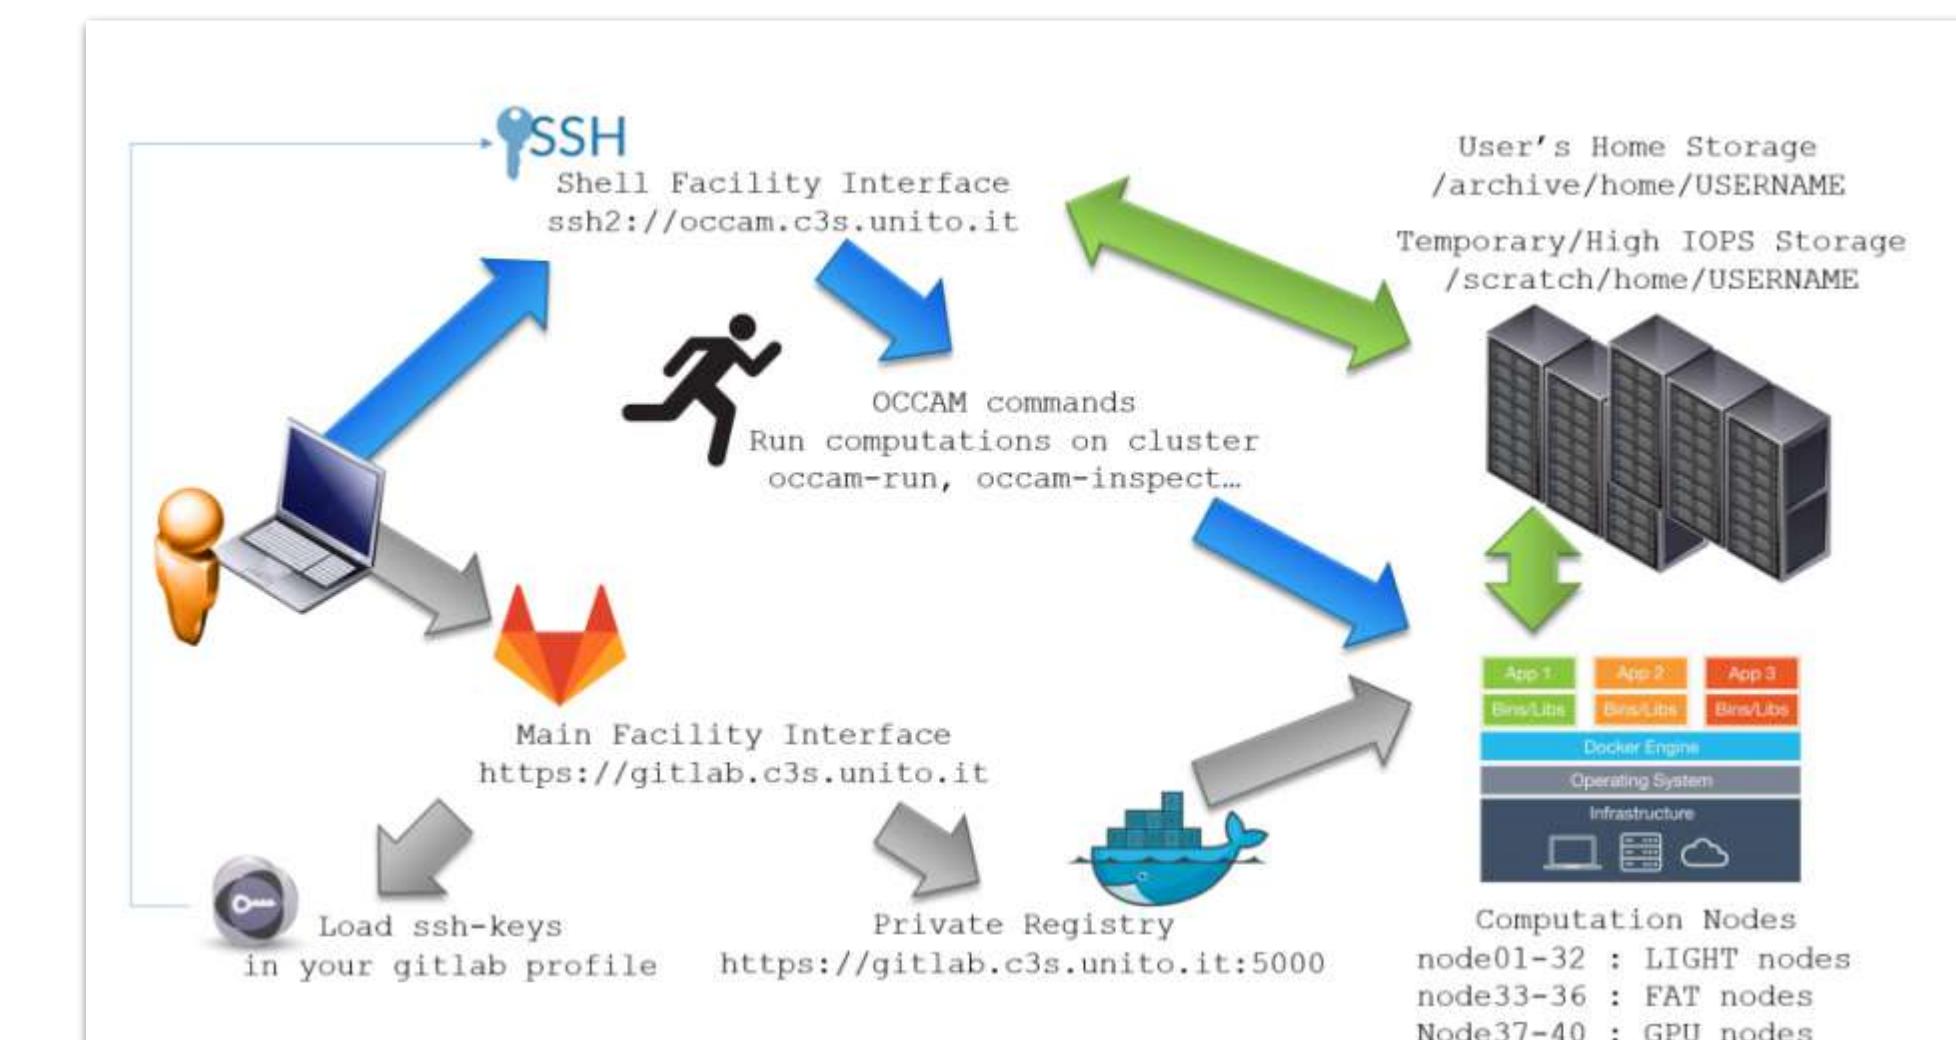
# **Condividere il nostro percorso**

## **dove le macchine sono solo il mezzo per concentrarci sulle persone**

- Mantenere e anzi accrescere la competenza sulle tecnologie
  - C<sub>3</sub>S e OCCAM
- Modernizzare le applicazioni: la portabilità, la composizionalità e il cloud
  - HPC4AI, il nostro laboratorio per la convergenza HPC-cloud-AI
- Mettere insieme le competenze
  - Il laboratorio HPC KTT del CINI
- Costruire una filiera nazionale del software
  - Il centro nazionale HPC del PNRR e il living lab “Software & Integration”

# C3S e OCCAM

## Il Centro di Competenza su Scientific Computing a UNITO (2016)



Bello Prof, ma si può entrare dentro come in HAL9000?

**METTIAMOCI DENTRO LE  
PERSONE !**



# Olivetti Elea9001 and Ettore Sottsass



Barbaricina, Pisa

# Olivetti Elea9001 and Ettore Sottsass

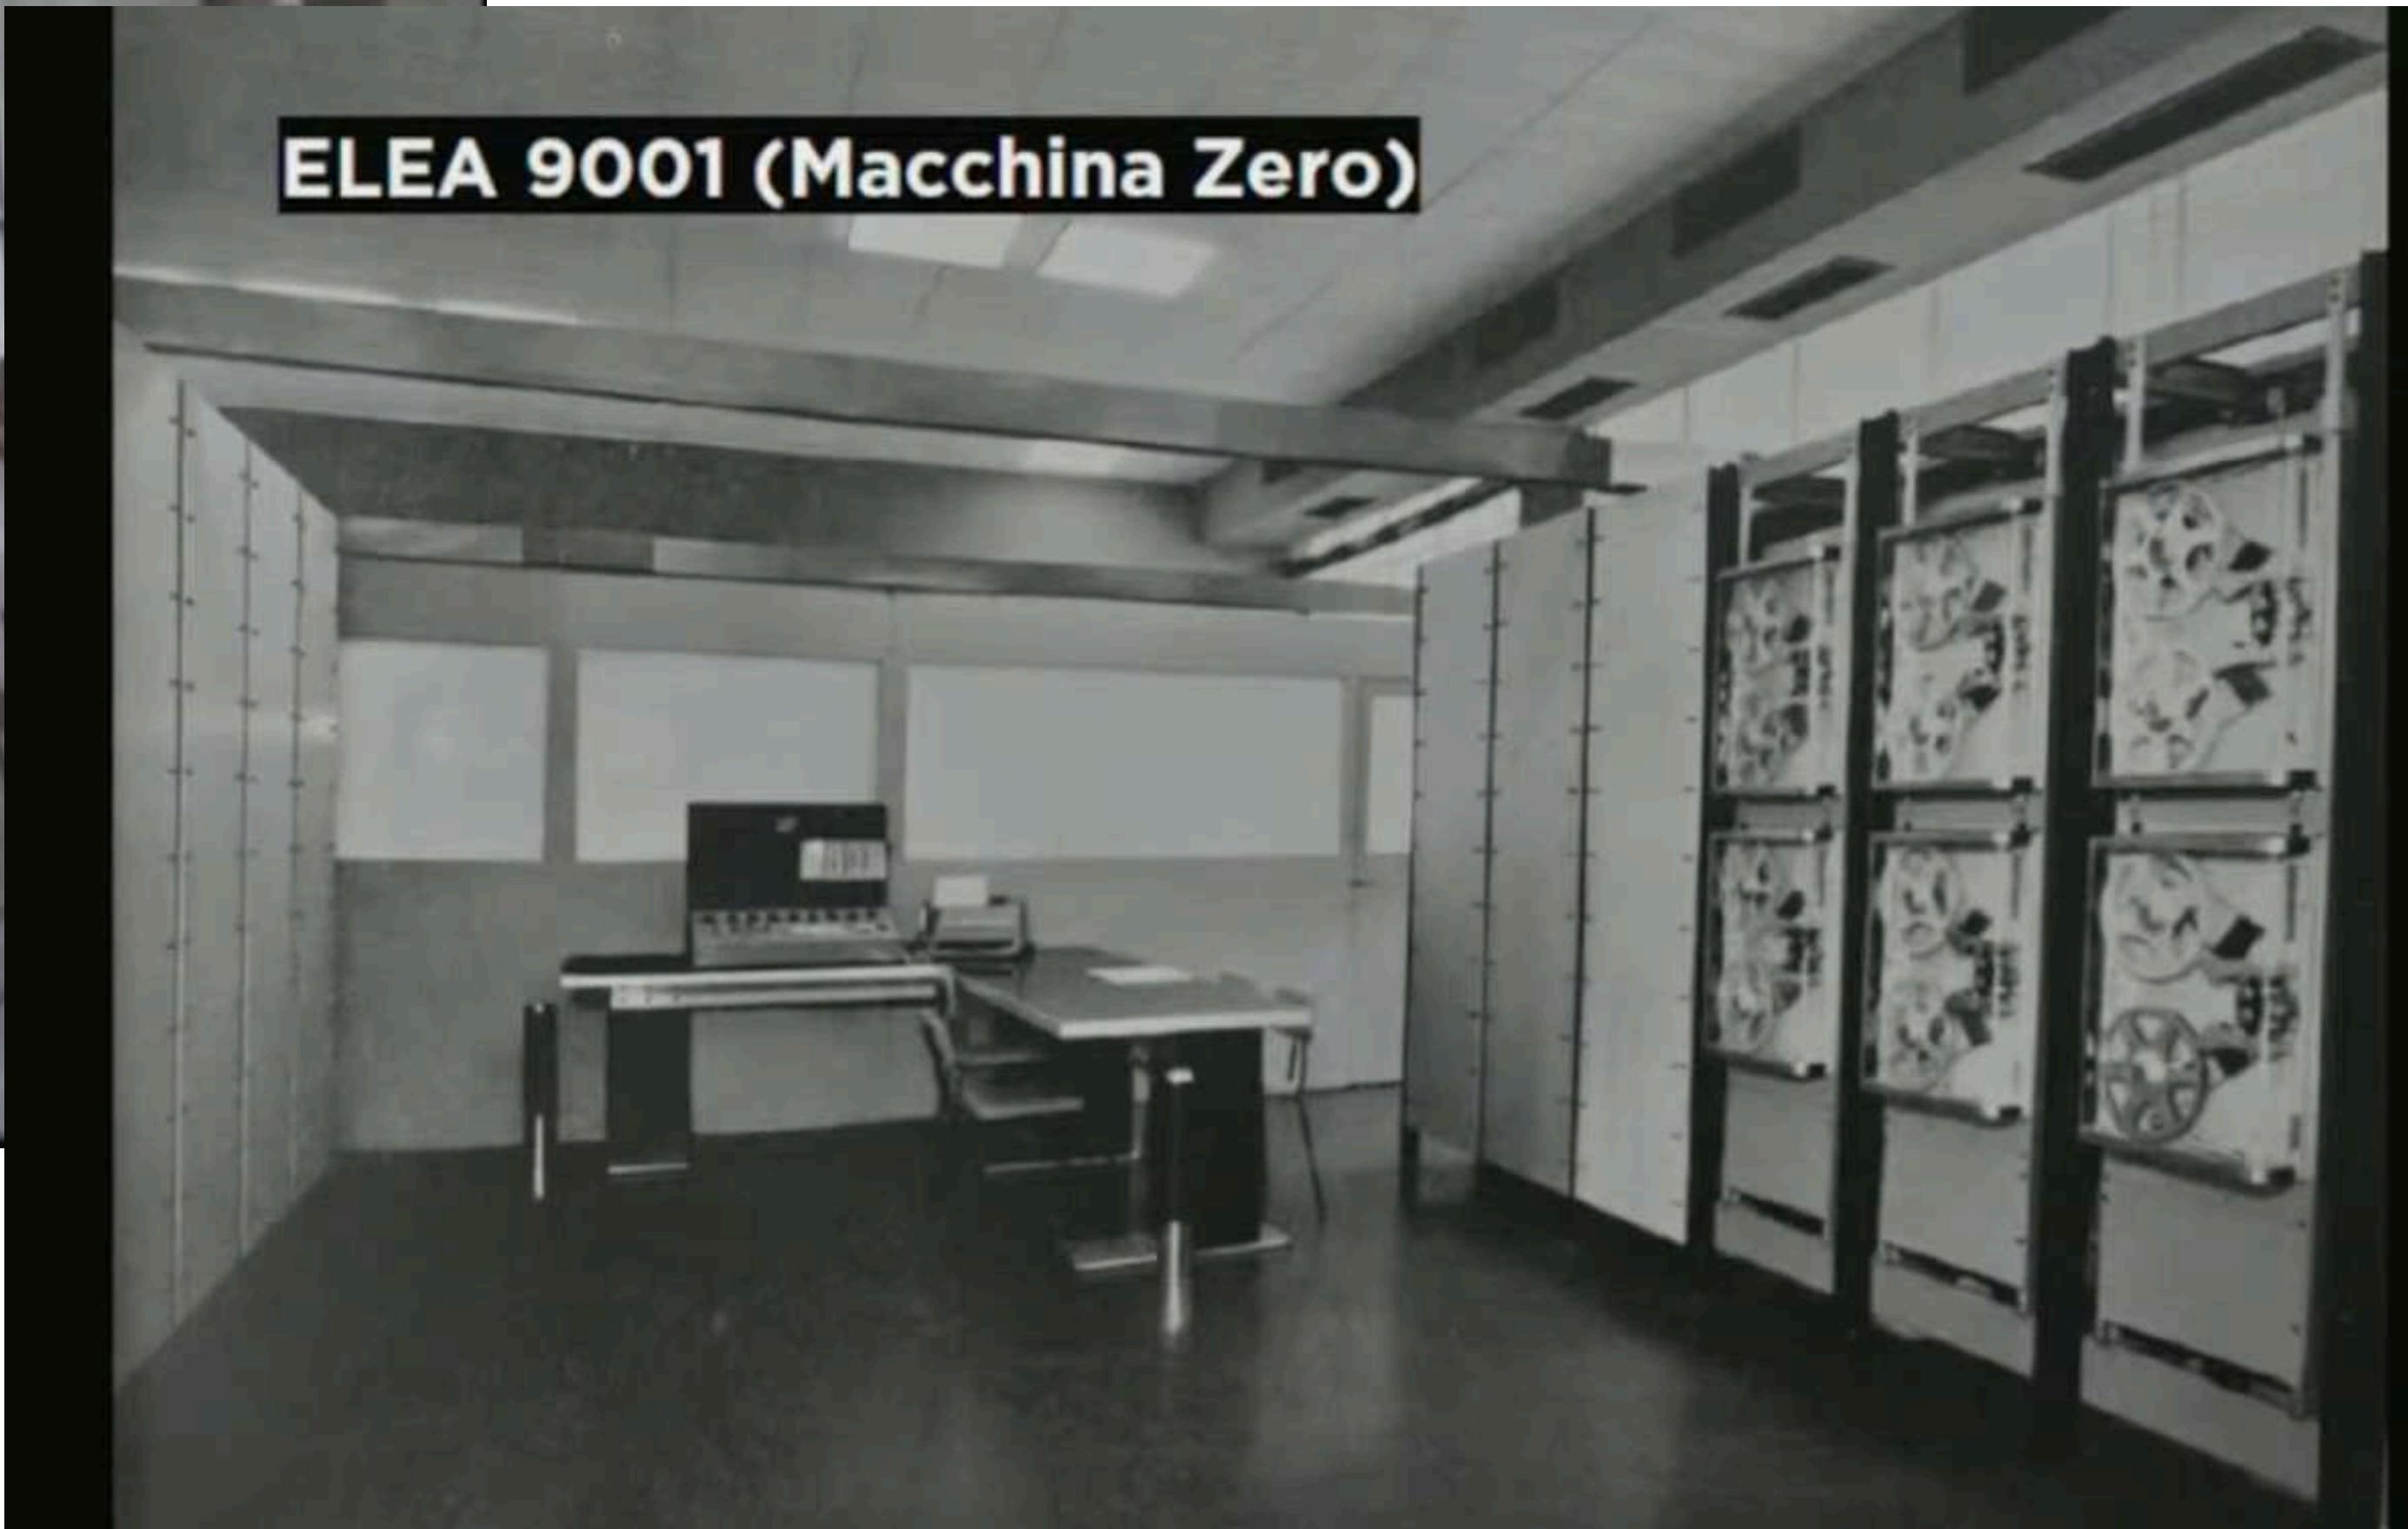
Macchina Zero



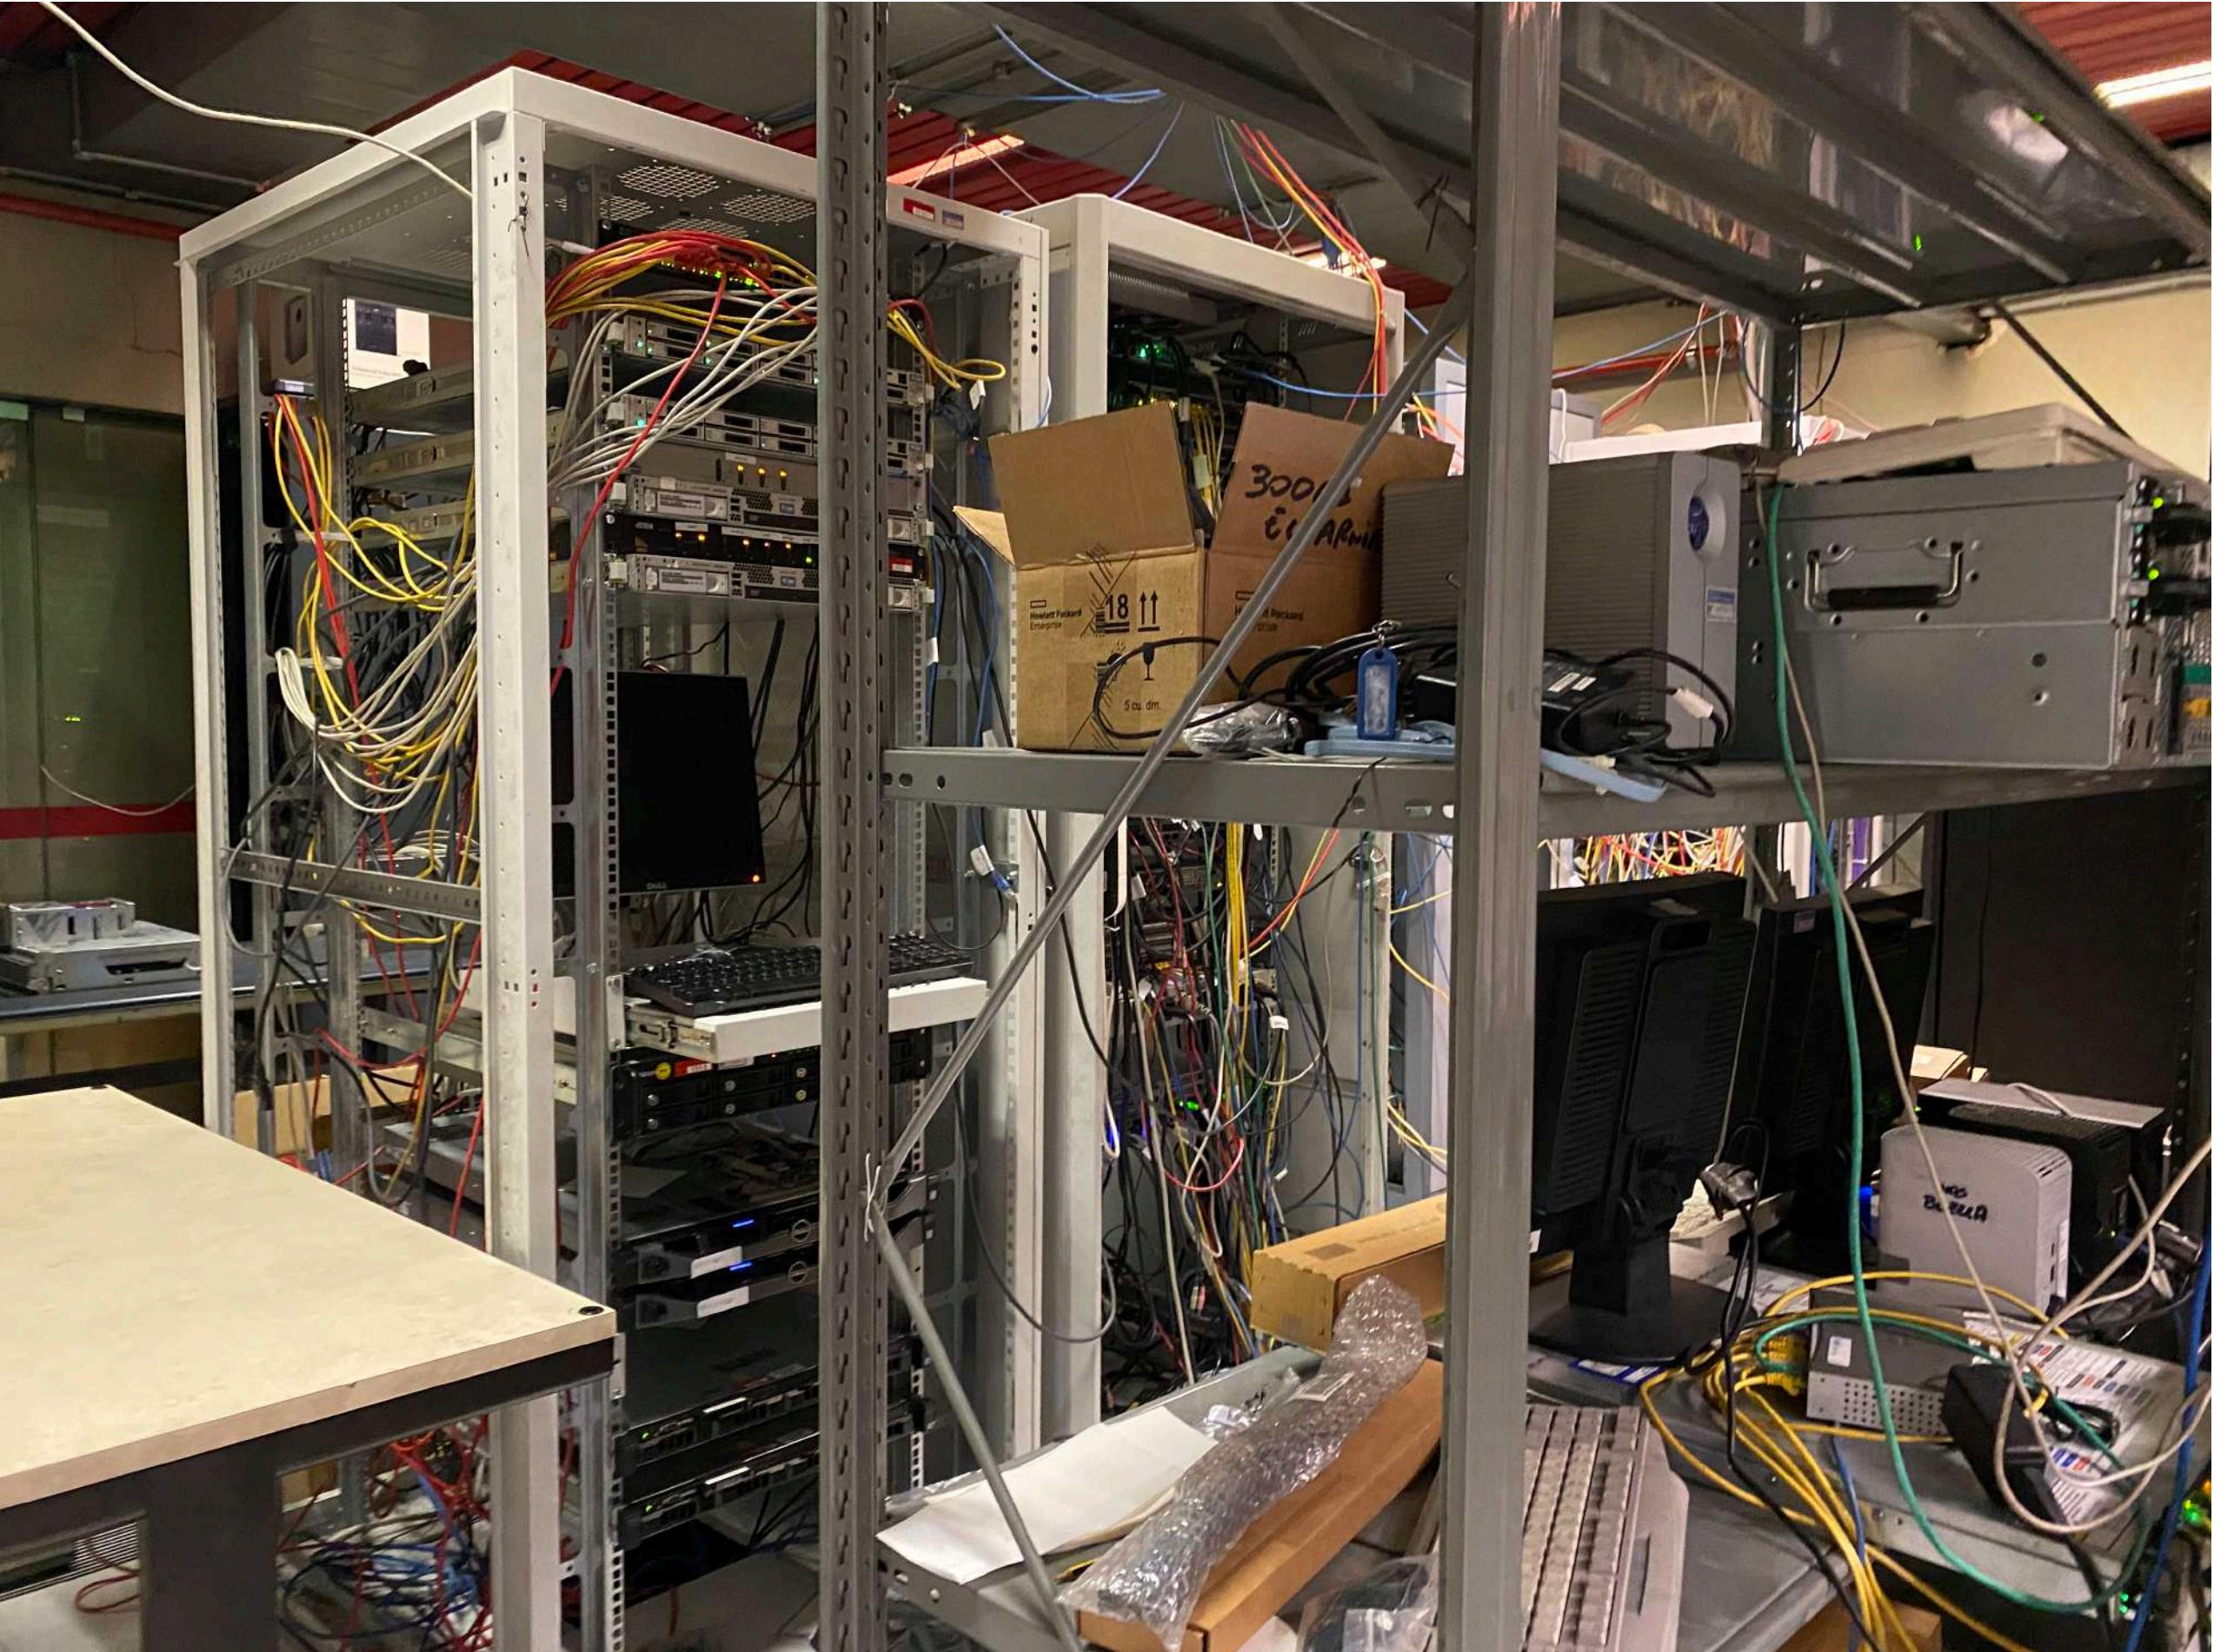
Barbaricina, Pisa

Ivrea

ELEA 9001 (Macchina Zero)

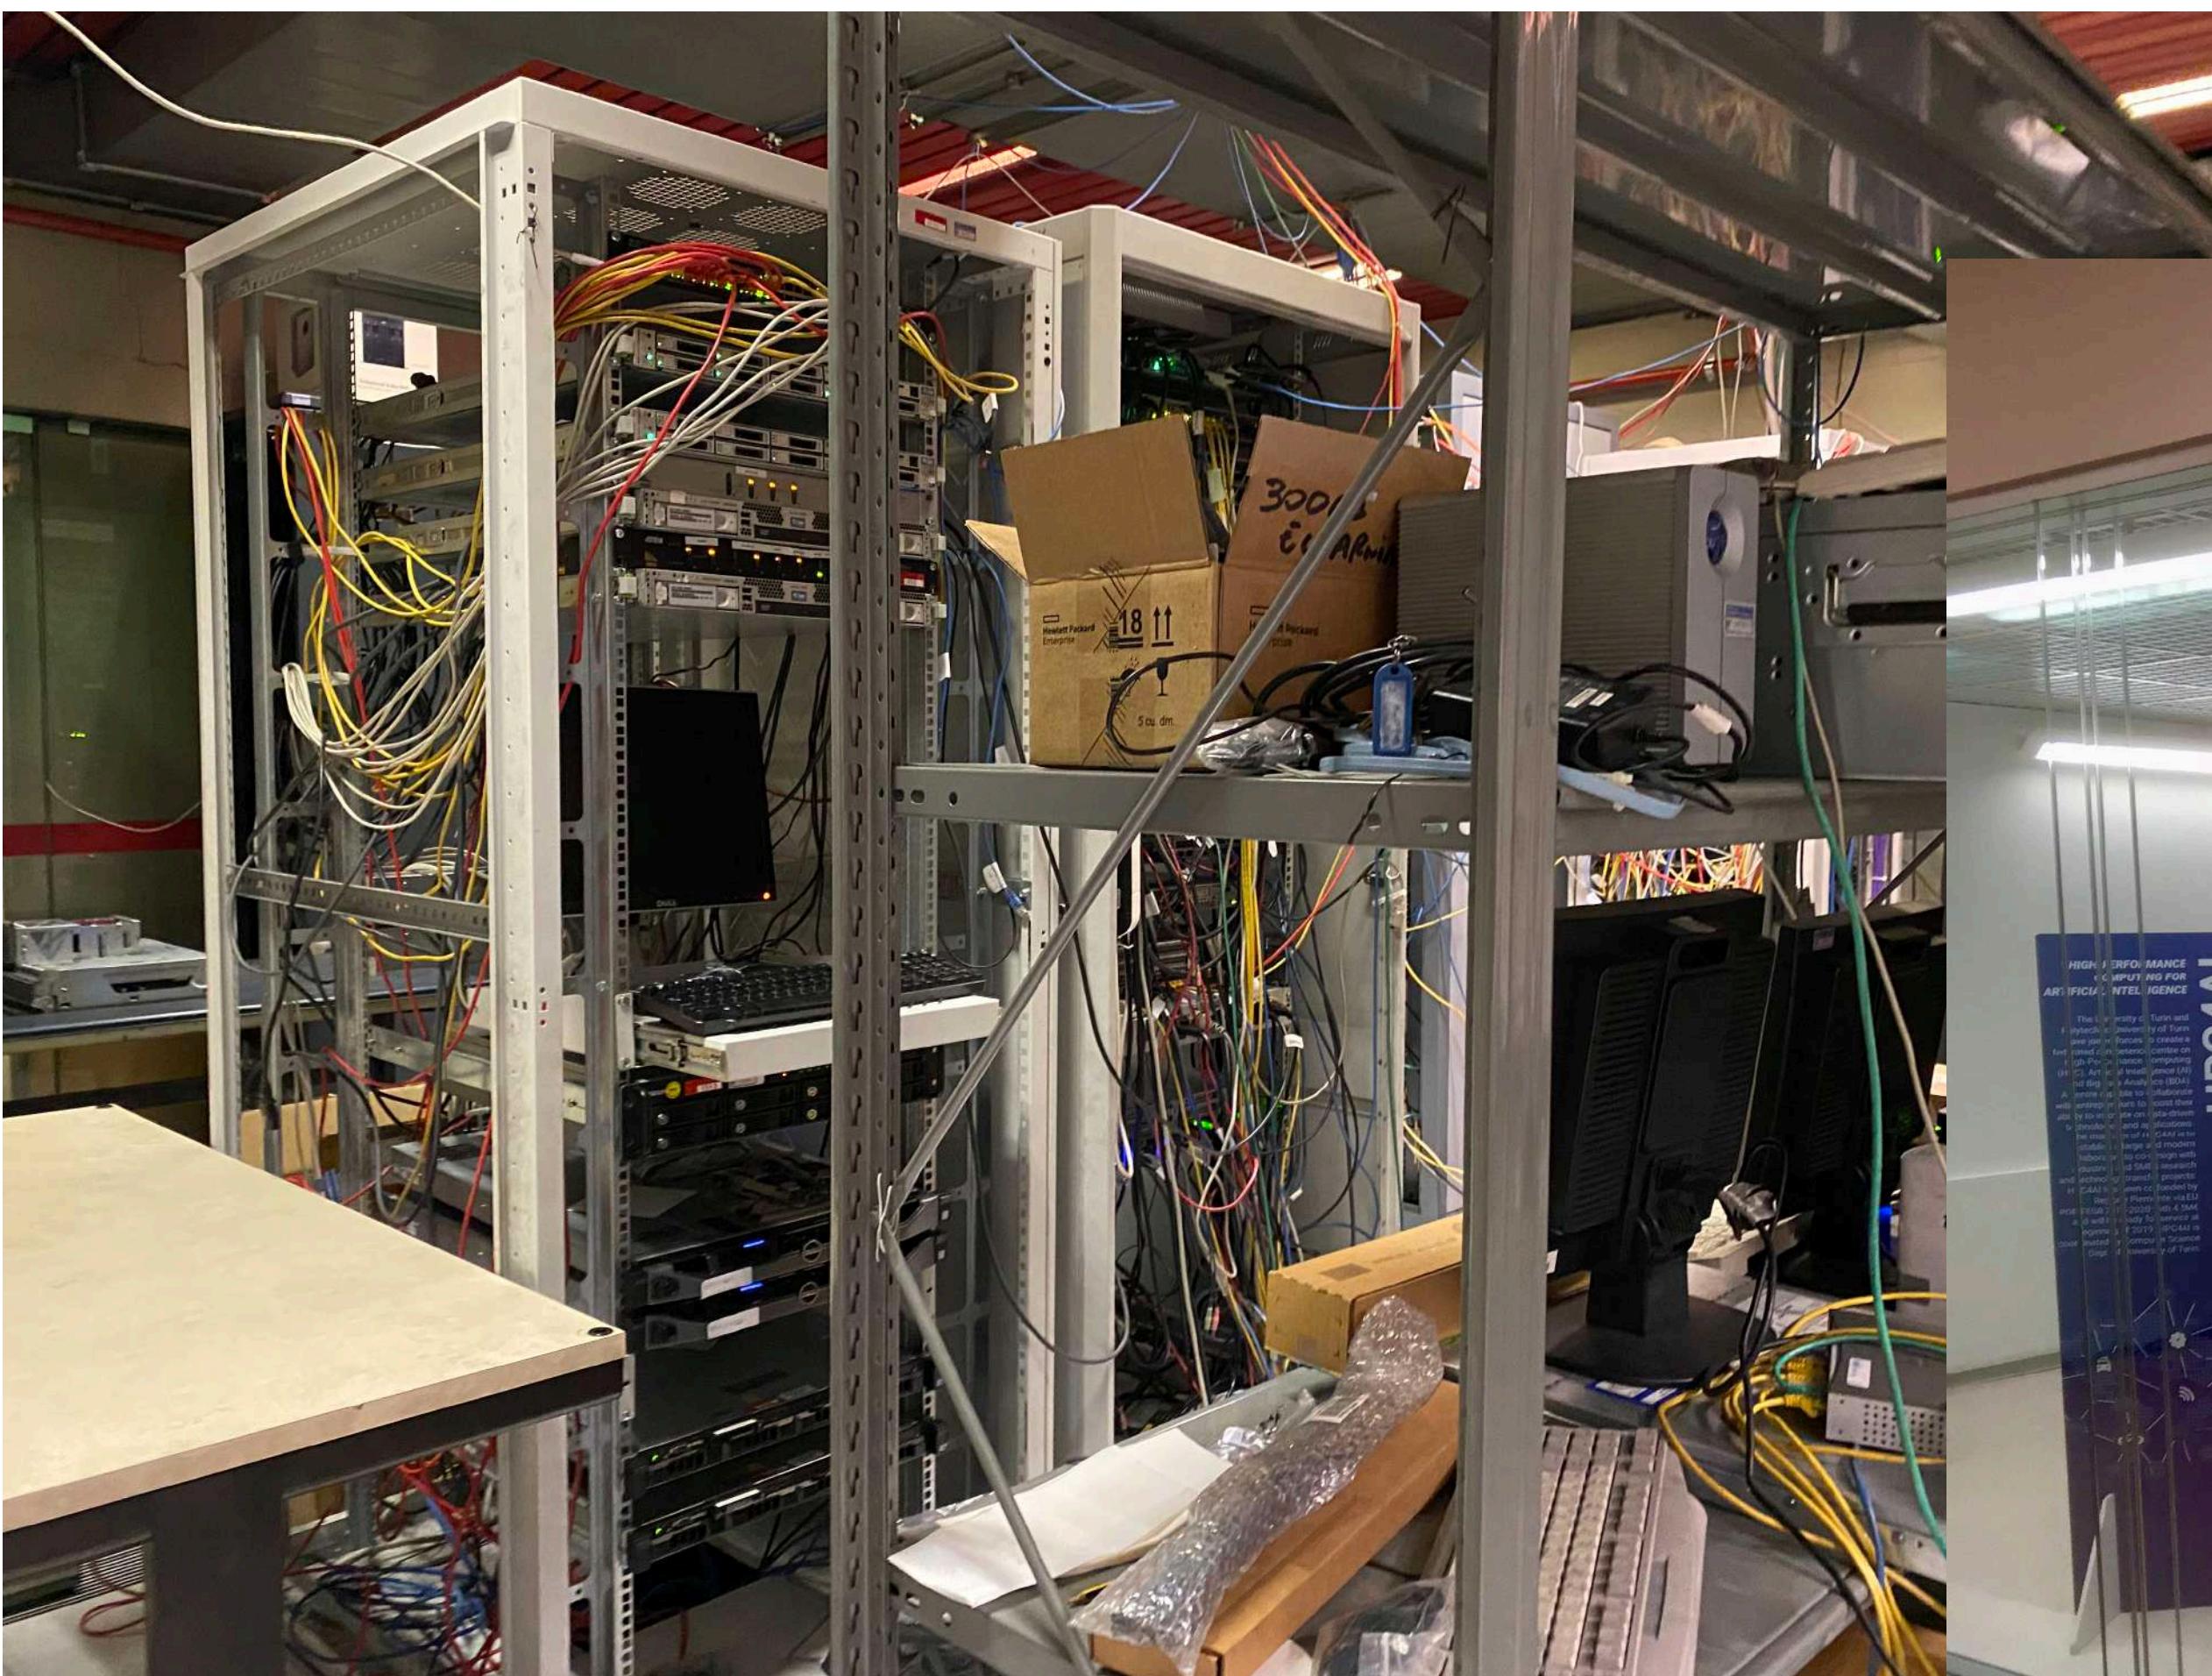


# Dip. Informatica and ...



Piero della Francesca, Torino  
(1 piano)

# Dip. Informatica and ...



Piero della Francesca, Torino  
(1 piano)

Piero della Francesca, Torino  
(3 piano)



# Ettore Sottsass for l'Olivetti 9003

"Il paesaggio elettronico", Compasso d'Oro 1959





<https://hpc4ai.unito.it>

Two systems “**DevOps**” (Production+Dev) each made by two clusters  
1 Cloud/OpenStack + 1 HPC/SLURM

Overall **11k+ cores** (Intel, ADM), **120 GPUs** (T4/V100/A100), 4 novel **Arm+A100+BlueField2** platforms, **RISCV Sifive**  
3PB in 5 systems: CEPH (SSD/rep3, Rot/ErasureCode), EMC2 (Unity+Avamar)

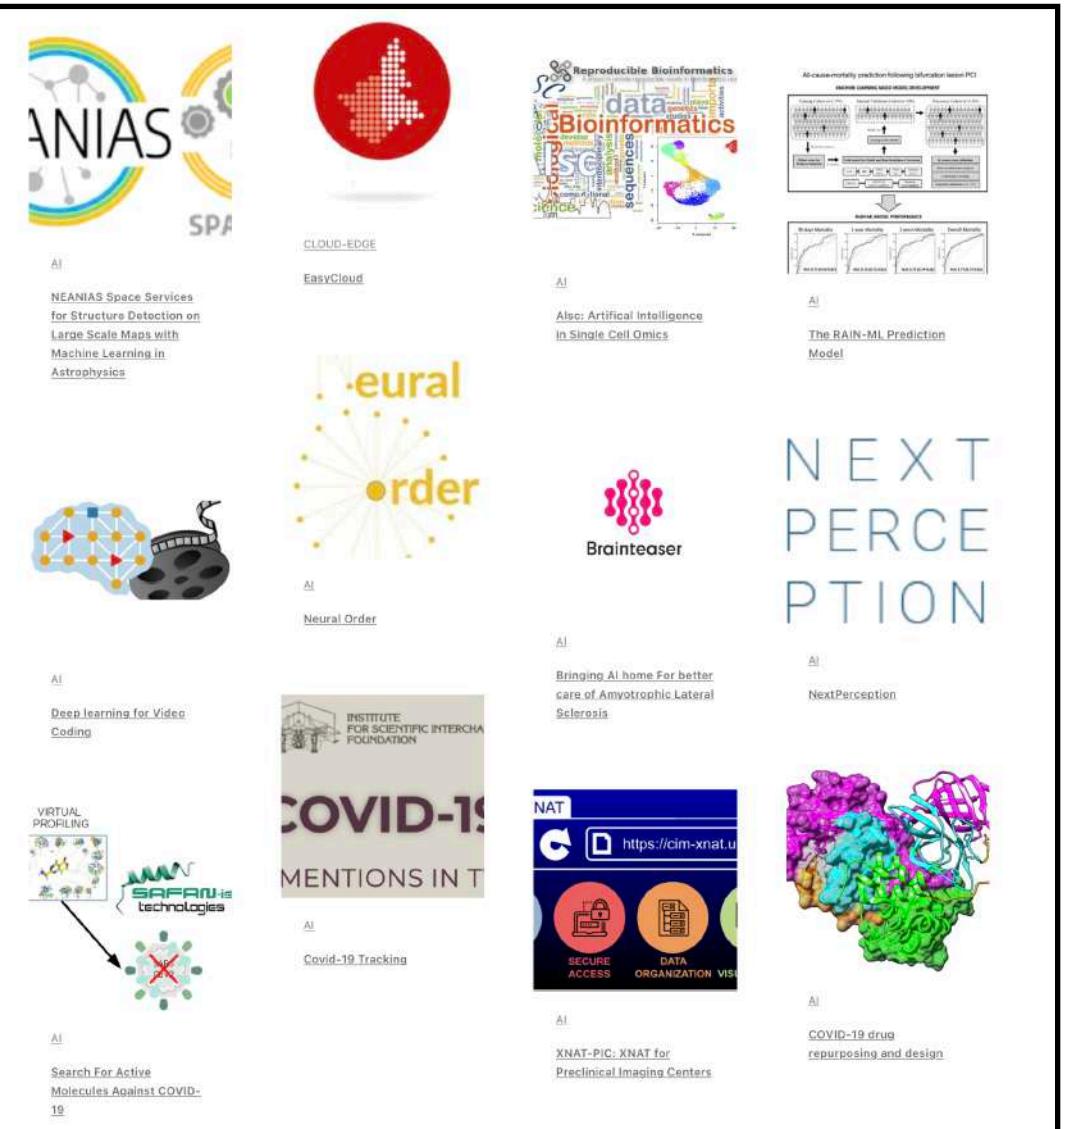
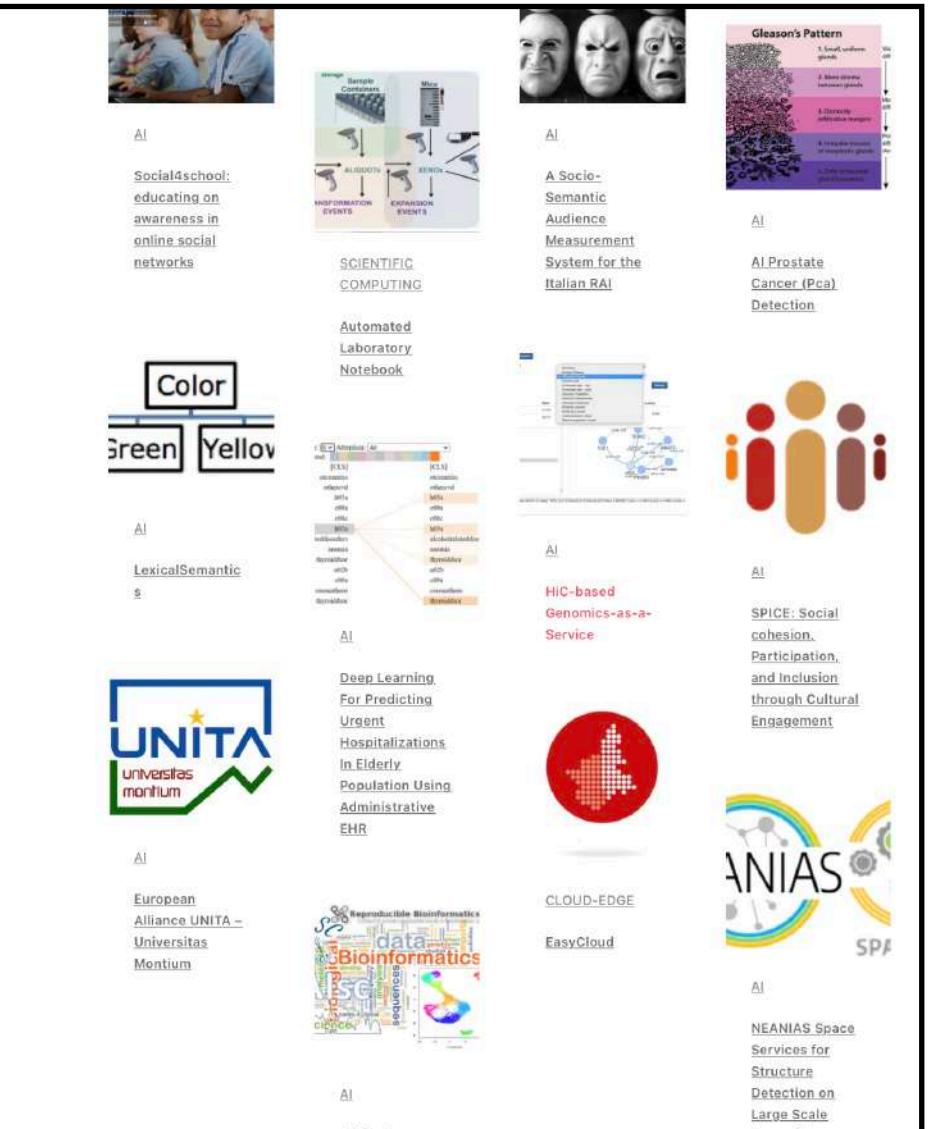
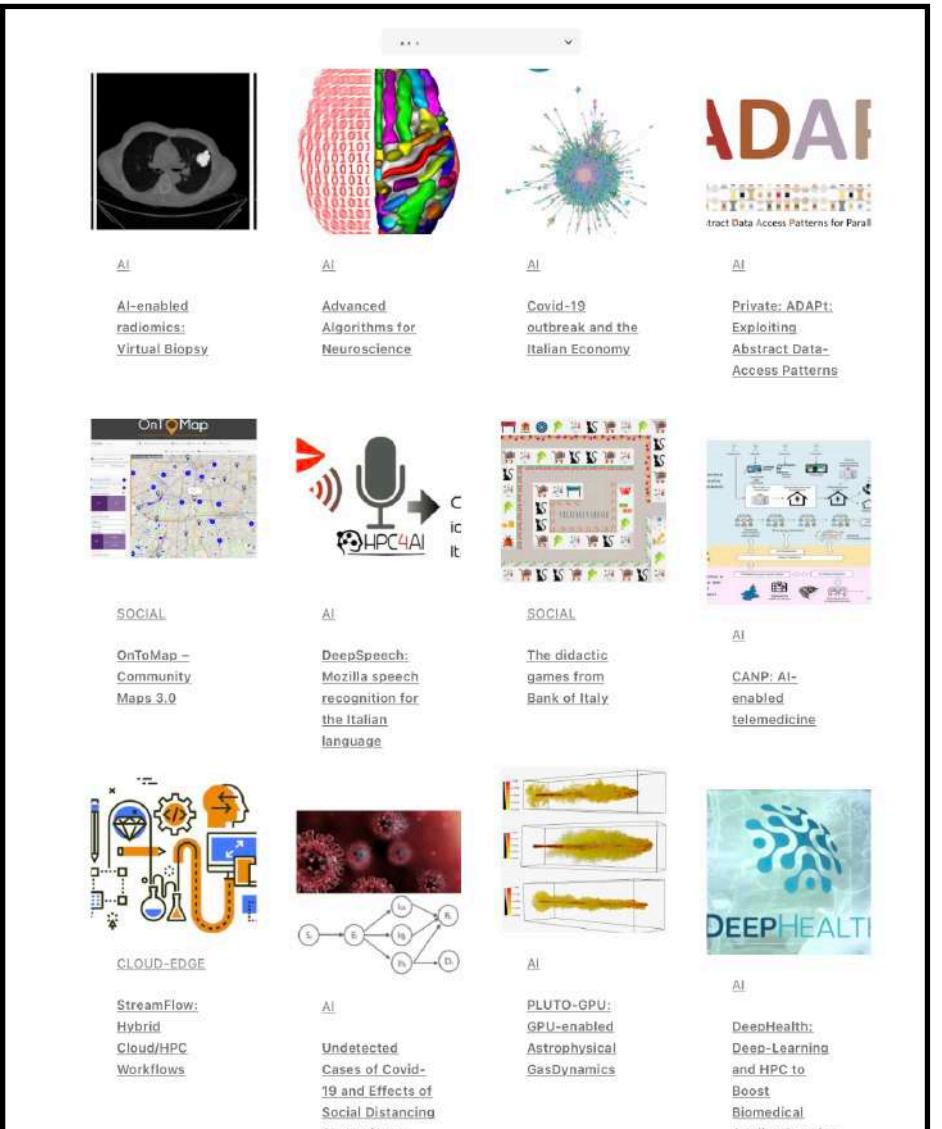
**60+ research projects running** (HPC, AI, ...), mostly with critical data (medical)

Several in-house PaaS: SLURM-calendar integration, Kubernetes-as-a-Service, **Jupyter-as-a-Service**, ...

Our “own” **250KVA Tier-3** datacenter - 1.04 P.PUE design (**PUE 1.1/year** measured)

Production system: 21 months with 0 “system-down” incidents

**We develop it and manage it!** Our “products” is knowledge and skilled students



# E4

COMPUTER  
ENGINEERING

HPC & Enterprise Solutions

E4 Computer Engineering SpA  
Via Martiri della Libertà, 66  
42019 Scandiano . Reggio Emilia . Italy

[www.e4company.com](http://www.e4company.com)

Tel. + 39 0522 99 18 11  
Fax + 39 0522 99 18 03  
e-mail: [info@e4company.com](mailto:info@e4company.com)

Scandiano, February 2nd, 2021

To: Prof. Marco Aldinucci  
Computer Science Department  
University of Torino  
Corso Svizzera 185, 10149 Torino – Italy

Subject: Letter of support for the application to the “Bando di Ateneo per l’acquisizione di Piccole e medie attrezzature per attività di ricerca – Anno 2020”

Dear Prof. Aldinucci,

E4 Computer Engineering is honored to support the proposal that you and your team will submit to the “Bando di Ateneo per l’acquisizione di Piccole e medie attrezzature per attività di ricerca – Anno 2020”.

E4 is a member of the EU-funded European Processor Initiative SGA1 and has joined the consortium who submitted the second phase of the project (SGA2). Within SGA2, E4 is tasked to develop a certain number of prototypes, i.e. fully functional servers based on the Rhea Reference Board. The objective of the prototypes is the validate the first version of the EPI processor (Rhea) in a realistic, datacenter-like environment and gather key information for the design of the second generation processor. It is currently envisioned in the proposal submitted for SGA2 that a prototype will be allocated for a certain amount of time to UniTO. UniTO has been selected because of its high-level skillset and domain-wide expertise as a developer and tester. It is envisioned that the tests and assessments performed by UniTO will provide key feedback in the techniques of application porting and optimization, in the analysis of the behaviour of the processor in realistic workload conditions and in the validation of the development tools. The outcome of this work will be the first functional prototype of the second generation processor of EPI.

E4 declares there is no conflict of interest between this letter of support and the preparation of competing proposals. Sincerely yours,  
Francesco Morsiani

CEO and Founder  
E4 Computer Engineering SpA



scolo sanitario elettronico a livello nazionale. «La sanità ha accelerato

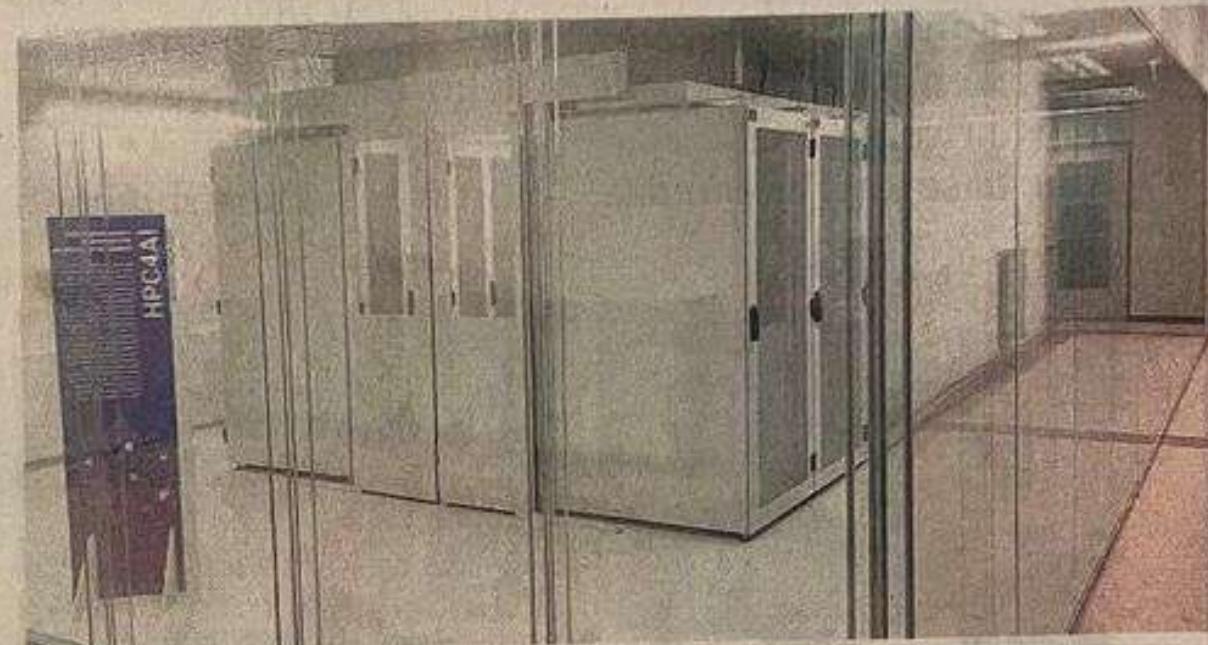
aziende sanitarie, società partecipate pubbliche, comuni e città metropolitane».

—F.Gre.  
© RIPRODUZIONE RISERVATA

# Università di Torino, il prototipo del supercalcolatore europeo

## Hpc Informatica e AI

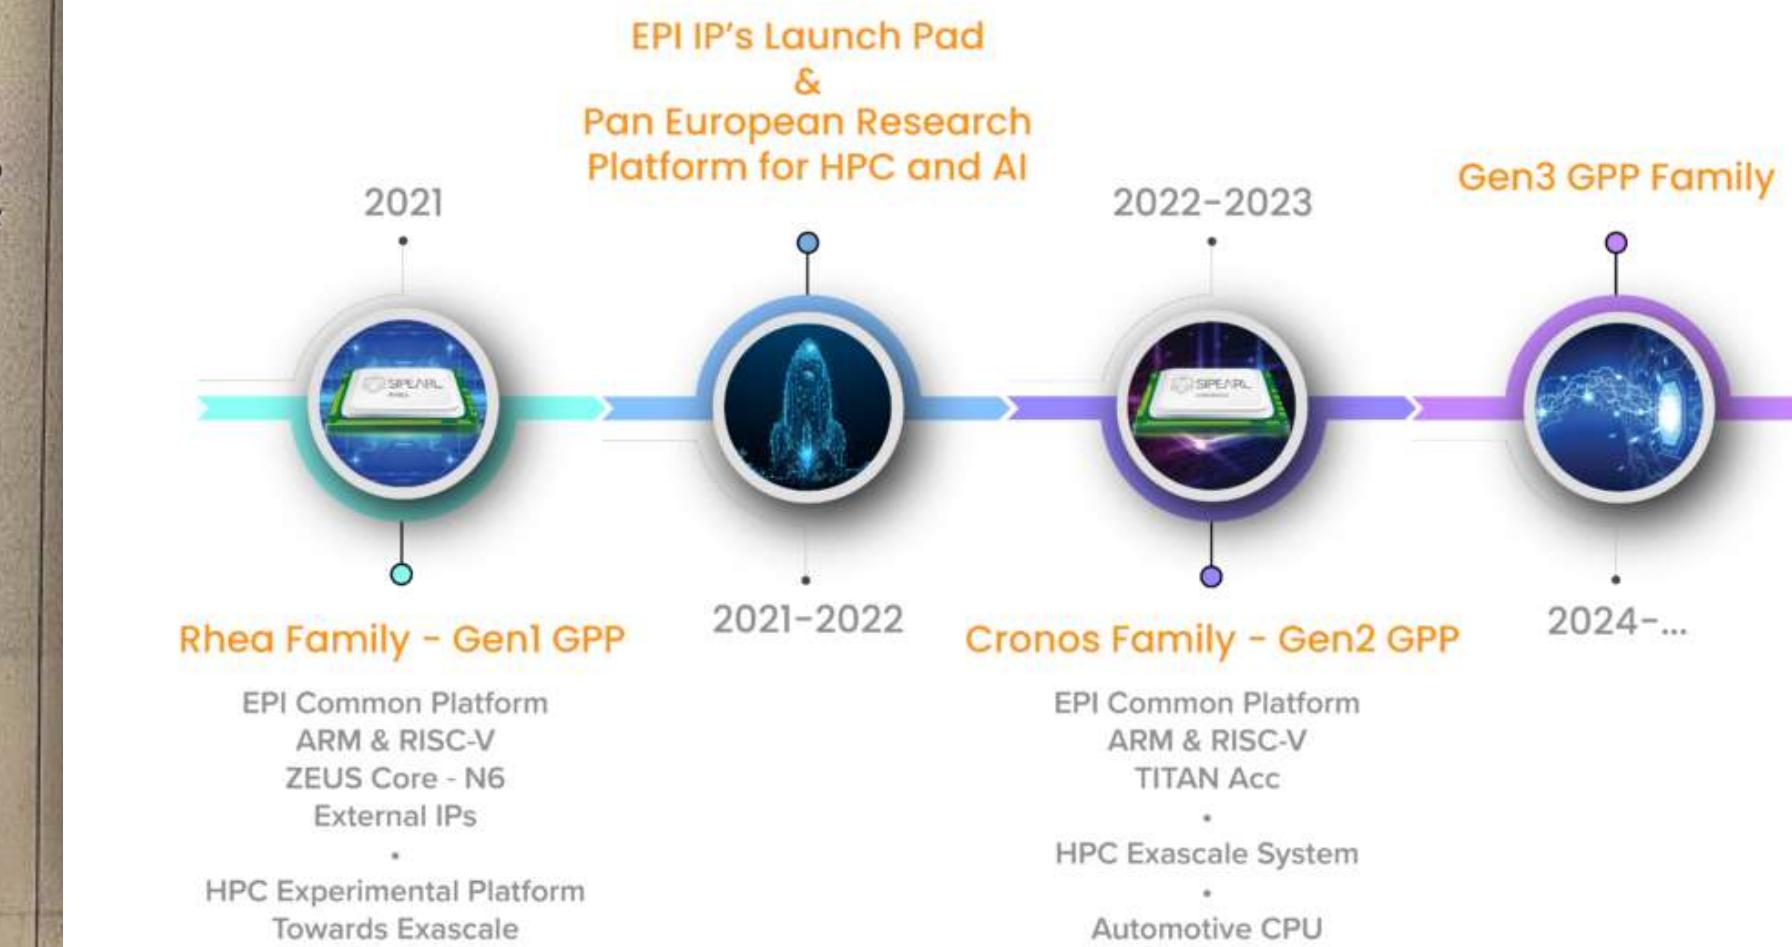
Università di Torino in prima fila nei progetti europei sul supercalcolo. «Siamo partiti un po’ in ritardo ma ora siamo inseriti nella rete di ricerca europea più avanzata» sottolinea Marco Aldinucci, responsabile del Dipartimento di Informatica dell’Unito. Punto di partenza è stato lo sviluppo, insieme al Politecnico di Torino, di HPC4AI, sistema di calcolo avanzato pensato per le applicazioni di intelligenza artificiale, attualmente in uso a ricercatori e imprese. «Abbiamo realizzato un Data Center – spiega Aldinucci – che ospita sistemi interamente



**Big data.** Il Data Center all’Università degli Studi di Torino

Il prossimo challenge, “exascale” (10 milioni di volte più veloce di un ordinario laptop), è atteso tra meno di tre anni. «L’Ue in questo momento storico se la sta giocando, e non era

ta dello sviluppo di sistemi di analisi ed elaborazione di dati provenienti da fonti diverse senza che i dati stessi debbano essere spostati insieme in un solo posto (data lake). Un modello



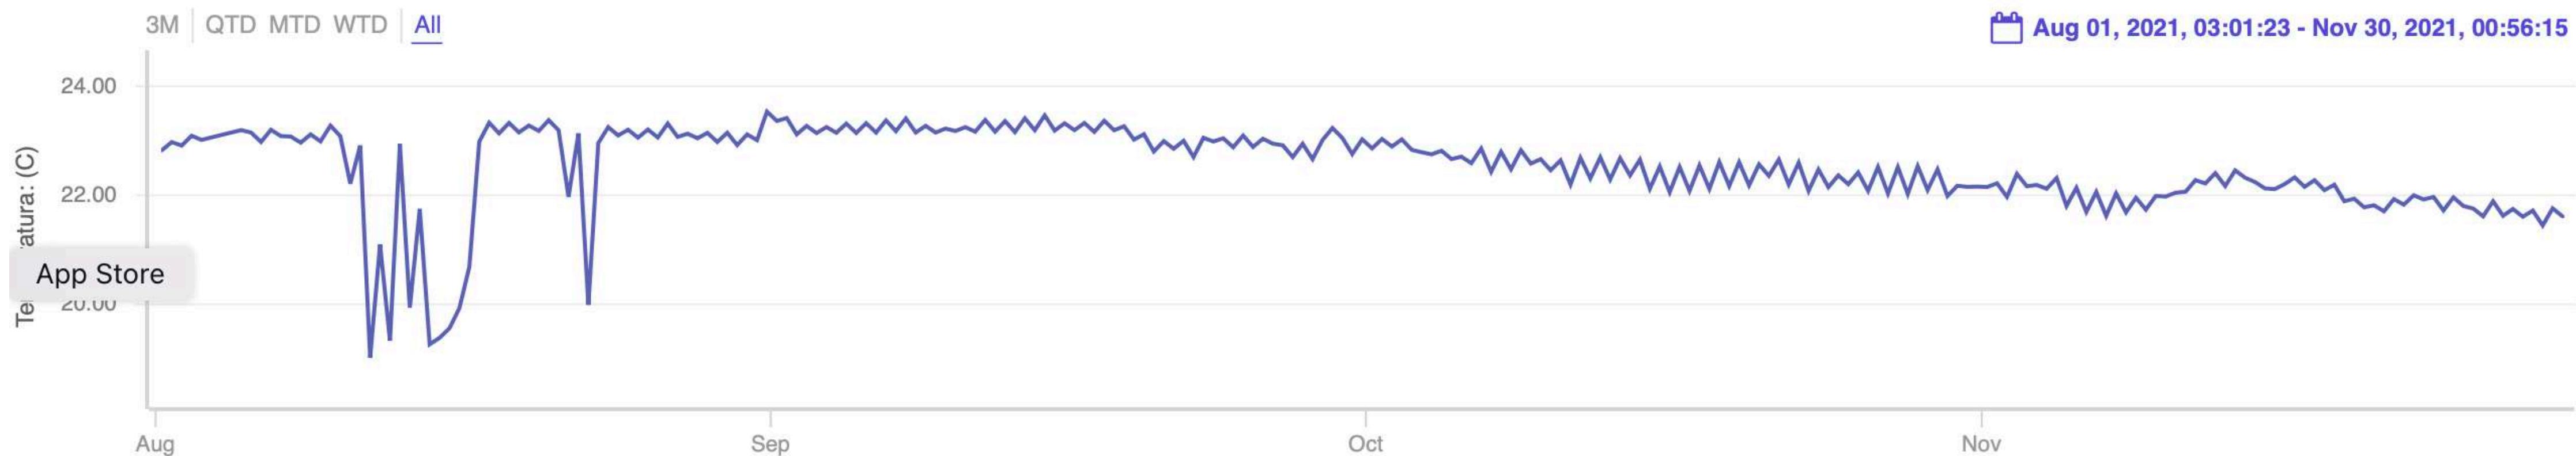
percalcolatori è una vera e propria corsa, con successivi livelli di potenza che aumentando di mille in mille,

a cui l’Università sta lavorando c’è il Federated Learning, che ha grande rilevanza a livello industriale. Si tratta

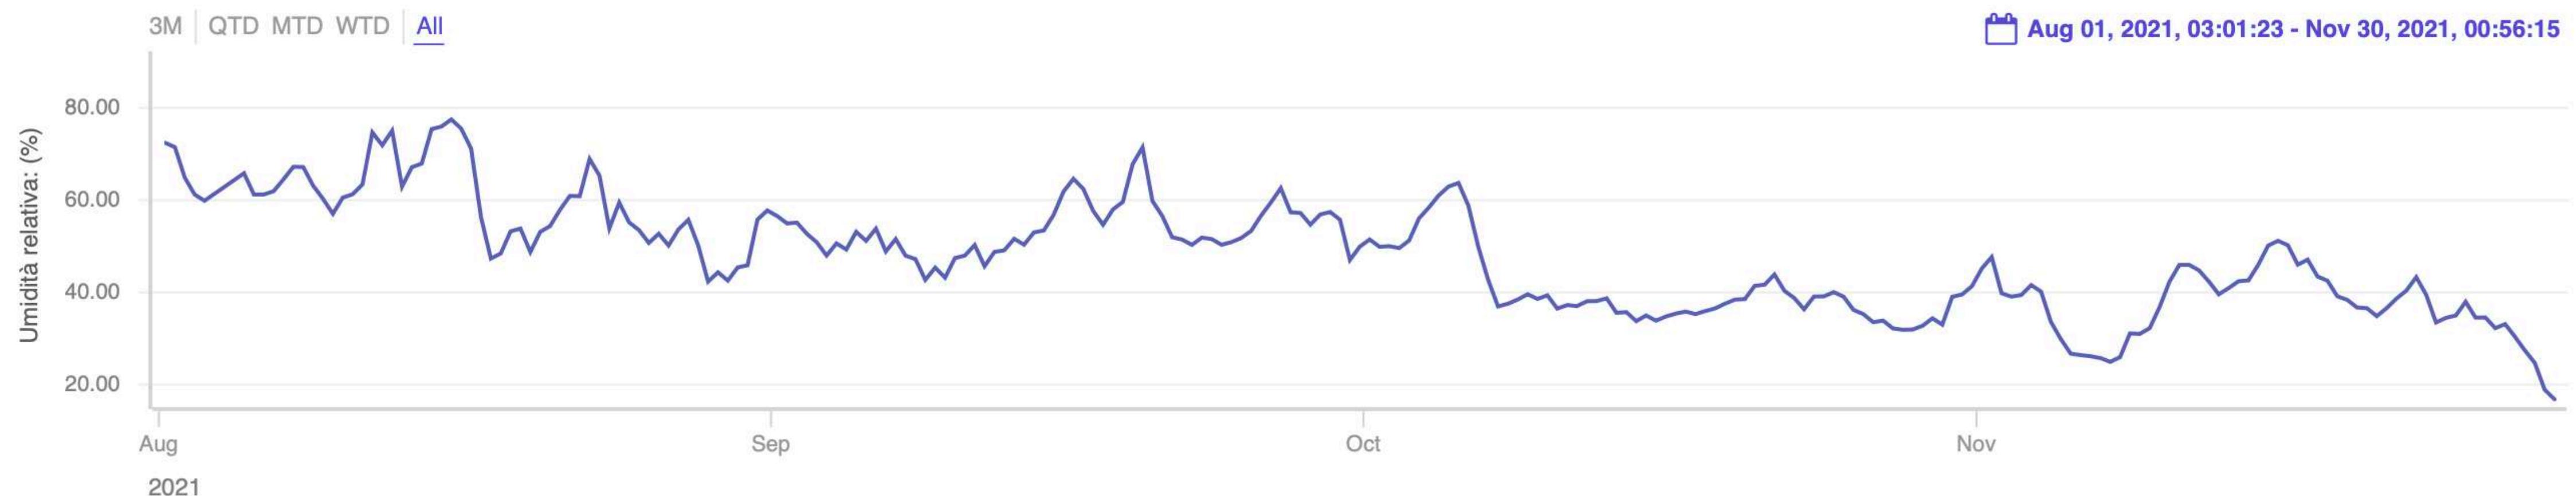
della catena del valore di molte filiere».

—F.Gre.

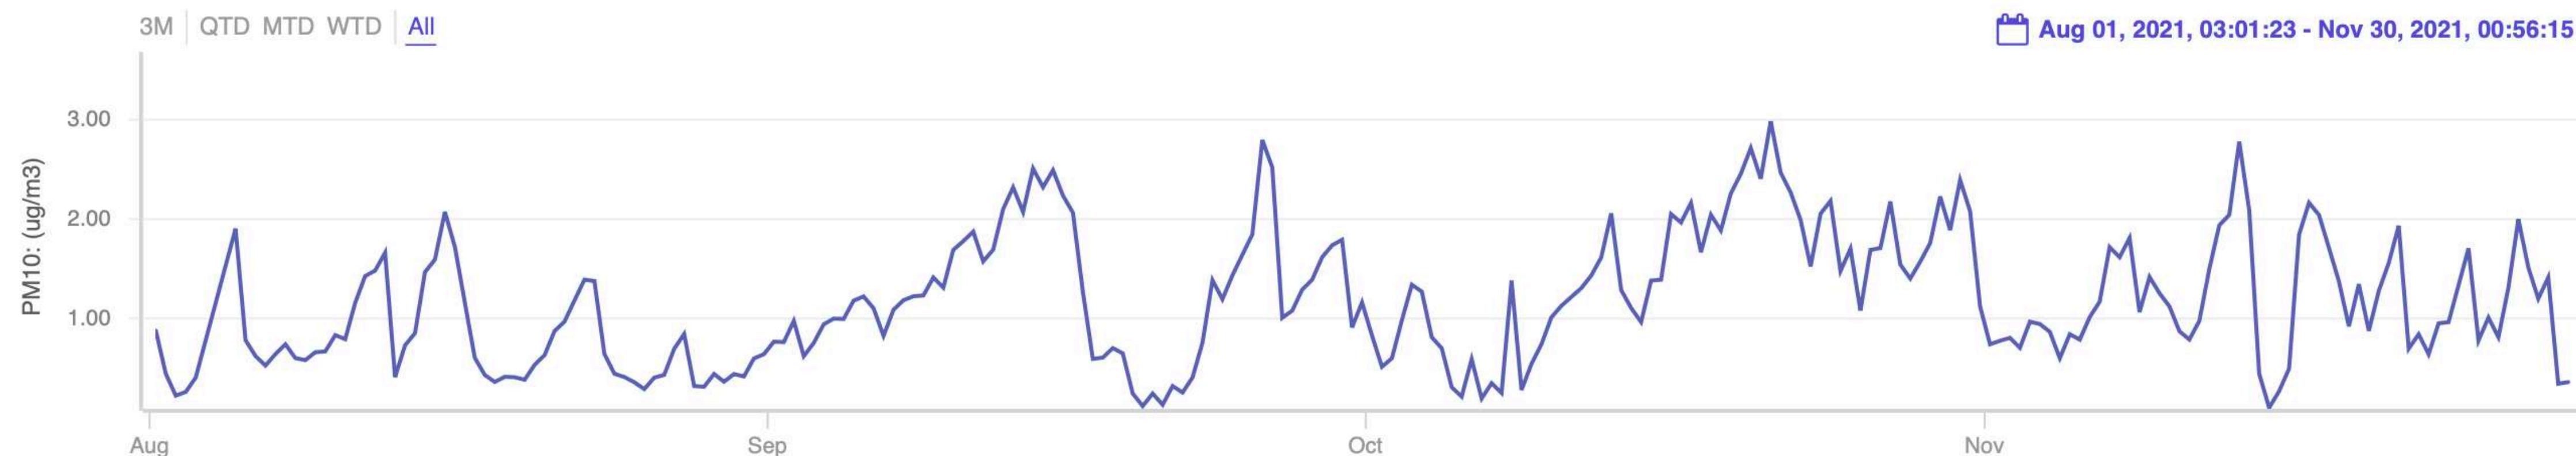
## Temperatura



## Umidità relativa



## PM10



## Temperatura



## Umidità relativa



## PM10



# Dip. Informatica UNITO - Piano TERZO - CED

Online

Temperatura: 21.62 C

021, 03:01:23 - Nov 30, 2021, 00:56:15

Temperatura

Umidità relativa

Pressione assoluta

Pressione relativa

Intensità luminosa

Rumore

VOC in CO<sub>2</sub> Equivalente

VOCT

CO<sub>2</sub>

NH<sub>3</sub>

NO<sub>2</sub>

CO

O<sub>3</sub>

PM1

PM2,5

PM4

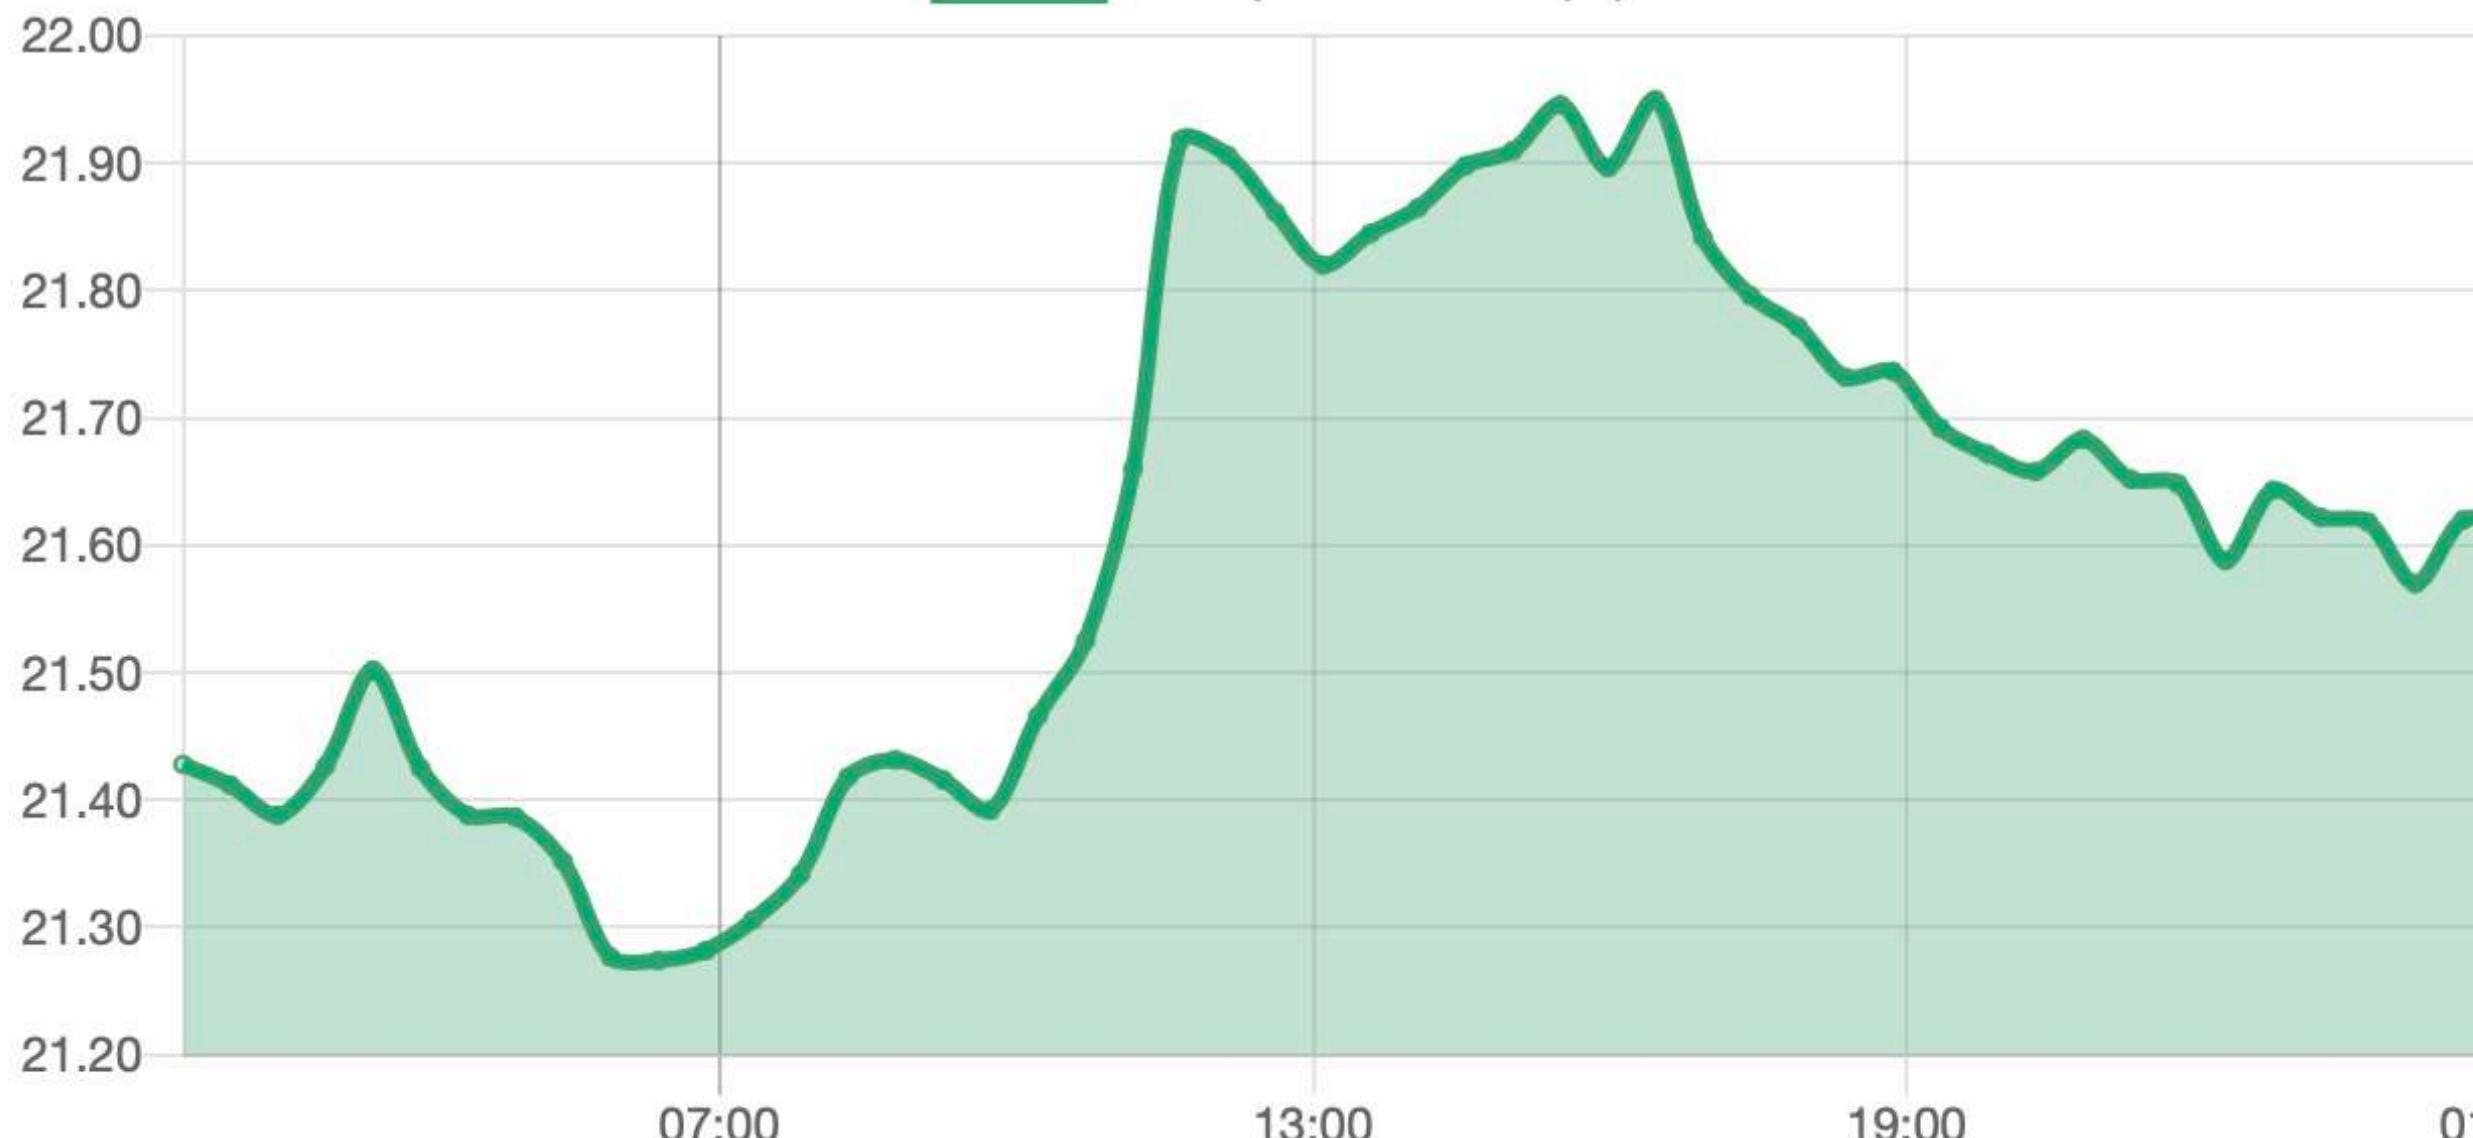
PM10

Magnetometro

Pressione

Temperatura - (C)

Temperatura - (C)



Ultima misura: 30/11/21, 01:01

021, 03:01:23 - Nov 30, 2021, 00:56:15

021, 03:01:23 - Nov 30, 2021, 00:56:15

021, 03:01:23 - Nov 30, 2021, 00:56:15

### Machine learning-based prediction of adverse events following an acute coronary syndrome (PRAISE): a modelling study of pooled datasets

Fabrizio D'Ascenzo, Ovidio De Filippo, Guglielmo Gallone, Gianluca Mittone, Marco Agostino Deriu, Mario Iannaccone, Albert Ariza-Solé, Christoph Liebtrau, Sergio Manzano-Fernández, Giorgio Quadri, Tim Kinnaird, Gianluca Campo, Jose Paulo Simao Henrques, James M Hughes, Alberto Dominguez-Rodriguez, Marco Aldinucci, Umberto Morbiducci, Giuseppe Patti, Sergio Raposeiras-Roubin, Emad Abu-Assi, Gaetano Maria De Ferrari, on behalf of the PRAISE study group

#### Summary

**Background** The accuracy of current prediction tools for ischaemic and bleeding events after an acute coronary syndrome (ACS) remains insufficient for individualised patient management strategies. We developed a machine learning-based risk stratification model to predict all-cause death, recurrent acute myocardial infarction, and major bleeding after ACS.

**Methods** Different machine learning models for the prediction of 1-year post-discharge all-cause death, myocardial infarction, and major bleeding (defined as Bleeding Academic Research Consortium type 3 or 5) were trained on a cohort of 19826 adult patients with ACS (split into a training cohort [80%] and internal validation cohort [20%]) from the BleeMACS and RENAMI registries, which included patients across several continents. 25 clinical features routinely assessed at discharge were used to inform the models. The best-performing model for each study outcome (the PRAISE score) was tested in an external validation cohort of 3444 patients with ACS pooled from a randomised controlled trial and three prospective registries. Model performance was assessed according to a range of learning metrics including area under the receiver operating characteristic curve (AUC).

**Findings** The PRAISE score showed an AUC of 0.82 (95% CI 0.78–0.85) in the internal validation cohort and 0.92 (0.90–0.93) in the external validation cohort for 1-year all-cause death; an AUC of 0.74 (0.70–0.78) in the internal validation cohort and 0.81 (0.76–0.85) in the external validation cohort for 1-year myocardial infarction; and an AUC of 0.70 (0.66–0.75) in the internal validation cohort and 0.86 (0.82–0.89) in the external validation cohort for 1-year major bleeding.

**Interpretation** A machine learning-based approach for the identification of predictors of events after an ACS is feasible and effective. The PRAISE score showed accurate discriminative capabilities for the prediction of all-cause death, myocardial infarction, and major bleeding, and might be useful to guide clinical decision making.

**Funding** None.

**Copyright** © 2021 Elsevier Ltd. All rights reserved.

#### Introduction

Patients with acute coronary syndrome (ACS) are at high risk for ischaemic and bleeding events, with both being drivers of adverse prognosis.<sup>1</sup> Careful evaluation of these risks plays a fundamental role in the clinical management of each patient, with important implications regarding the choice of optimal medical therapy for secondary prevention.<sup>2–6</sup>

To this aim, several predictive tools have been developed to estimate ischaemic and bleeding risks following an



CrossMark

praise.hpc4ai.it

adverse events following an acute coronary syndrome (PRAISE). a modelling study of pooled datasets, *The Lancet*, Volume 397, Issue 10270, 2021, Pages 199–207, ISSN 0140-6736.

DOI: [https://doi.org/10.1016/S0140-6736\(20\)32519-8](https://doi.org/10.1016/S0140-6736(20)32519-8)

Single patient analysis    Multiple patients analysis

**Single patient analysis**

In order to run a single patient analysis with PRAISE it is necessary provide all the clinical, therapeutic, angiographic and procedural data available for the patient, then press the **SUBMIT** button. The result will be shown at the bottom of the page, showing the calculated **score** for death, ReAMI and BARC MB events with the corresponding **risk class** (low/intermediate/high). The score is calculated as a probability, so it is always included between 0 and 1.

Note that the score will be calculated independently from the number of variables provided; nonetheless it is worth noting that the more information are provided the more accurate the prediction will be.

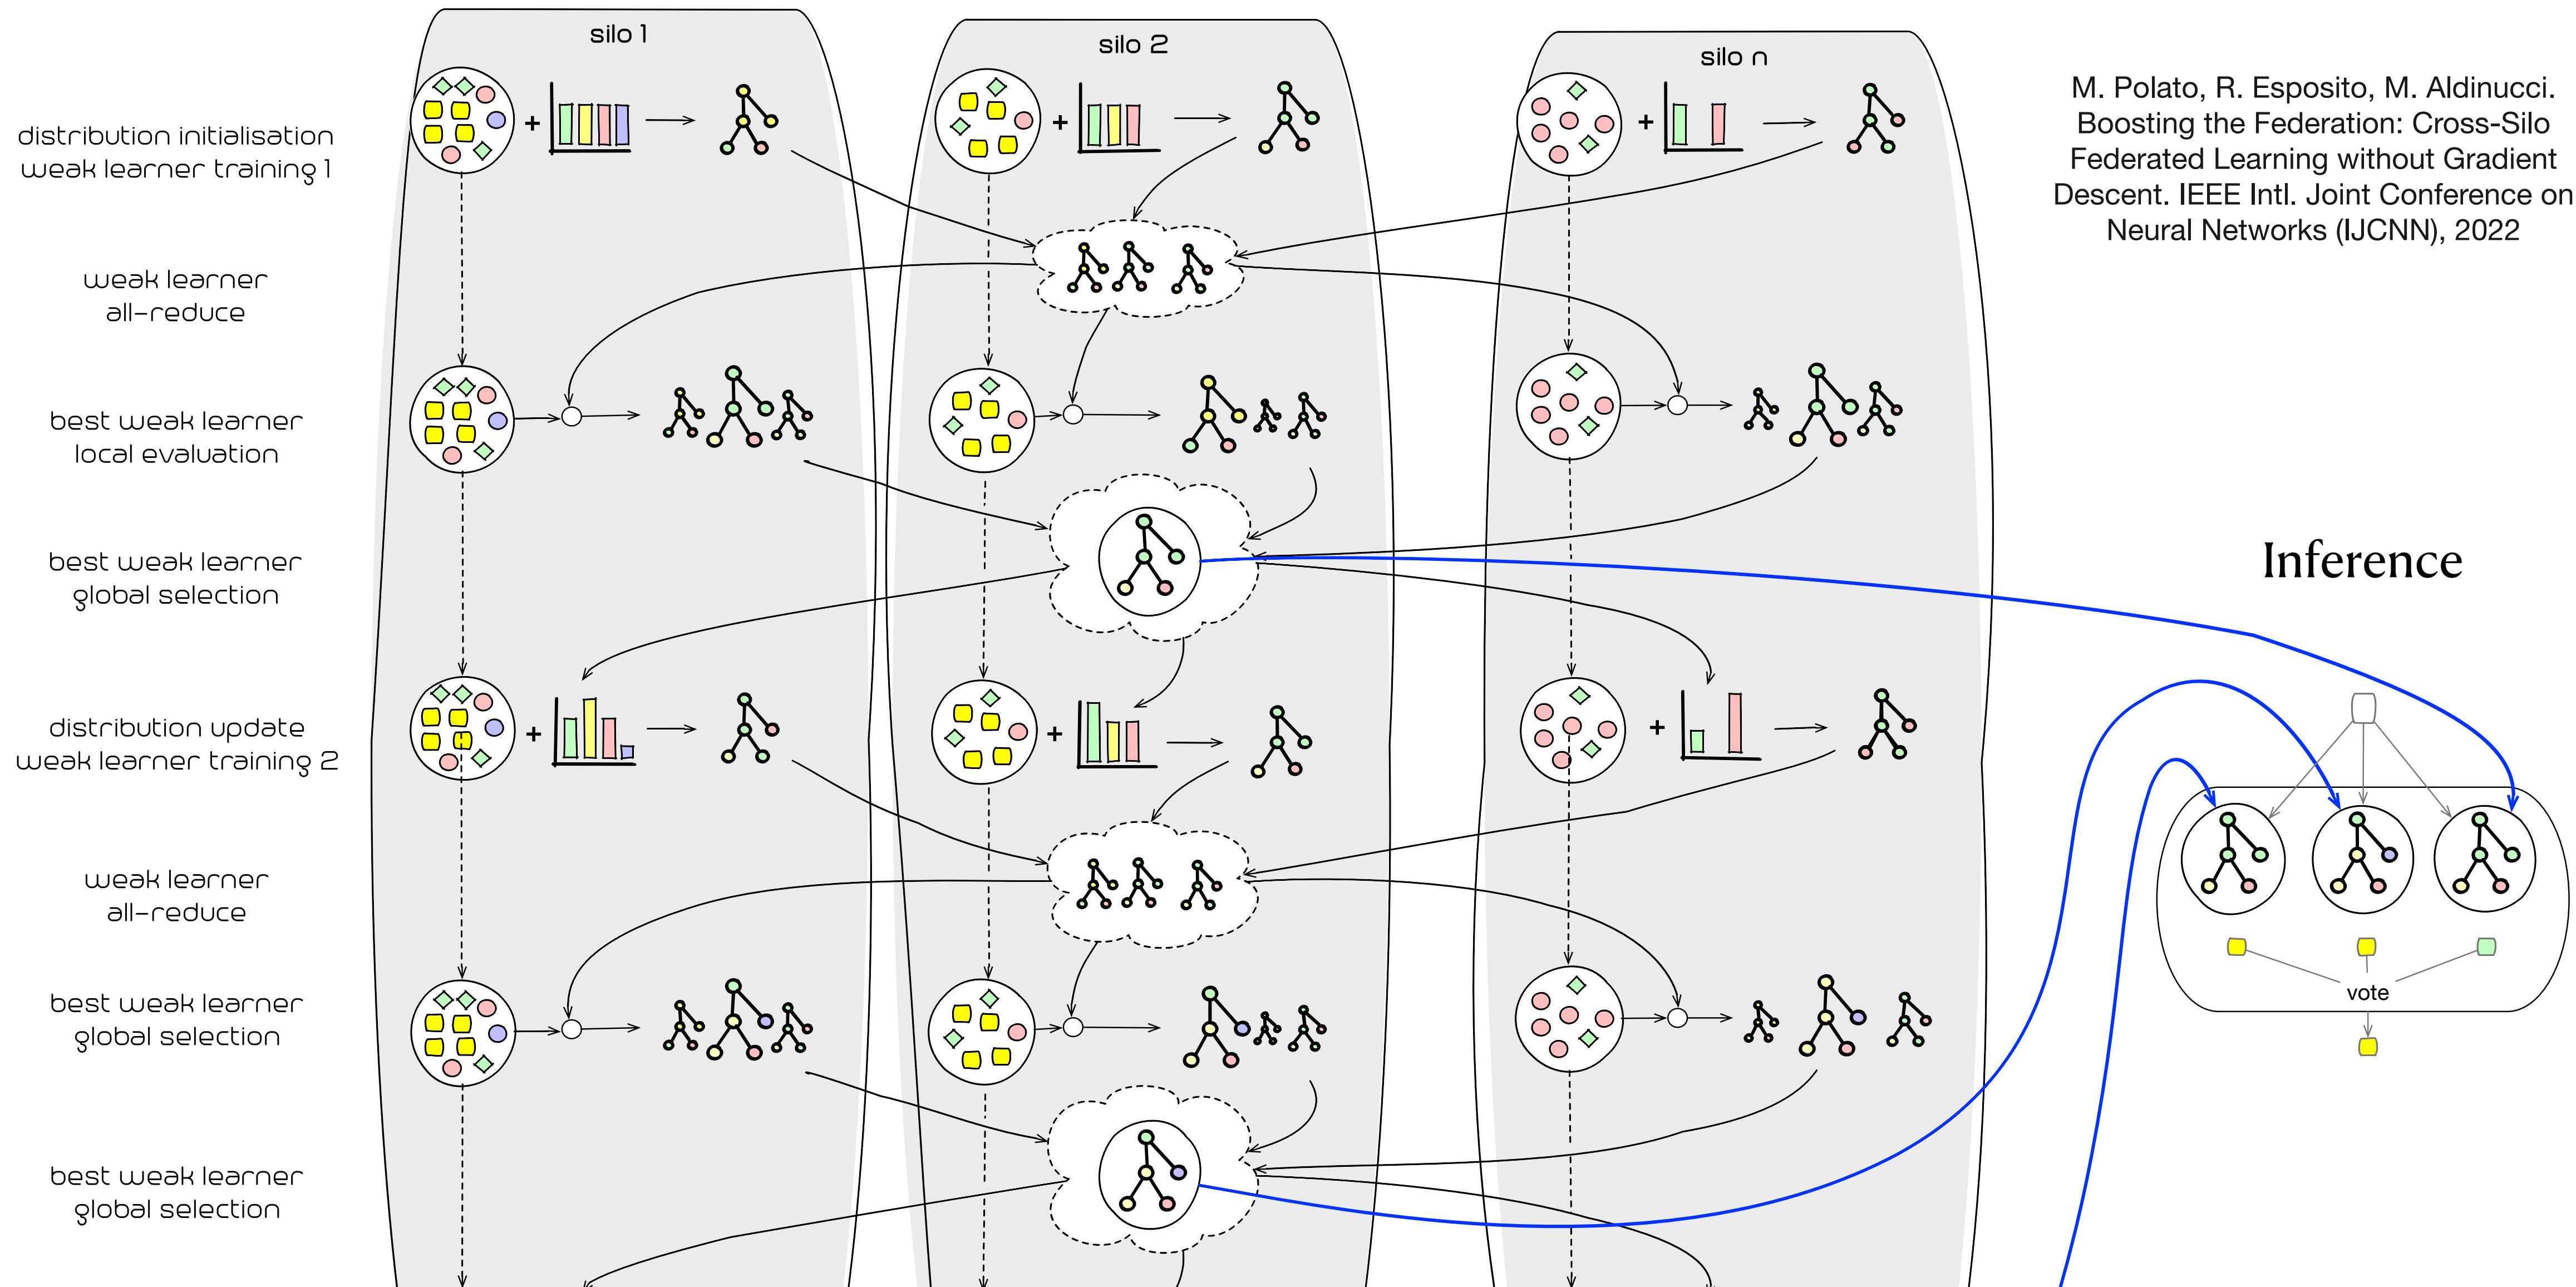
**Clinical variables**

Age	Hemoglobin (g/dl)	LVEF (%)	eGFR (MDRD)
Sex	Hypertension	Hyperlipidemia	Peripheral Artery Disease
<input type="radio"/> Male <input type="radio"/> Female <input checked="" type="radio"/>	<input type="radio"/> No <input checked="" type="radio"/> Yes <input type="radio"/>	<input type="radio"/> No <input checked="" type="radio"/> Yes <input type="radio"/>	<input type="radio"/> No <input checked="" type="radio"/> Yes <input type="radio"/>
Unknown	Unknown	Unknown	<input type="radio"/> No <input checked="" type="radio"/> Yes <input type="radio"/>
Prior AMI	Prior CABG	Prior stroke	Prior bleeding
<input type="radio"/> No <input checked="" type="radio"/> Yes <input type="radio"/>	<input type="radio"/> No <input checked="" type="radio"/> Yes <input type="radio"/>	<input type="radio"/> No <input checked="" type="radio"/> Yes <input type="radio"/>	<input type="radio"/> No <input checked="" type="radio"/> Yes <input type="radio"/>
Unknown	Unknown	Unknown	<input type="radio"/> Unknown
Malignancy	STEMI	NSTEMI	<input type="radio"/> Prior bleeding
<input type="radio"/> No <input checked="" type="radio"/> Yes <input type="radio"/>	<input type="radio"/> No <input checked="" type="radio"/> Yes <input type="radio"/>	<input type="radio"/> No <input checked="" type="radio"/> Yes <input type="radio"/>	<input type="radio"/> No <input checked="" type="radio"/> Yes <input type="radio"/>



DEEPHEALTH

# Cross-silo non-SGD federated boosting



# Cross-silo non SGD federated learning

distribution initial  
weak learner tra

weak learner  
all-reduce

best weak learner  
local evalua

best weak learner  
global selecti

distribution upda  
weak learner tra

weak learner  
all-reduce

best weak learner  
global selecti

best weak learner  
global selection

**About**

Extension to the OpenFL framework for non gradient descent learning

**Readme**

0 stars

6 watching

0 forks

**Releases**

No releases published

[Create a new release](#)

**Packages**

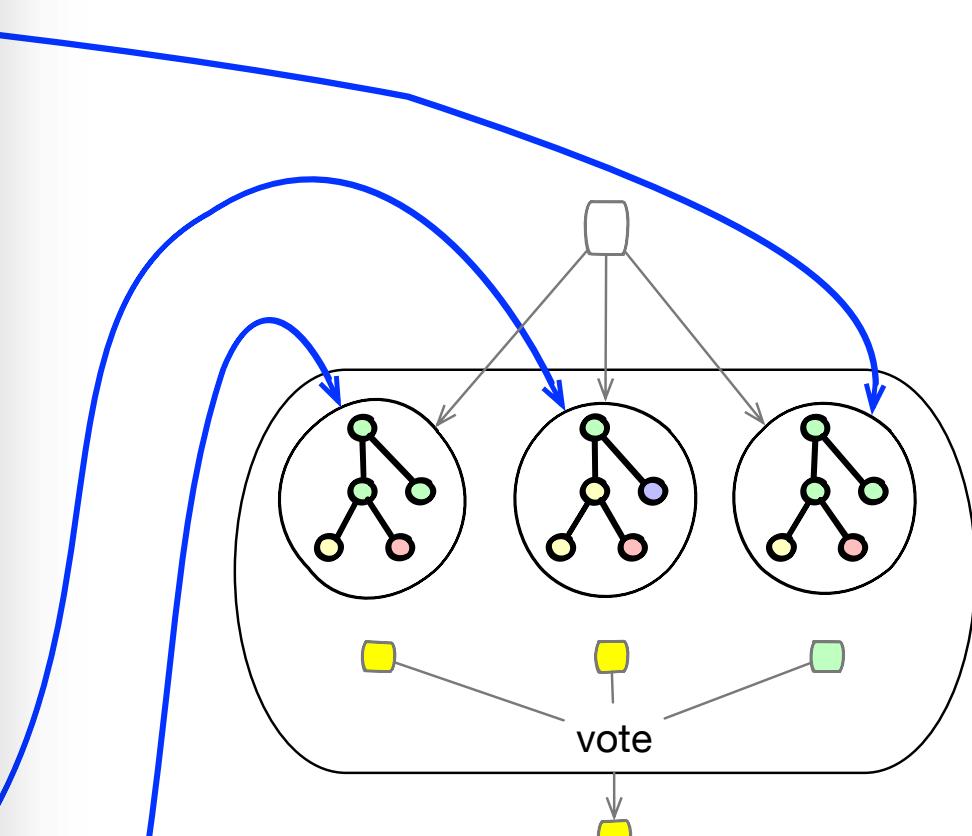
No packages published

[Publish your first package](#)

**Contributors** 2

M. Polato, R. Esposito, M. Aldinucci.  
Boosting the Federation: Cross-Silo  
Federated Learning without Gradient  
Descent. IEEE Intl. Joint Conference on  
Neural Networks (IJCNN), 2022

## Inference



# HPC4AI e il trasferimento tecnologico

## anche se in realtà il mercato di HPC4AI è la ricerca non i servizi

- **RAI:** Nuovo Auditel socio-semantico
- **Cerved group:** Analisi probabilità di fallimento post-covid mediante DeepLearning
- **Mozilla foundation:** Speech-to-text per la lingua Italiana
- **Università “L. Bocconi”:** Supporto al dipartimento di data science
- **Cliniche universitarie Molinette e Candiolo:** Oncologia, Cardiologia, Patologia
- **Synesthesia:** Deep Learning
- **SAFAN:** bioinfo pipelines
- ...
- **Circular Health** (Compagnia di San Paolo, 2021, 36 months, total cost **1M€**)
- **DeepHealth** (EC H2020 IA, ICT-2018-11): Deep-Learning and HPC to Boost Biomedical Applications for Health (2019, 36 months, total cost **14.8M€**, G.A. 825111)
- **Brainteaser** (EC H2020 RIA, SC1-DTH-2020-1): BRinging Artificial INTelligencE home for a better cAre of amyotrophic lateral sclerosis and multiple SclERosis (2020, 36 months, total cost **5.9M€**, G.A. 101017598)
- **European Processor Initiative** (EC H2020 RIA, EuroHPC-02-2020): EPI SGA-2 (2022, 48 months, total cost **70M€**)
- **EPI-TO** - the European Processor Iniziative at the University of TOriino - 250k€

Prof, le persone che prendono sul serio le applicazioni  
dovrebbero progettare i propri sistemi, che si sa?

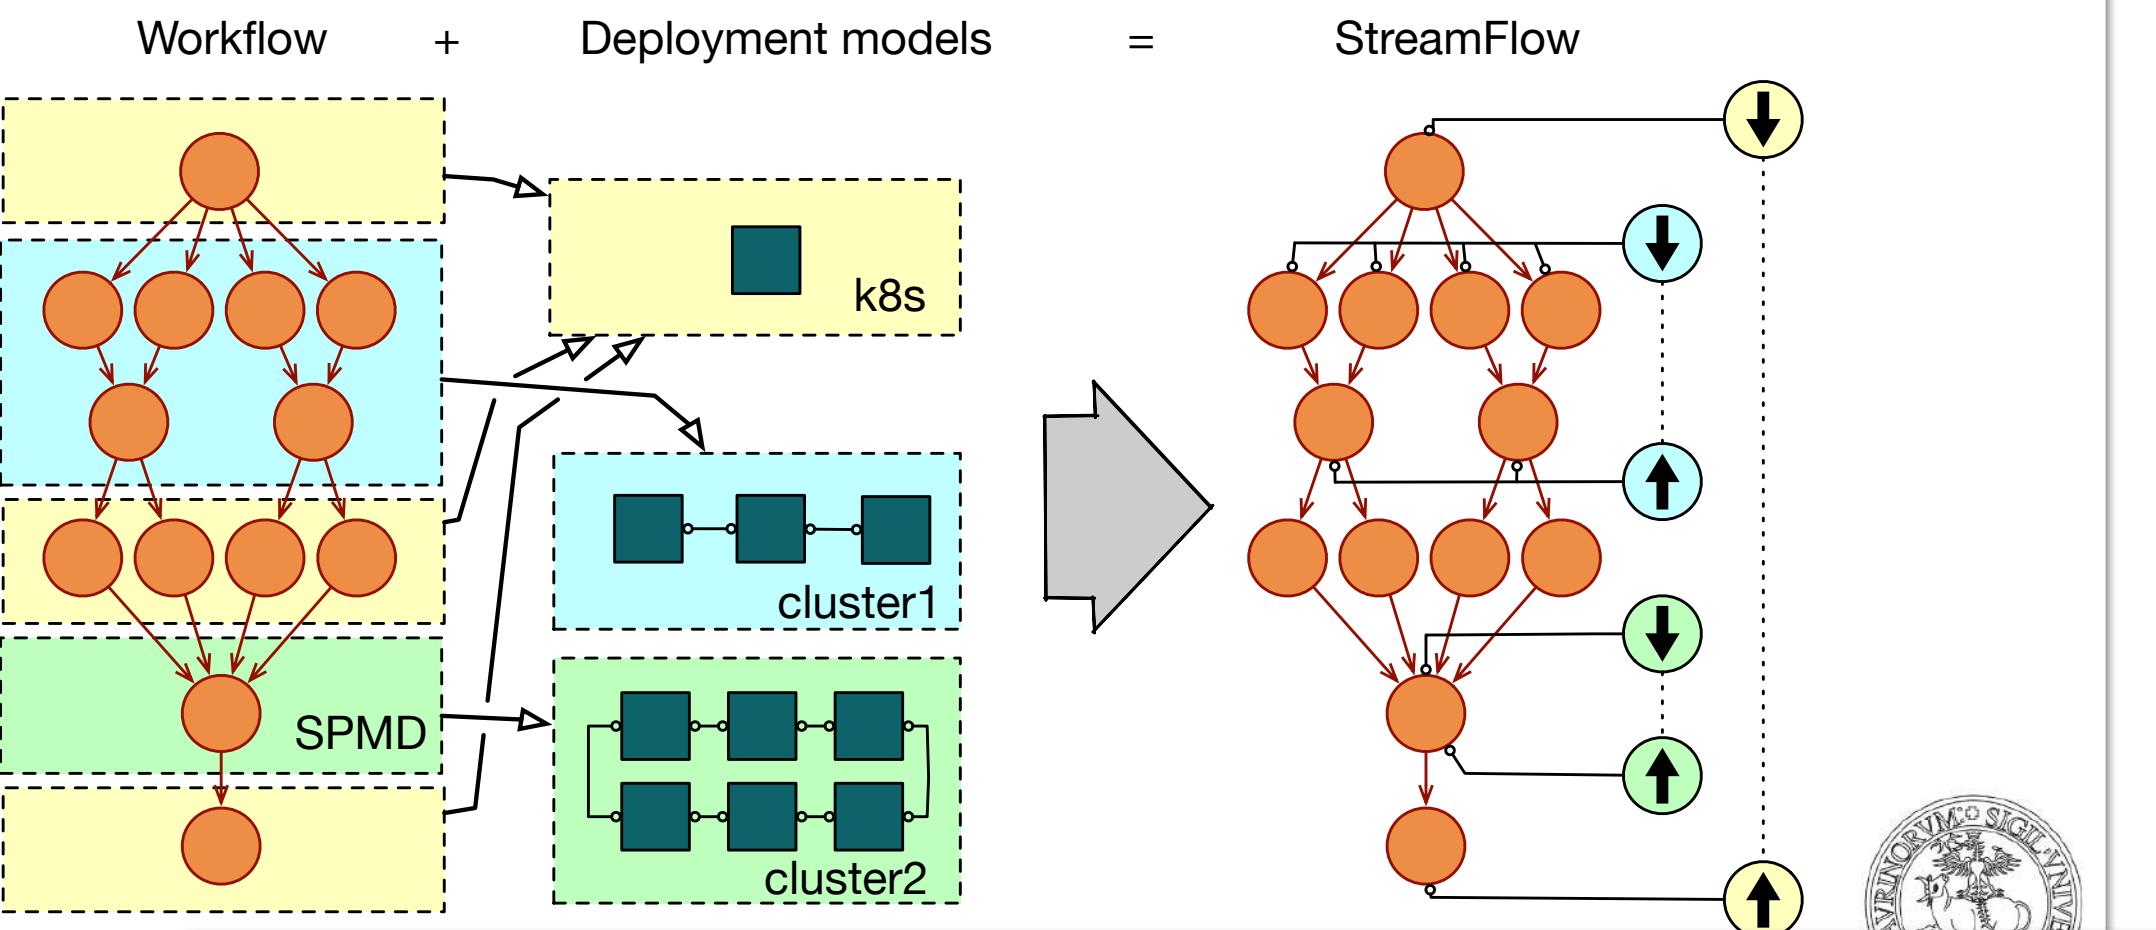
**PROGETTIAMO  
TUTTO!**



# Inostri software

- Streamflow WMS
- Jupyter-workflow & Jupyter-as-a-Service (dossier)
- Portable Secure Tenant PaaS
- UrgentSLURM
- Extended-OpenFL and AI-enclave

# StreamFlow



46

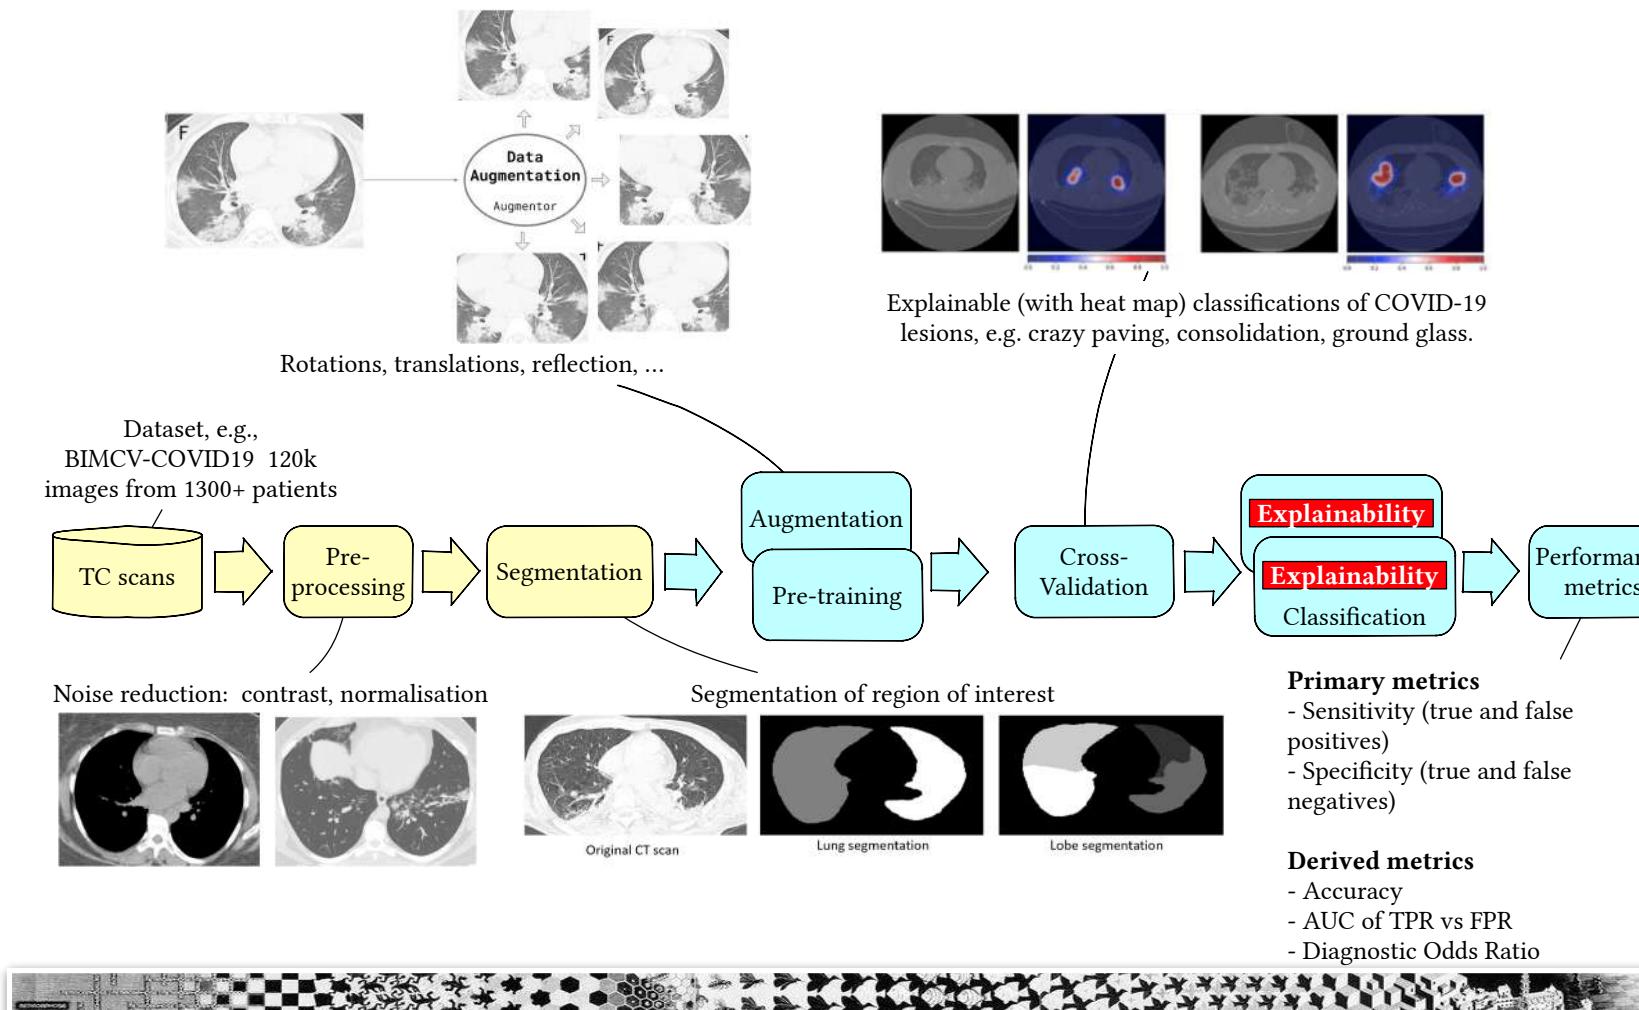
**In Production**

Software	Description	Self-Reported Compliance	Platform support
cwltool	Reference implementation of CWL	CWL v1.0 - v1.2	Linux, OS X, Windows, local execution only
Arvados	Distributed computing platform for data analysis on massive data sets.	CWL v1.0 - v1.2 required 100% <a href="#">Using CWL on Arvados</a>	AWS, GCP, Azure, Slurm, LSF
Toil	Toil is a workflow engine entirely written in Python.	CWL v1.0 - v1.2	AWS, Azure, GCP, Grid Engine, HTCondor, LSF, Mesos, OpenStack, Slurm, PBS/Torque
CWL-Airflow	Package to run CWL workflows in Apache-Airflow (supported by BioWardrobe Team, CCHMC)	CWL v1.0 - v1.2	Linux, OS X
<b>StreamFlow</b>	Workflow Management System for hybrid HPC-Cloud infrastructures	CWL v1.0 - v1.2 required 100% (and nearly all optional features)	Kubernetes, HPC with Singularity (PBS, Slurm), <a href="#">Occam</a> , multi-node SSH, local-only (Docker, Singularity)

<https://streamflow.di.unito.it>

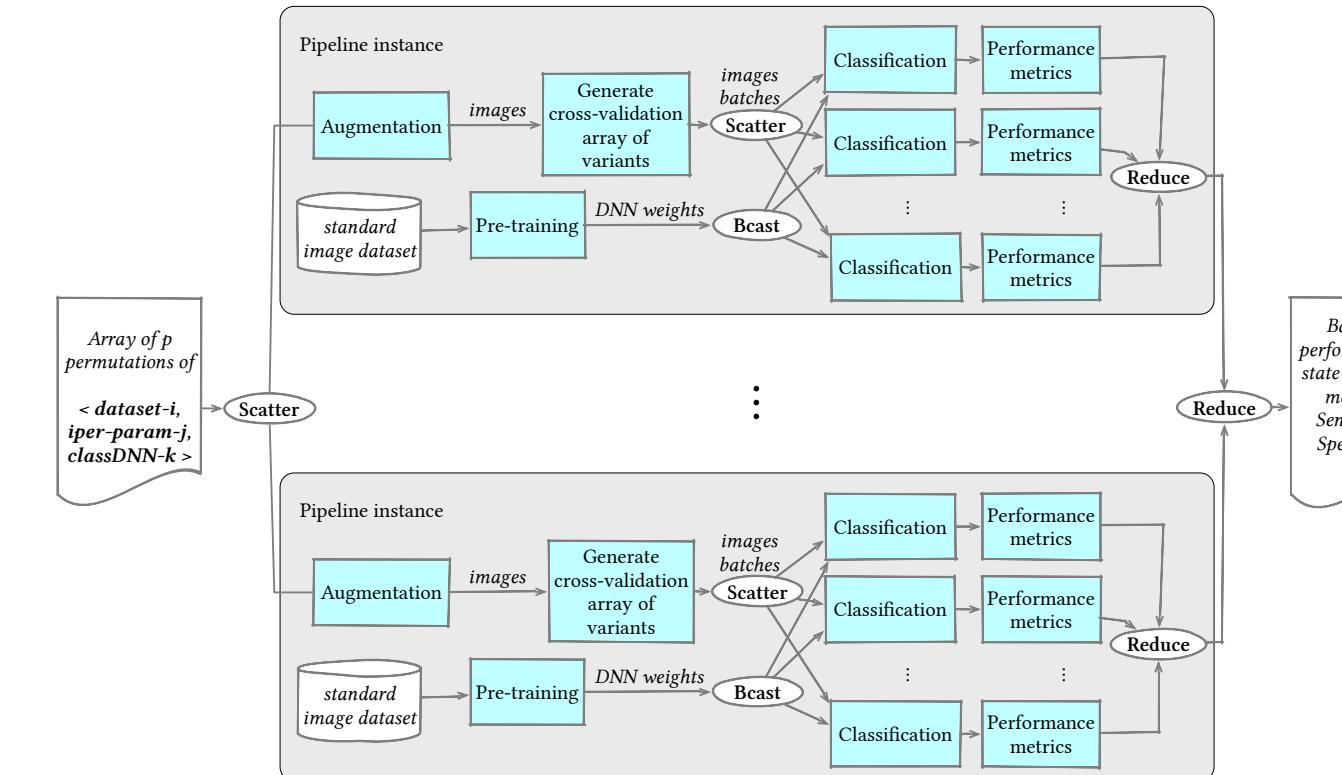
41

# Case study: CLAIRE COVID-19 pipeline



50

# Case study: CLAIRE COVID-19 pipeline



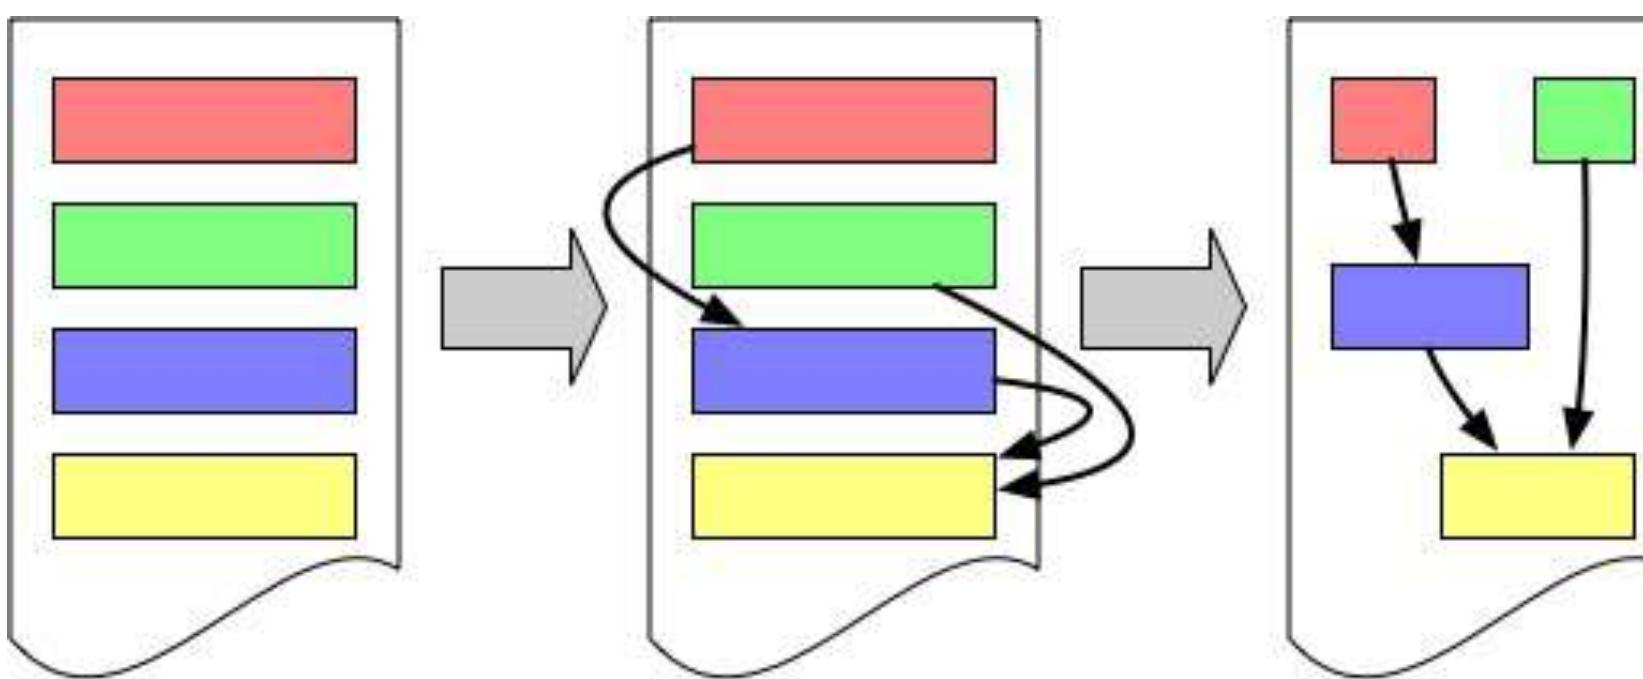
51

<https://streamflow.di.unito.it>



# Hybrid literate workflows

Use Notebook metadata as a **coordination language** to express distributed workflows



63

```

train_loader = torch.utils.data.DataLoader(dataset1,**train_kwargs)
test_loader = torch.utils.data.DataLoader(dataset2, **test_kwargs)

model = Net().to(device)
optimizer = optim.Adadelta(model.parameters(), lr=args.lr)

scheduler = StepLR(optimizer, step_size=1, gamma=args.gamma)
for epoch in range(1, args.epochs + 1):
    train(args, model, device, train_loader, optimizer, epoch)
    test(model, device, test_loader)
    scheduler.step()

if args.save_model:
    torch.save(model.state_dict(), "mnist_cnn.pt")

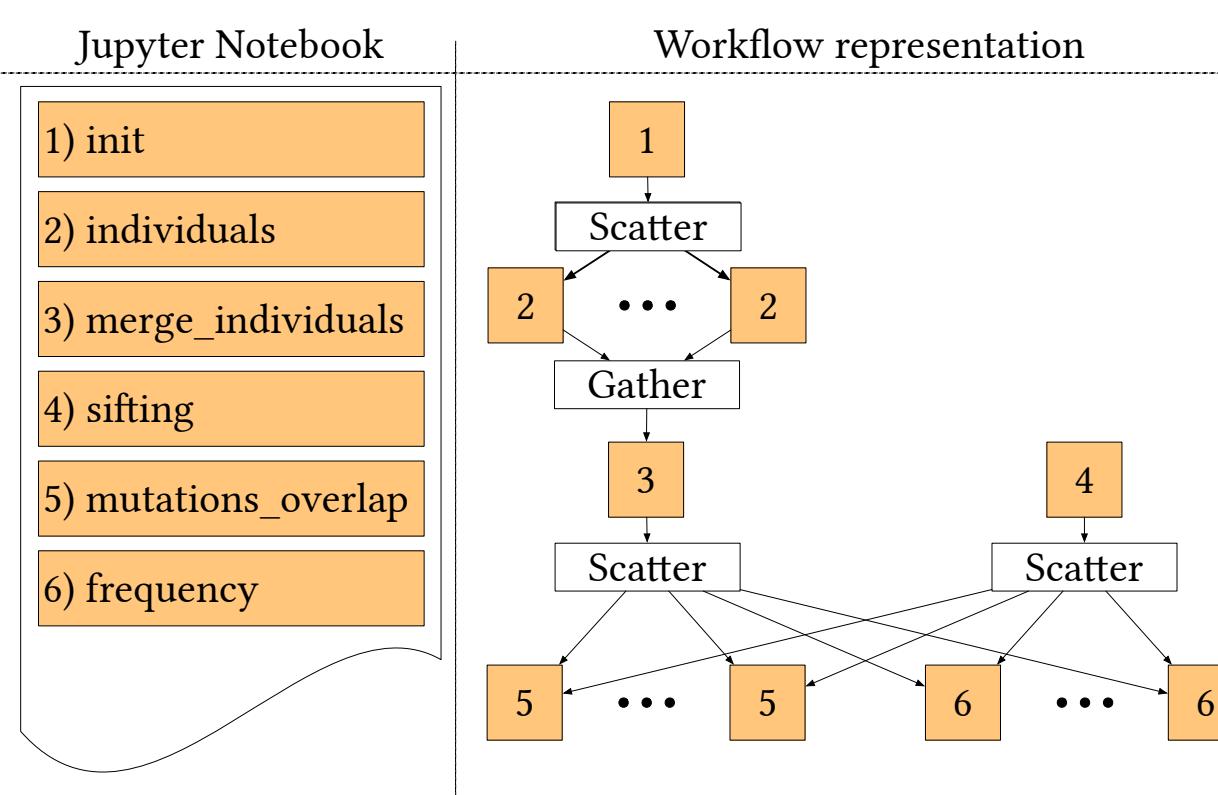
if __name__ == '__main__':
    main()

```

Simple 0 1 Python on Kubernetes | Idle Mode: Command Ln 137, Col 11 claire-covid.ipynb

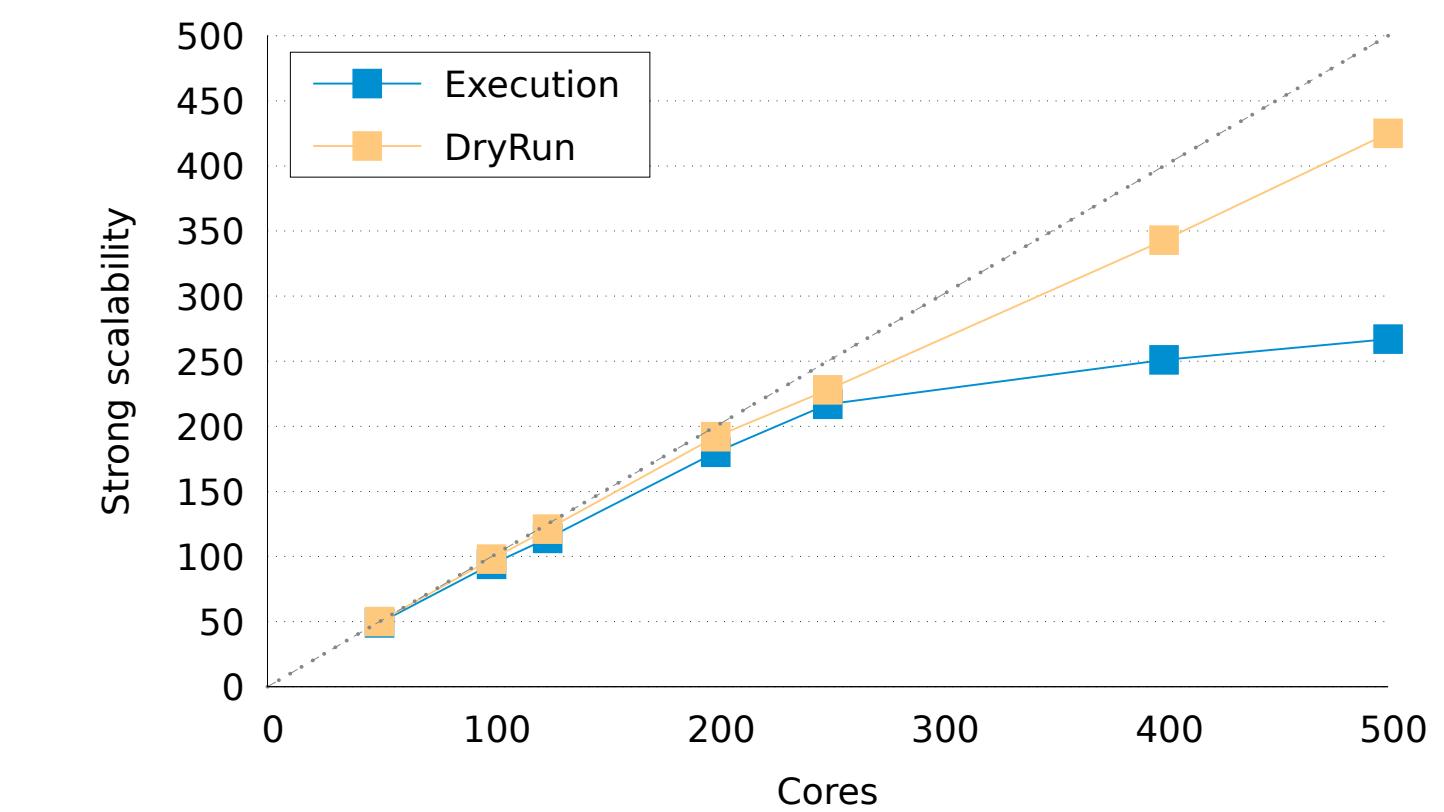


## Case study: 1000-genome Notebook



102

## Case study: 1000-genome Notebook

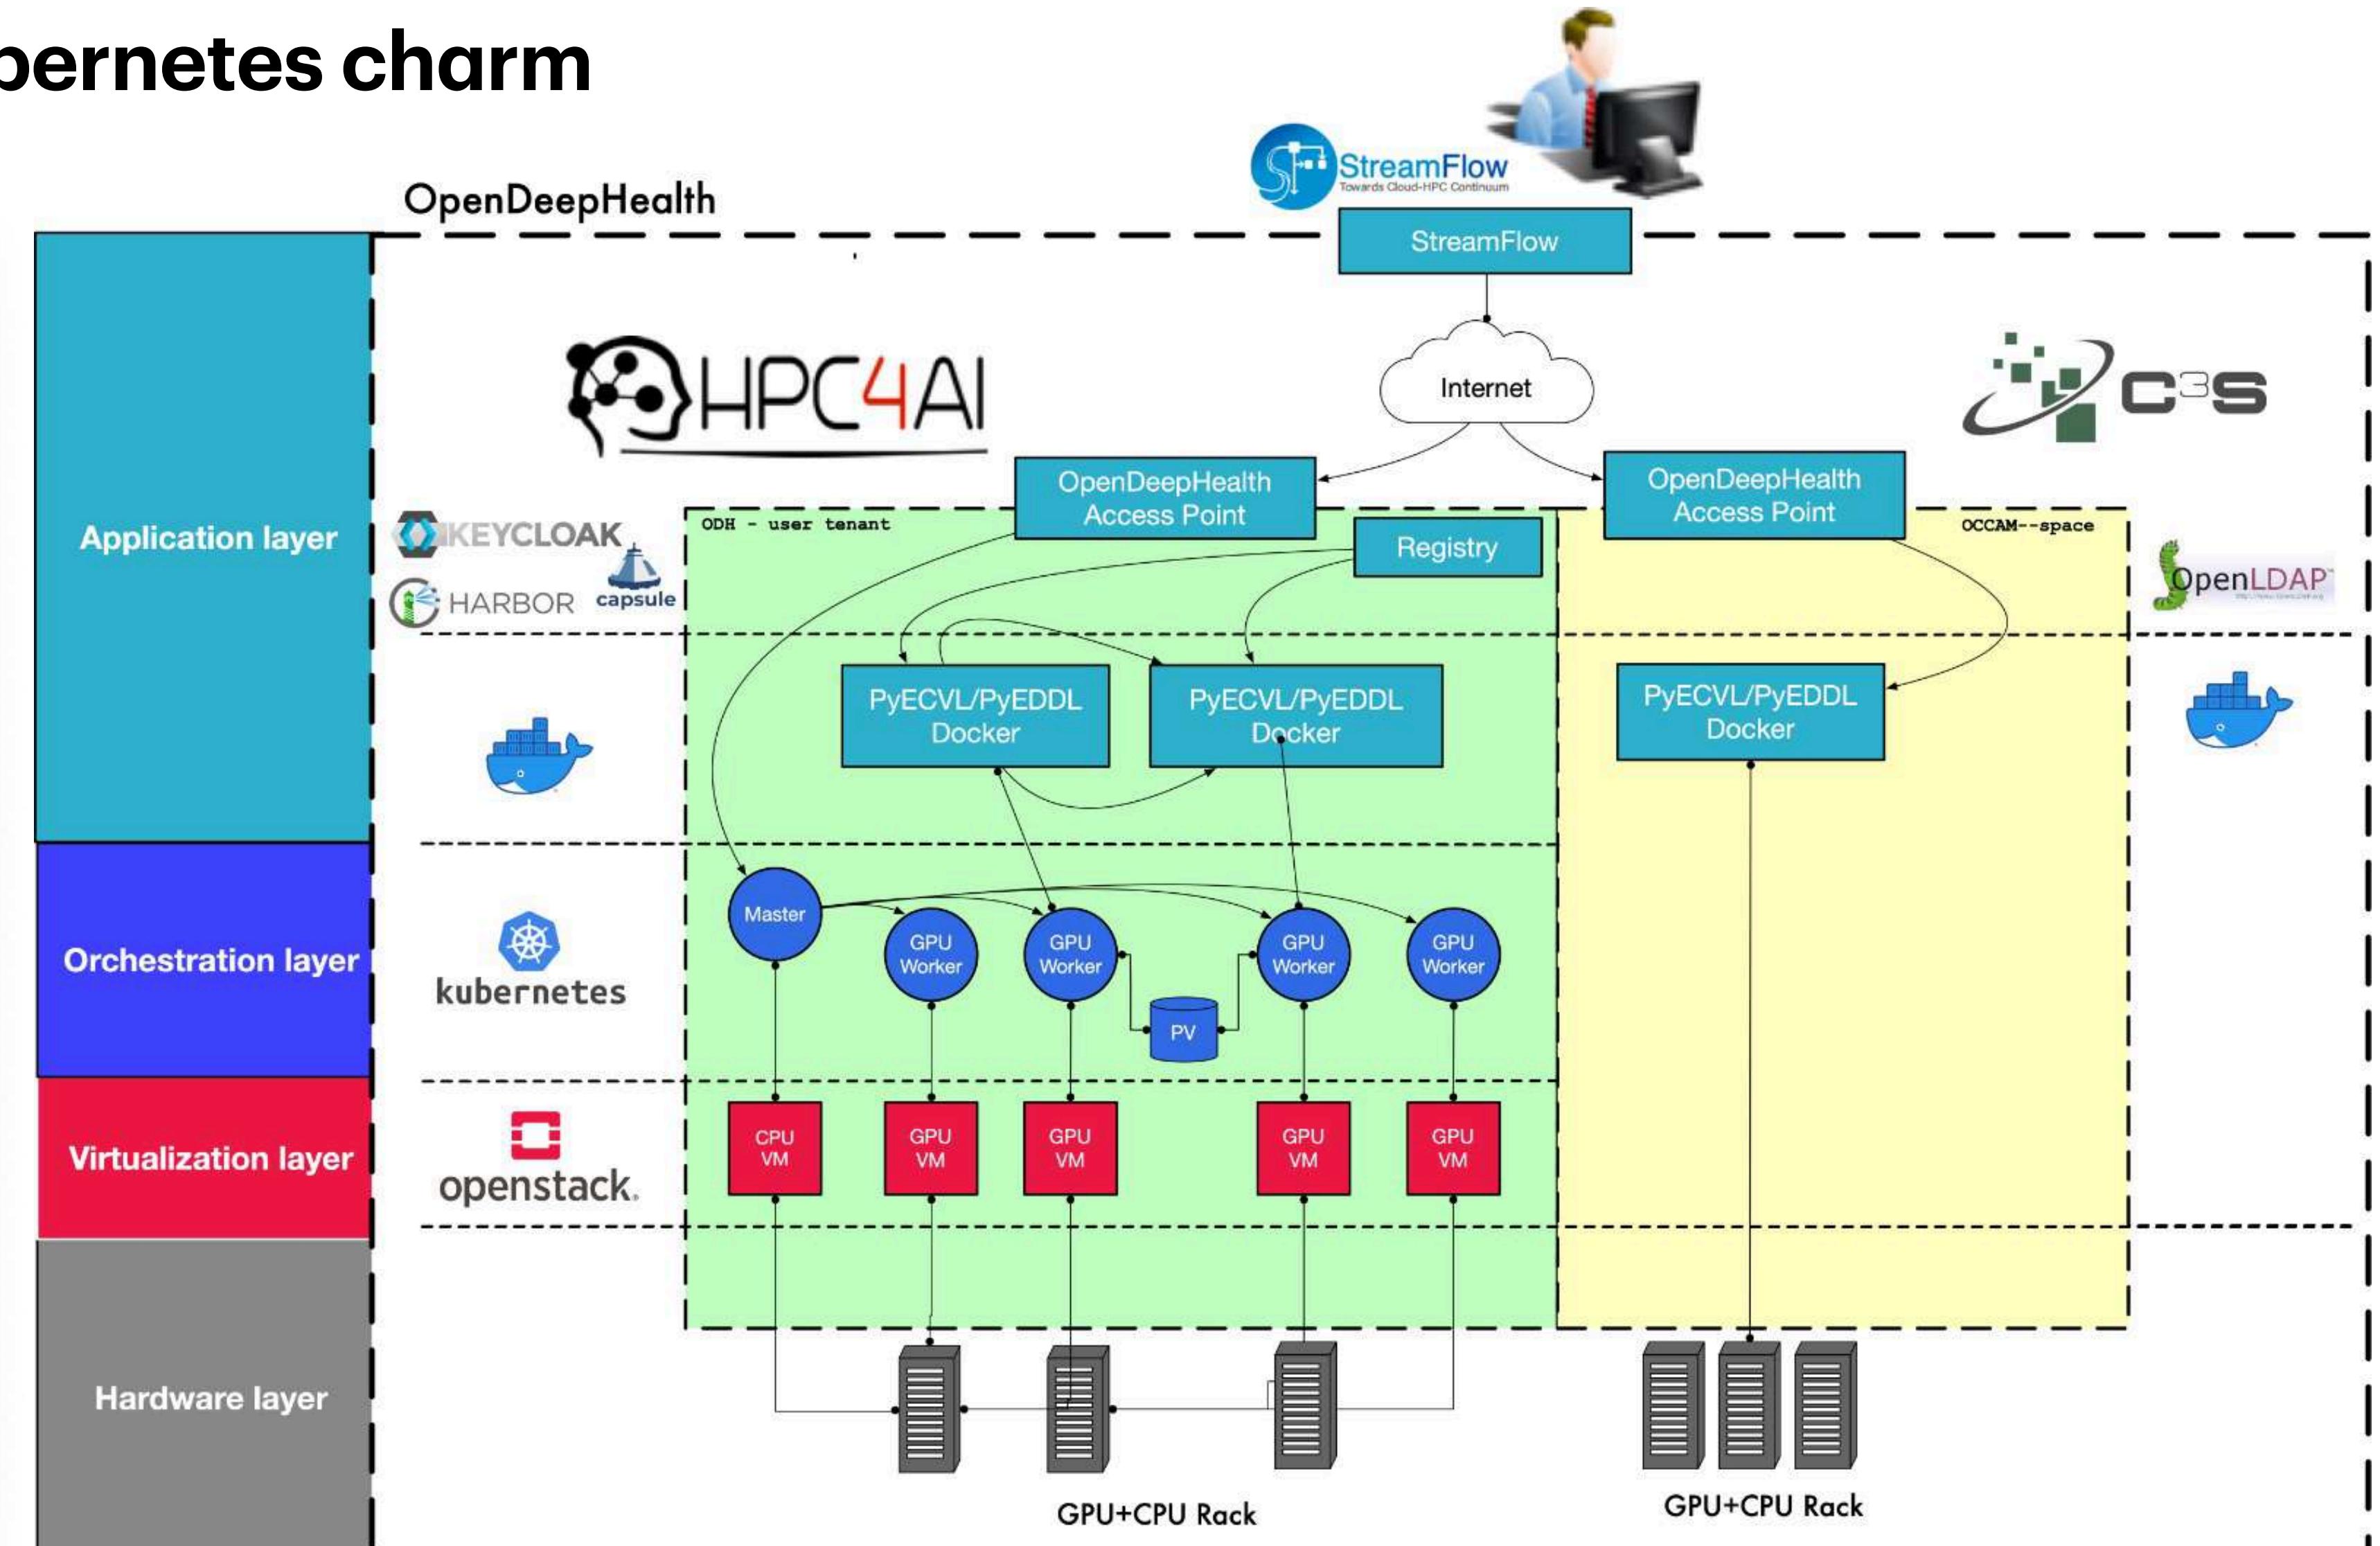
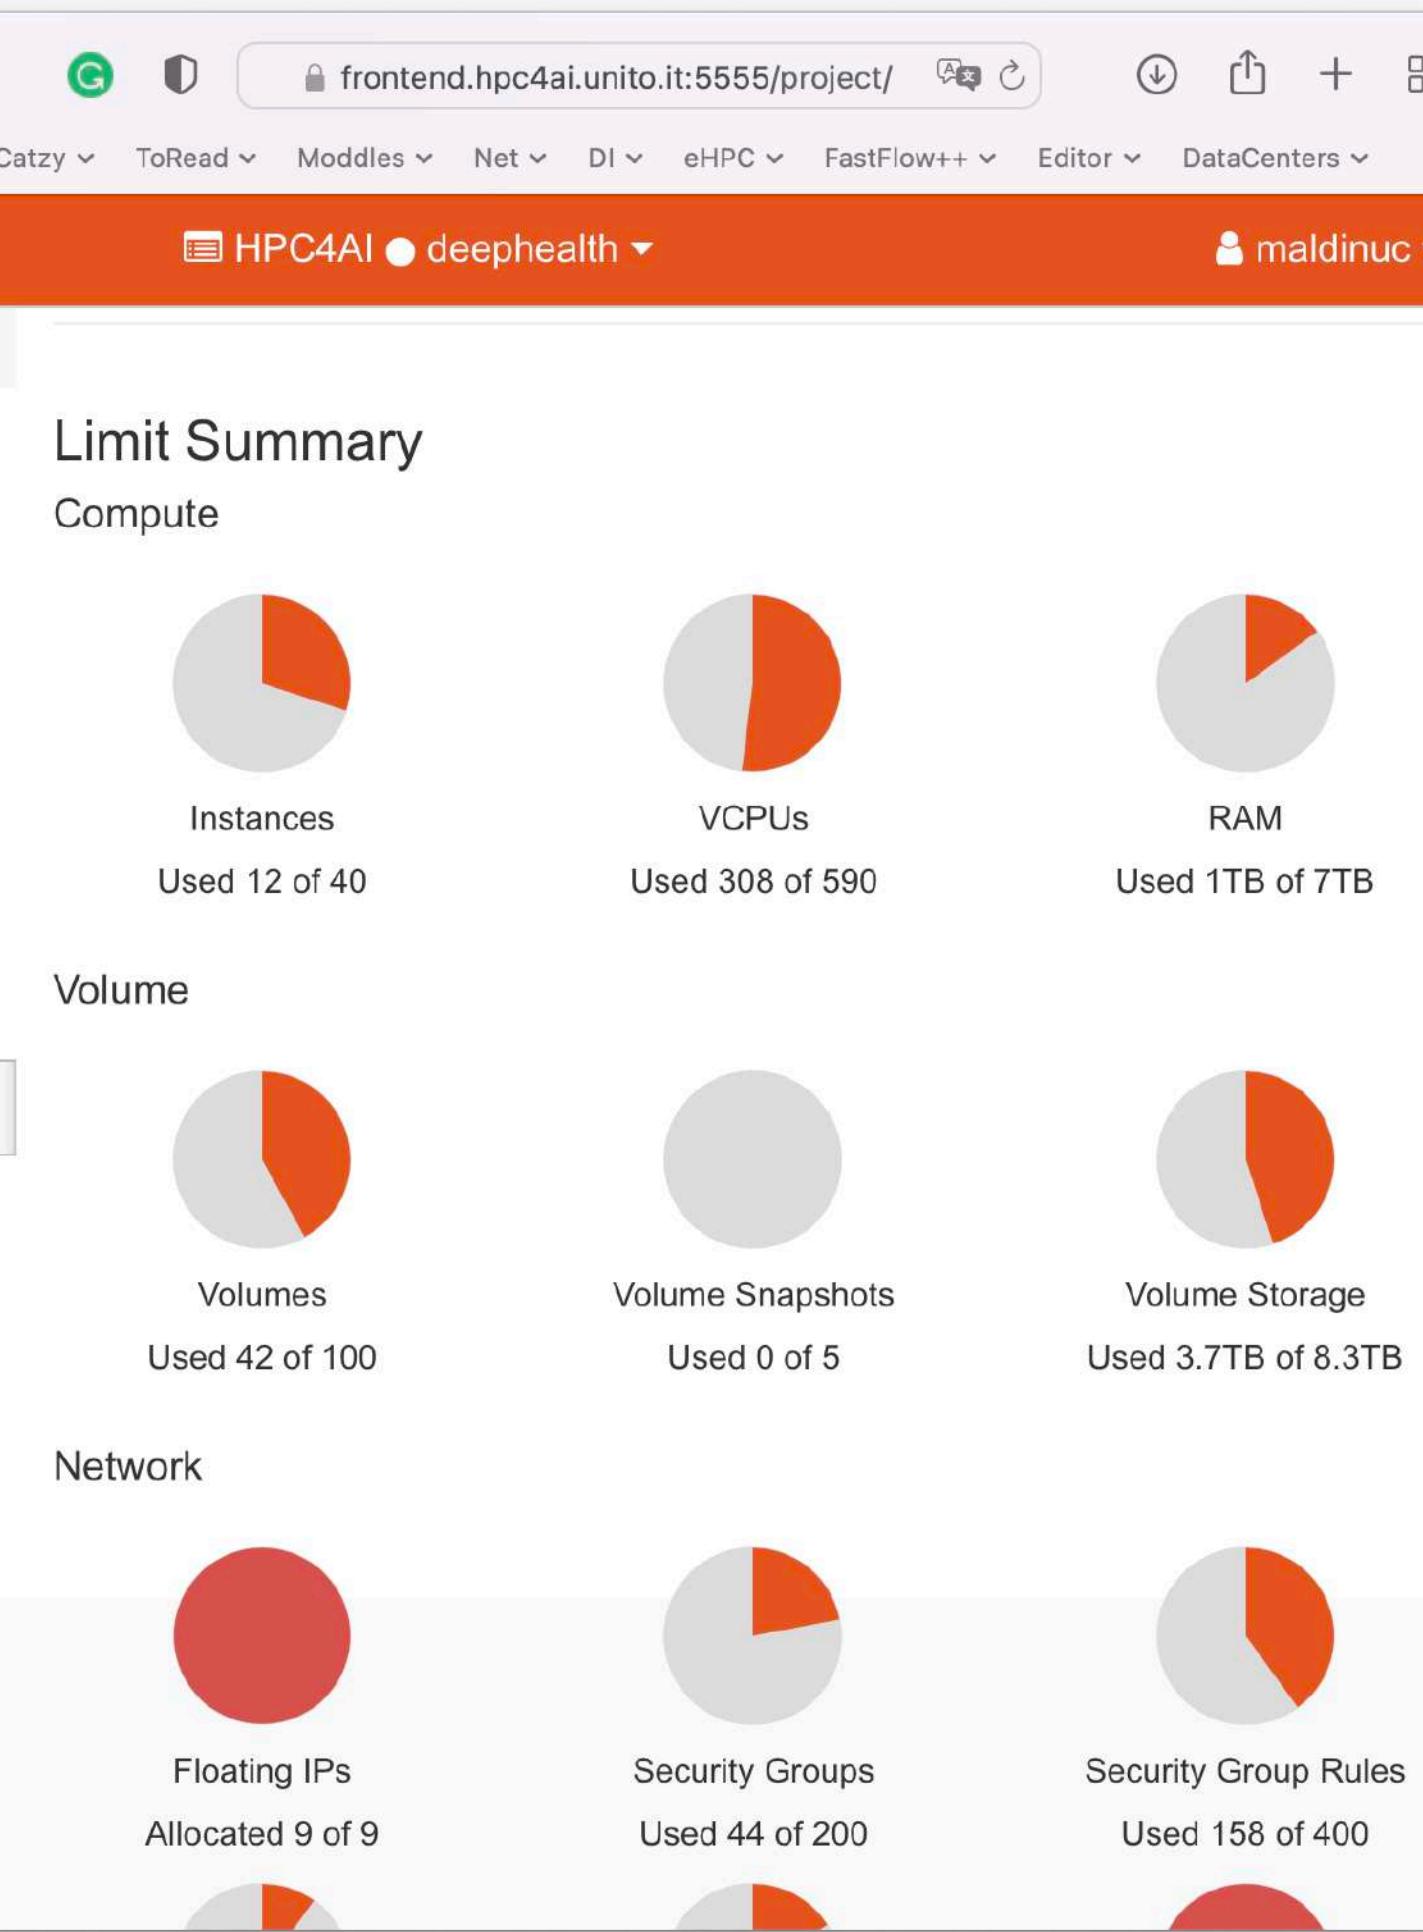


104

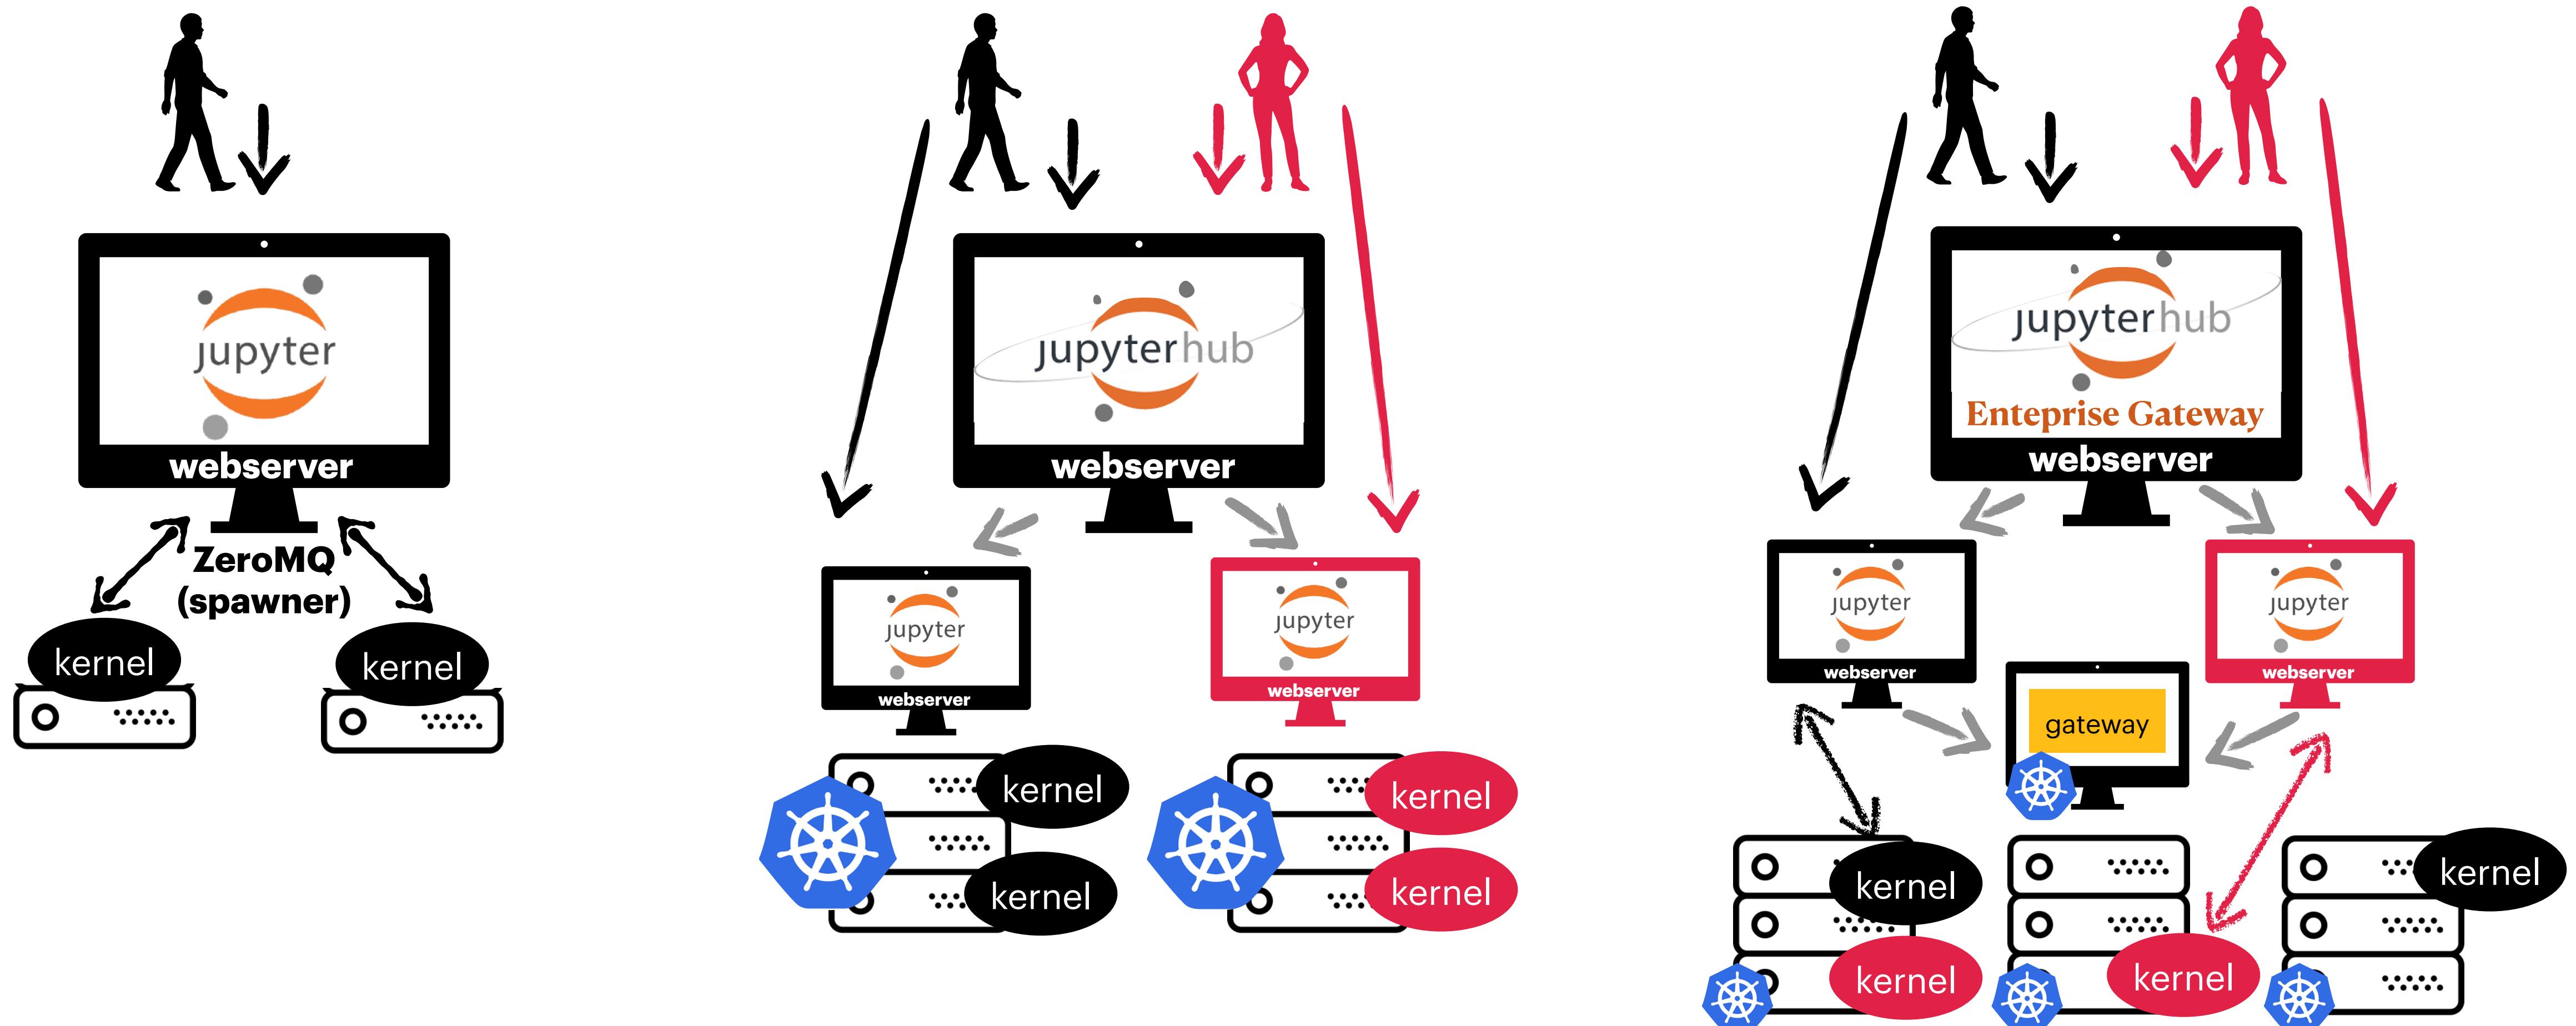


# Portable Secure Tenant (PaaS)

extends GARR/JUJU Kubernetes charm



# Jupyter-as-a-Service (dossier)



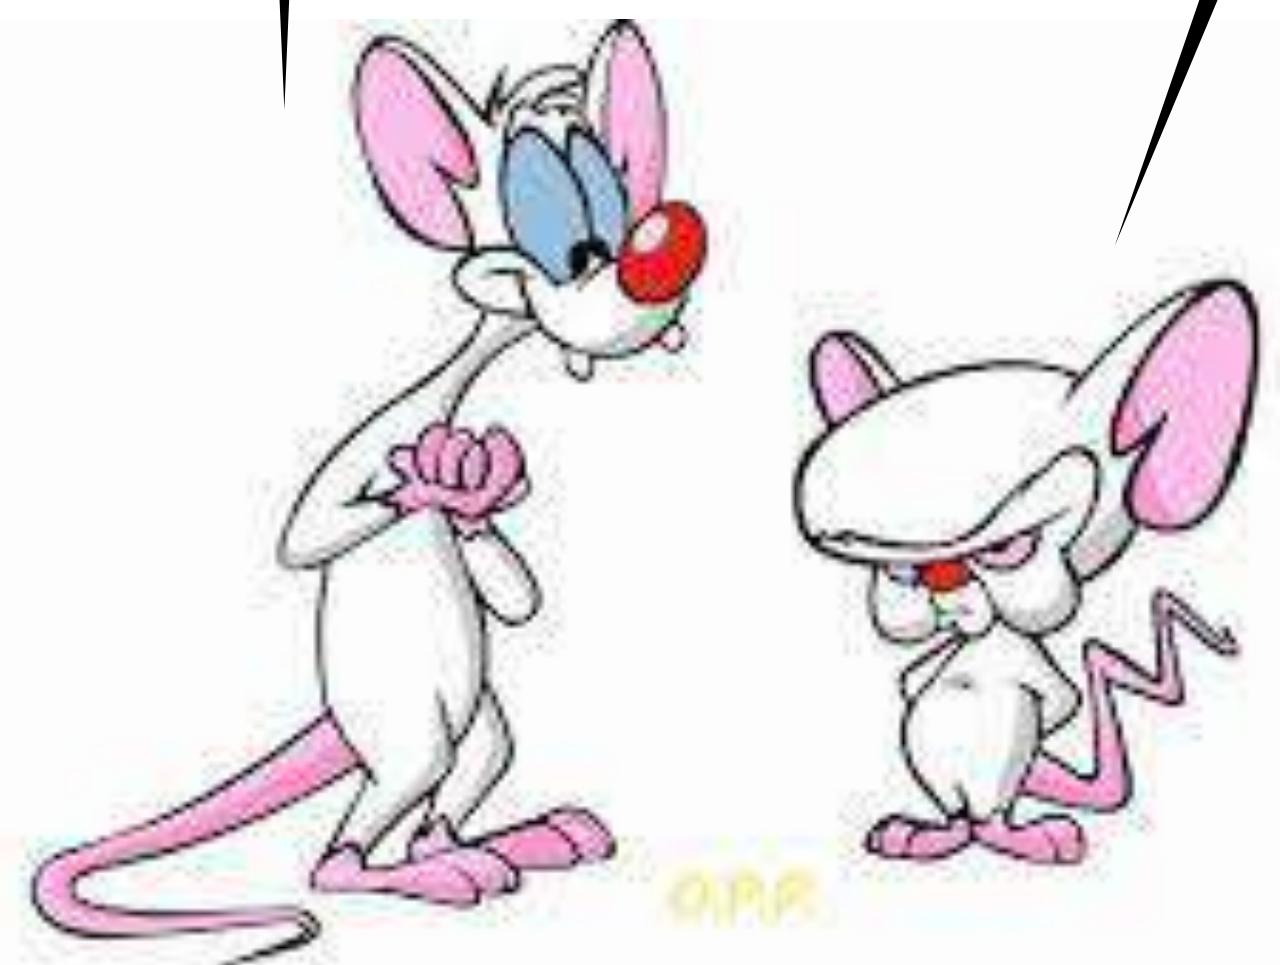
Single user  
Single node (\*)

Multi-tenant  
Single node per user (\*)

Multi-tenant  
Multiple nodes per user

Prof ma quanti siamo a fare tutte queste cose?

**FEDERIAMO  
TUTTI!**



# CINI HPC Key Technologies & Tools



21

- Avviato ufficialmente nel Marzo 2021 - WG attivo Sett. 2019
  - Direttore Marco Aldinucci (UNITO)
- 30 università e 8 istituti CNR,
  - Promuove la strategia nazionale degli “*informatici*” in ambito HPC
- Partecipa 6 progetti EuroHPC di grande dimensione
  - Aggrega competenze solide ma sparse
- Nasce con una precisa visione al livello nazione
  - Infrastrutture (CINECA, ENEA, etc.)
  - **Technologies & Tools**
  - Applicazioni (non computer science)
- Noi siamo **Technologies & Tools** e vogliamo collaborare con Applicazioni e Infrastrutture

## The Italian Research on HPC Key Technologies across EuroHPC

### Invited Paper

Marco Aldinucci  
DI - University of Torino  
CINI HPC-KTT laboratory  
Torino, Italy  
marco.aldinucci@unito.it

Claudio A. Ardagna  
Università degli Studi di Milano  
CINI HPC-KTT laboratory  
Milano, Italy  
claudio.ardagna@unimi.it

Biagio Cosenza  
University of Salerno  
CINI HPC-KTT laboratory  
Salerno, Italy  
bcosenza@unisa.it

William Fornaciari  
DEIB - Politecnico di Milano  
CINI HPC-KTT laboratory  
Milano, Italy  
william.fornaciari@polimi.it

Raffaele Montella  
DiST - University of Naples  
“Parthenope”  
CINI HPC-KTT laboratory  
Napoli, Italy  
raffaele.montella@uniparthenope.it

Daniele Simoni  
DIME - University of Genova  
CINI HPC-KTT laboratory  
Genova, Italy  
daniele.simoni@unige.it

Giovanni Agosta  
DEIB - Politecnico di Milano  
CINI HPC-KTT laboratory  
Milano, Italy  
agosta@acm.org

Andrea Bartolini  
DEI - Università di Bologna  
CINI HPC-KTT laboratory  
Bologna, Italy  
a.bartolini@unibo.it

Marco Danelutto  
University of Pisa  
CINI HPC-KTT laboratory  
Pisa, Italy  
marco.danelutto@unipi.it

Roberto Giorgi  
DIISM - University of Siena  
CINI HPC-KTT laboratory  
Siena, Italy  
giorgi@unisi.it

Mauro Olivieri  
Sapienza University of Rome  
CINI HPC-KTT laboratory  
Roma, Italy  
mauro.olivieri@uniroma1.it

Massimo Torquati  
University of Pisa  
CINI HPC-KTT laboratory  
Pisa, Italy  
massimo.torquati@unipi.it

Antonio Andreini  
DIEF - University of Florence  
CINI HPC-KTT laboratory  
Firenze, Italy  
antonio.andreini@unifi.it

Alessandro Cilardo  
University of Naples Federico II  
CINI HPC-KTT laboratory  
Napoli, Italy  
acilardo@unina.it

Roberto Esposito  
DI - University of Torino  
CINI HPC-KTT laboratory  
Torino, Italy  
roberto.esposito@unito.it

Davide Lengani  
DIME - University of Genova  
CINI HPC-KTT laboratory  
Genova, Italy  
devide.lengani@unige.it

Sergio Saponara  
DII-University of Pisa  
CINI HPC-KTT laboratory  
Pisa, Italy  
sergio.saponara@unipi.it

### ABSTRACT

High-Performance Computing (HPC) is one of the strategic priorities for research and innovation worldwide due to its relevance for industrial and scientific applications. We envision HPC as composed

of three pillars: infrastructures, applications, and key technologies and tools. While infrastructures are by construction centralized in large-scale HPC centers, and applications are generally within the purview of domain-specific organizations, key technologies fall in an intermediate case where coordination is needed, but design and development are often decentralized. A large group of Italian researchers has started a dedicated laboratory within the National Interuniversity Consortium for Informatics (CINI) to address this challenge. The laboratory, albeit young, has managed to succeed in its first attempts to propose a coordinated approach to HPC research within the EuroHPC Joint Undertaking, participating in the calls 2019-20 to five successful proposals for an aggregate total cost of 95M€. In this paper, we outline the working group's scope and

# I principi fondativi

22

- **Non esclusività:** il gruppo non vuole impedire che i proponenti partecipino da soli quando ne hanno la forza. Il laboratorio non erige barriere ma si propone come una opportunità per raccogliere competenze che sono spesso troppo distribuite per fare massa critica.
- **Inclusività:** il gruppo che già ha iniziato lavorare ai progetti si propone di coinvolgere altri nodi, allargando il gruppo. Modello bottom-up.
- **Coordinamento:** il gruppo intende ricoprire il pillar “informatica” (HW e SW), cercando collaborazioni con infrastrutture e applicazioni per la costruzione di un sano sistema nazionale. Vogliamo lavorare con CINECA ed ENEA e con i gruppi di eccellenza sulle applicazioni “non informatici”, ad esempio con i Center of Excellence in HPC.
- **Missione:** progetti, didattica (MSc e dottorato), rete di reclutamento, trasferimento tecnologico.

# I progetti HPC-KTT

23

- **ACROSS** (EC H2020 IA, EuroHPC-01-2019): HPC Big Data Artificial Intelligence cross-stack platform toward exascale (2021, 36 months, total cost **8M€**, G.A. n. 955648)
- **ADMIRE** (EC H2020 RIA, EuroHPC-01-2019): Adaptive multi-tier intelligent data manager for Exascale (2021, 36 months, total cost **8M€**, G.A. n. 956748)
- **TEXTAROSSA** (EC H2020 RIA, EuroHPC-01-2019): Towards EXtreme scale Technologies and Accelerators for euROhpc hw/Sw Supercomputing Applications for exascale (2021, 36 months, total cost **6M€**, G.A. n. 956831)
- **EUPEX** (EC H2020 RIA, EuroHPC-02-2020): European Pilot for Exascale (2022, 48 months, total cost **41M€**, G.A. n. 101033975)
- **The European Pilot** (EC H2020 RIA, EuroHPC-02-2020, 42 months, total cost **30M€**, G.A. n. 101034126)
- **EUMaster4HPC** (EC H2020 RIA, EuroHPC-03-2020): Training and Education on HPC (2022, 48 months, total cost **7M€**)

# Tool bag (1)

- **StreamFlow** container-native Workflow Management System (WMS)
- **FastFlow** C++ programming library targeting multi/many-cores
- **CAPIO** research-grade software layer capable of coordinating legacy modules of data-driven large-scale software pipelines that cooperate using the I/O storage as a temporary communication buffer
- **Barbeque (BBQ)** run-time resource manager (RTRM) that enables node-level resource management across multiple applications
- **Celerity** run highly parallel applications on a cluster of accelerator nodes (extends SYCL)
- **DF-Threads** low level API which allows an efficient management of thread-level data flow
- **TAFFO** plugins for the LLVM Compiler Framework that enables automated precision tuning



# Tool bag (2)

- **COUNTDOWN** library for performance-neutral energy saving in MPI applications
- **GVirtuS**. Enables the execution of CUDA (and OpenCL) kernels on physical or virtual machines unprovided of general-purpose GPU acceleration
- **DagOnStar** enables the execution of direct acyclic graph (DAG) jobs on anything, ranging from the local machine to virtual HPC clusters hosted on private, public or hybrid clouds
- **CppPosits** supports mixed-precision and Posits arithmetic (LLVM compliant)
- **CRFlex** library of HW accelerator IPs
- **U-Therm3D** multiphysics-multiscale CFD methodology to handle complex phenomena such as the turbulent combustion occurring in gas turbine combustors
- **Vector Processing Hardware Acceleration** RTL implementation of vector processing units



Prof, arrivano un sacco di soldi col PNRR, li distribuiamo a pioggia a università e industria, vero?

**INVESTIAMO  
TUTTO!**

Centro Nazionale HPC,  
Big Data e Quantum Computing



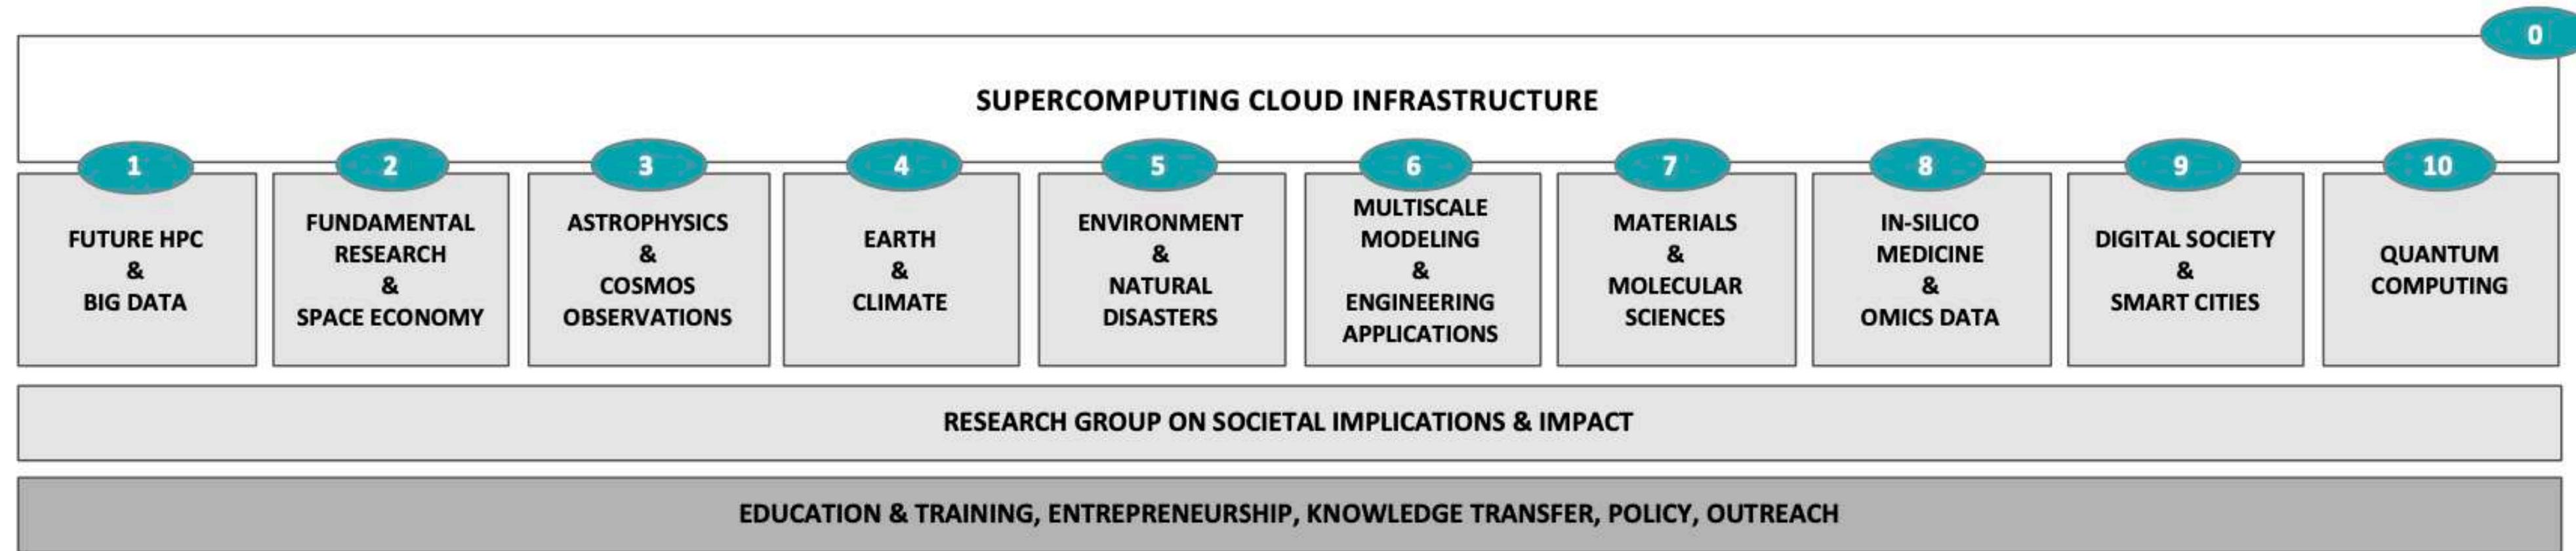
ICSC

XC

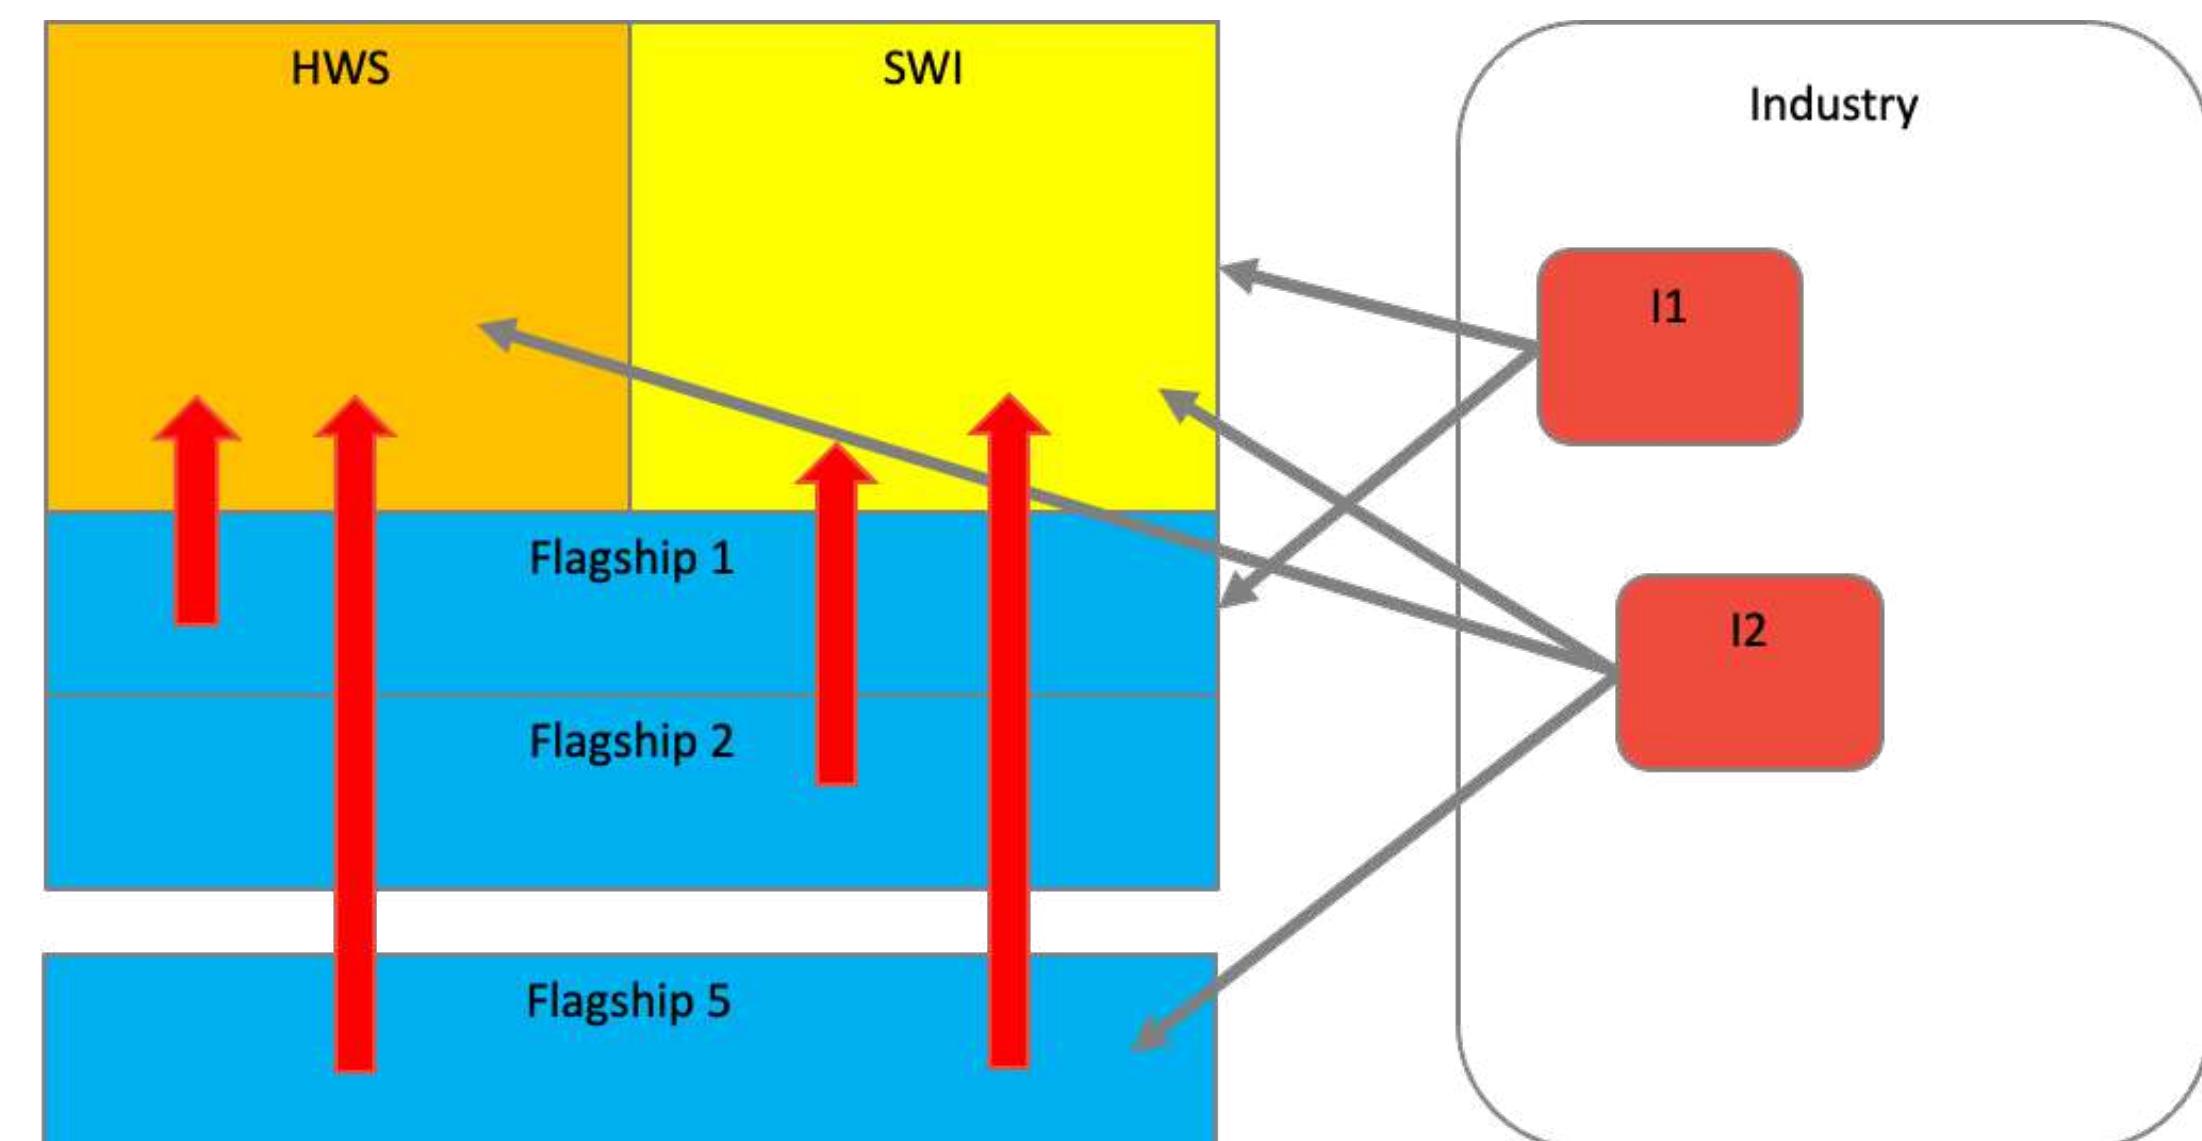
INFN

# FutureHPC spoke

## The technological spoke

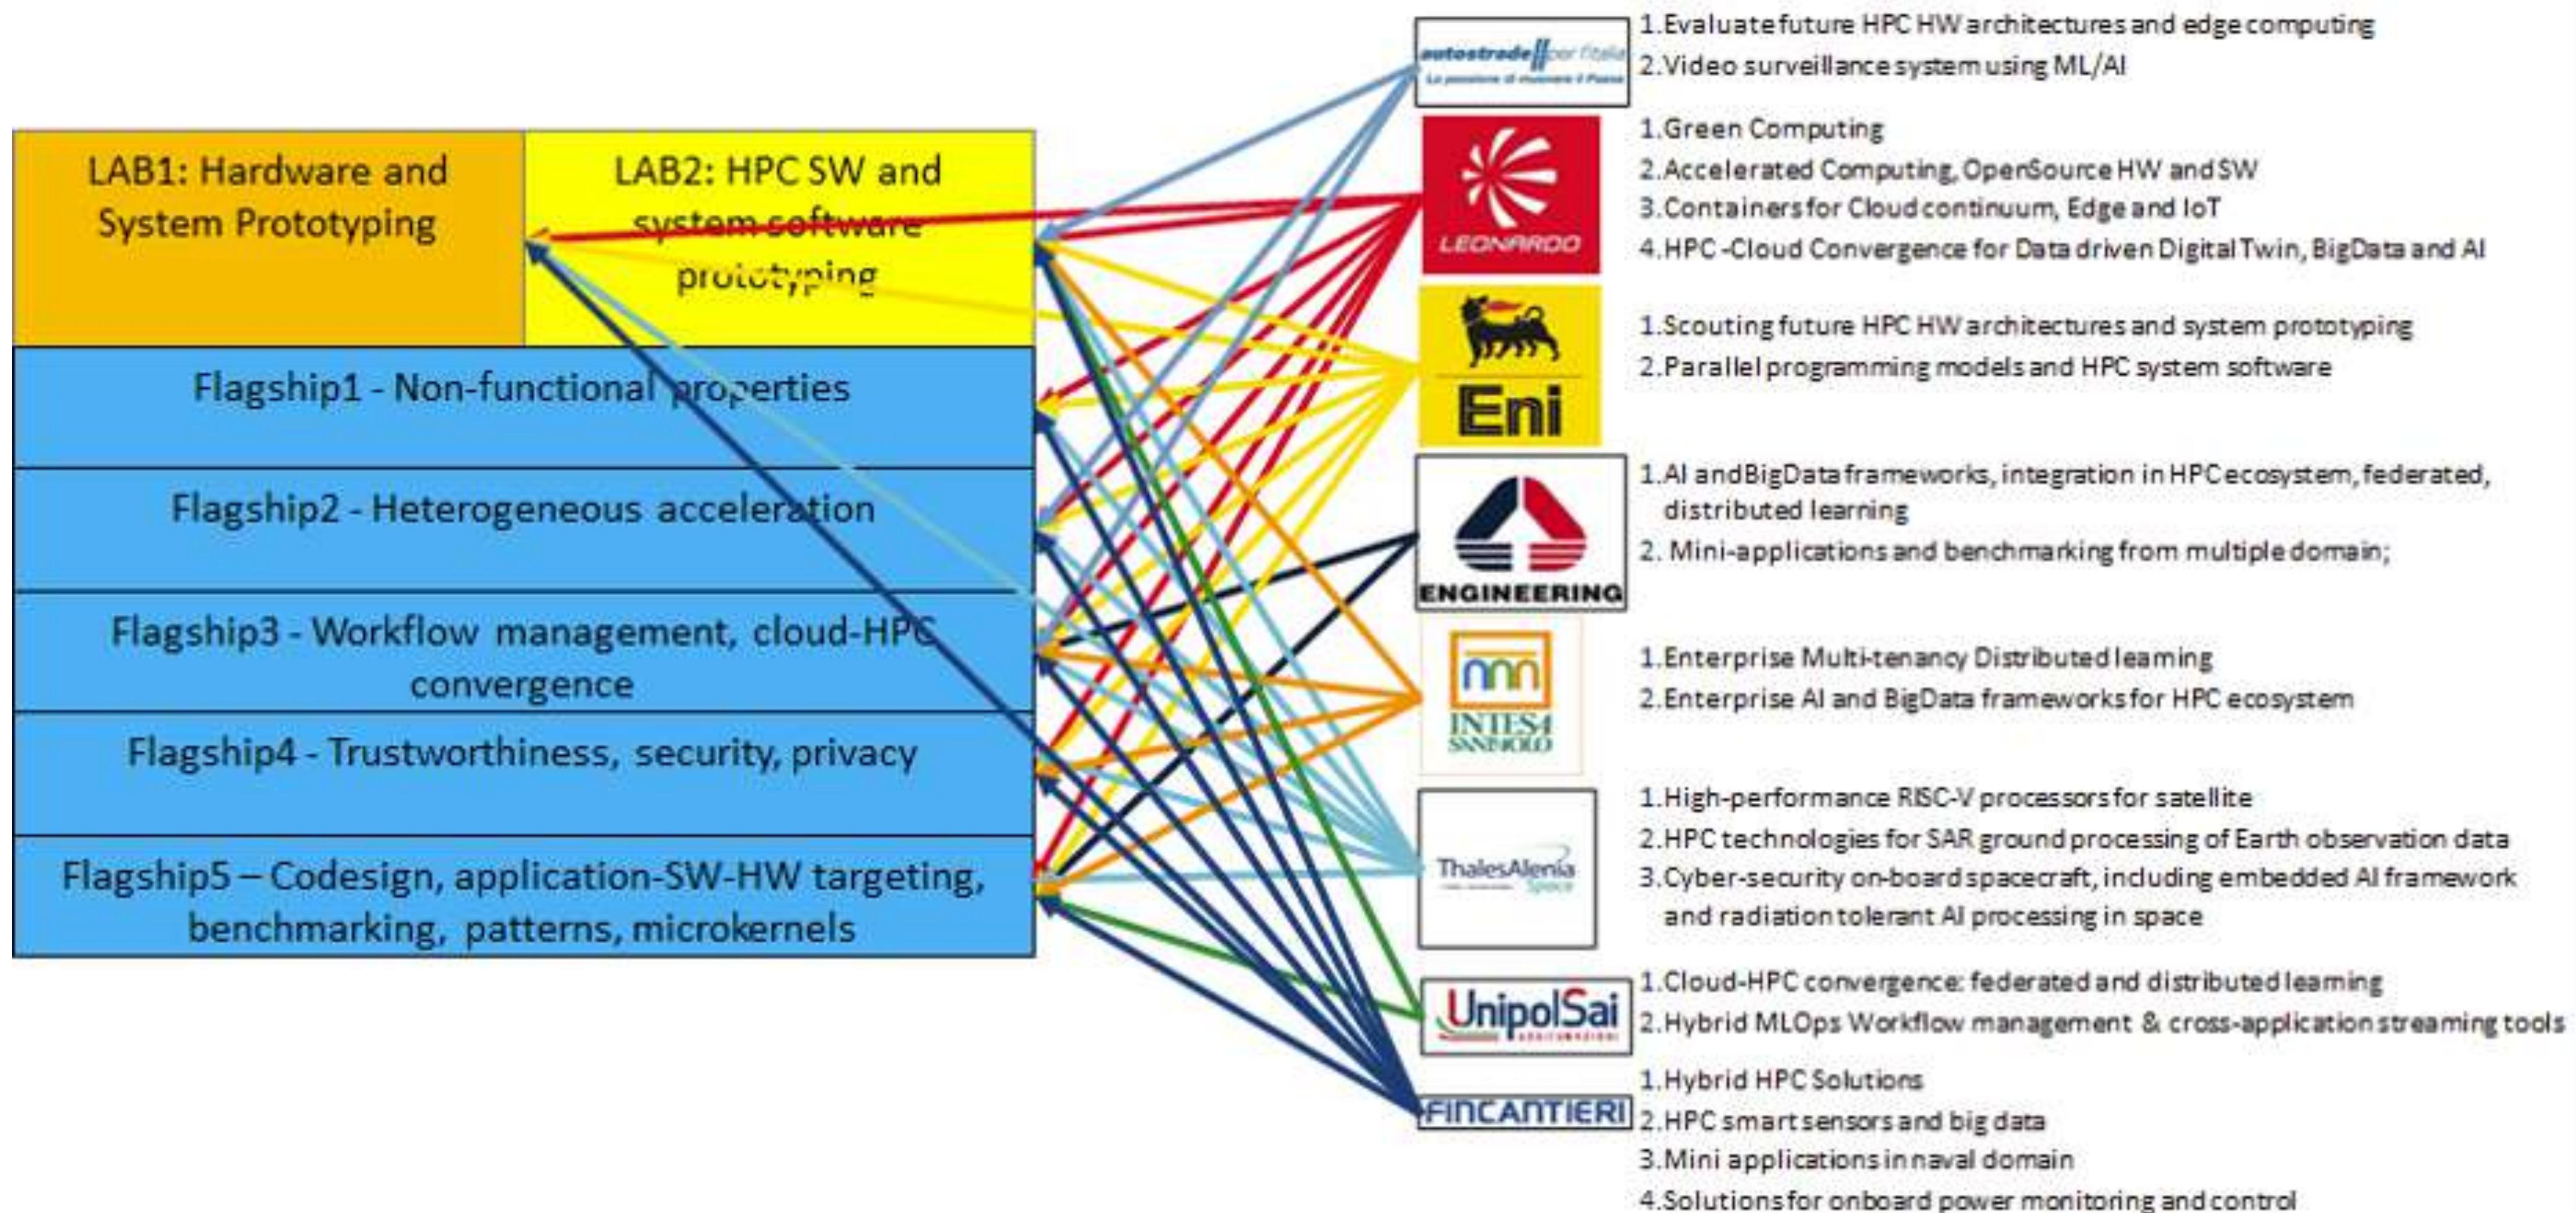


- Two “Spoke labs”: research + infrastructure investment
  - HWS Bologna (@Tecnopolis): **future HPC HW**
  - SWI Torino (@UNITO): **cloud-HPC SW & Integration**
- 5 Multi-partner “Flagship Initiatives”
  - Similar to EU projects (many partners, milestones, KPIs)
  - Funding for technology transfer and innovation projects
  - Cascade funding (beneficiaries excluded)



# Spoke1: FutureHPC

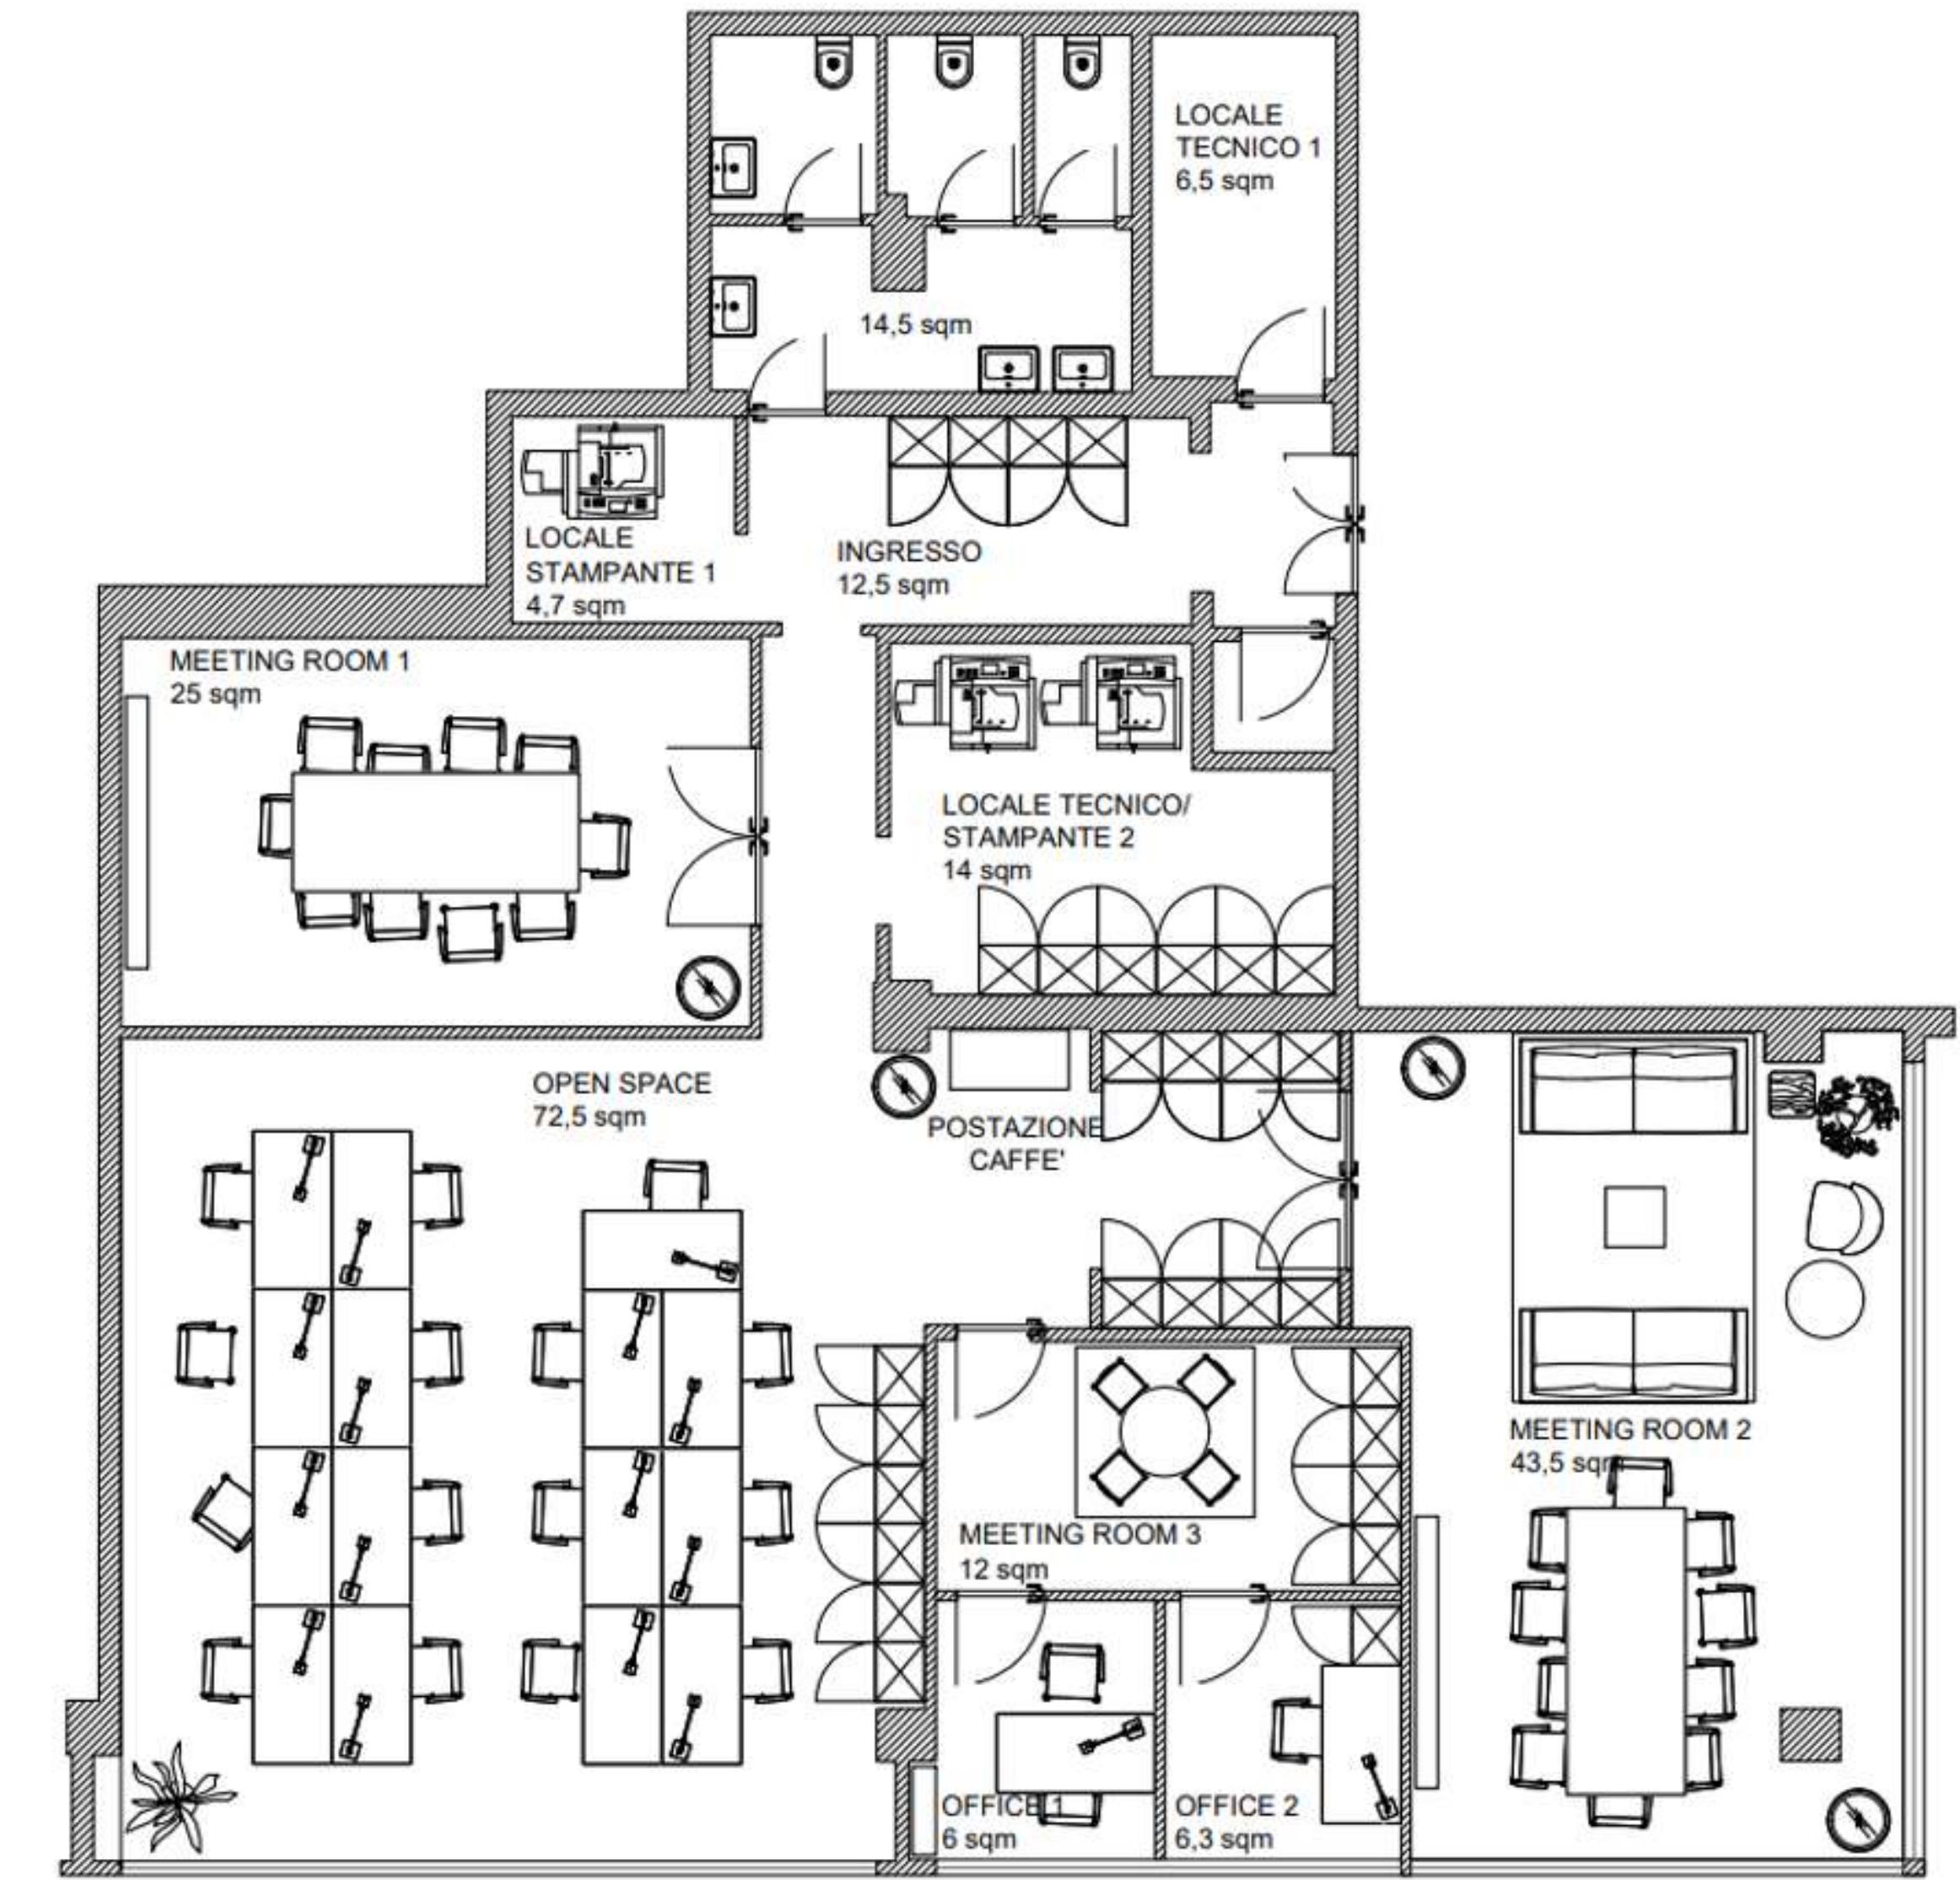
## A strong industrial partnership



# SWI

## SoftWare & Integration

- Progetto “lanterna” per UNITO strutturato come una "bottega" rinascimentale specializzata nel “High-Performance software engineering & integration”
- 350mq presso centro Piero della Francesca, in locali attigui al dipartimento di Informatica del Università di Torino
- Fra 3 anni, al Valentino ...



UFFICI CENTRO NAZIONALE INFORMATICA - UNITO

# FutureHPC

## Ricerca e formazione

- **La ricerca:** di sviluppare i metodi e le competenze per la progettazione di codici ad alte prestazioni che possano mantenere il loro valore nel tempo attraversando le normali evoluzioni tecnologiche (temi chiave: portabilità e la portabilità delle prestazioni del software).
- **La formazione:** La promessa che facciamo alle grandi aziende è quella di aiutali a trasformare alcuni dei loro servizi di calcolo avanzati (HPC, AI, etc) da beni da acquisire dalle compagnie tecnologiche Over-The-Top (cioè OPEX) in robusti asset industriali in grado di generare innovazione di prodotto. SWI avrà degli “allievi”, in parte dottorandi/assegnisti/RTDA UNITO, in parte dipendenti delle industrie che partecipano (sono 8 le grandi aziende già ingaggiate).

Prof, cosa facciamo adesso?

**QUELLO CHE FACCIAMO  
SEMPRE. MIGNOLO.**

**ANDIAMO A PRANZO!**

