Finanziato dall'Unione europea NextGenerationEU
Ministero dell'Università e della Ricerca
Italiadomani PIANO NAZIONALE DI RIPRESA E RESILIENZA
ICSC Centro Nazionale di Ricerca in HPC, Big Data and Quantum Computing

ICSC
Centro Nazionale di Ricerca in HPC, Big Data and Quantum Computing

Quasi interactive analysis of big data with high throughput:

Initial steps and future perspectives

Tommaso Diotalevi (UniBO), Francesco G. Gravili (UniSalento)
on behalf of WP2 UC2.2.2 members

ICSC Spoke 2 Annual Meeting - 19 December 2023
Casalecchio di Reno (BO)

WP1
WP2
WP3
WP4
WP5
WP6
ICSC-SPOKE2
Centro Nazionale di Ricerca in HPC,
Big Data and Quantum Computing

Want to know more?

More details in the general WP2 talk by **P.Lenzi** and **A.Annovi**

Experimental High Energy Physics

WP1 **WP2** WP3 WP4 WP5 WP6

Flagship usecase

UC2.2.2

ICSC-SPOKE2
Centro Nazionale di Ricerca in HPC, Big Data and Quantum Computing

Finanziato dall'Unione europea NextGenerationEU
Ministero dell'Università e della Ricerca
Italiadomani PIANO NAZIONALE DI RIPRESA E RESILIENZA
ICSC

Quasi interactive analysis of big data with high throughput

Spoke	2
WP	2, 5
Use case short name	Quasi interactive analysis of big data with high throughput
Use case ID	UC2.2.2
Expected Completion	31/8/2025

Approval workflow

Status	Version	Date	Submitter	Note	Signature
Draft	1.0	03/07/23	WP Leaders	First version	
Final Version	1.1	1/9/2023	WP Leaders		
Approved by Spoke Leaders	1.1	11/9/2023	Spoke Leaders		

cn1-spoke2-wp2-analysisfacility@lists.infn.it

Principal Investigators:

Tommaso Diotalevi

Francesco G. Gravili

Motivation

The upcoming high-luminosity phase at the **CERN Large Hadron Collider (LHC)** will require an increasing amount of computing resources.

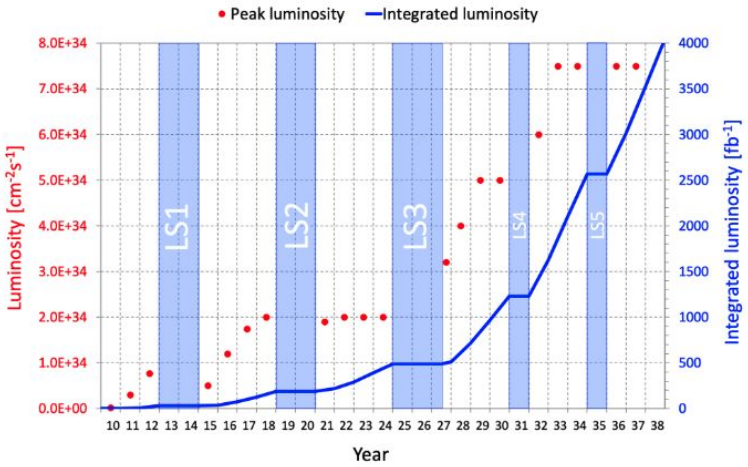

link

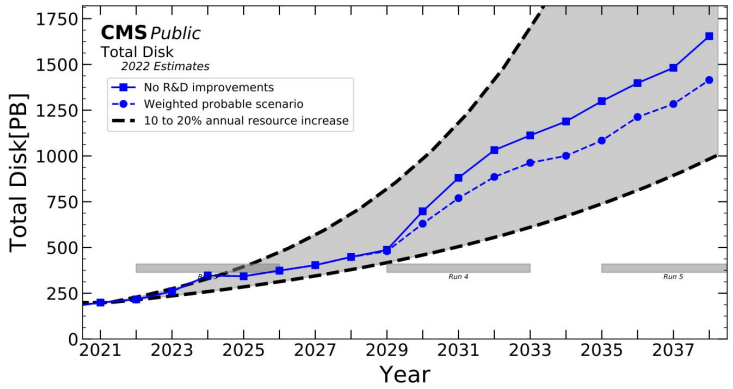

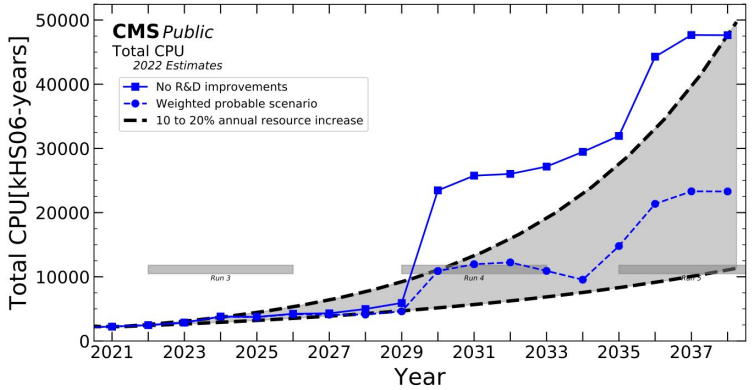

Higher rates of collision events → Higher demand for computing and storage resources

To better analyse this increasing amount of Big Data:

- Optimize the usage of CPU and storage;
- Promote the usage of better data formats;
- **Develop new analysis paradigms**!

Motivation

The upcoming high-luminosity phase at the **CERN Large Hadron Collider (LHC)** will require an increasing amount of computing resources.

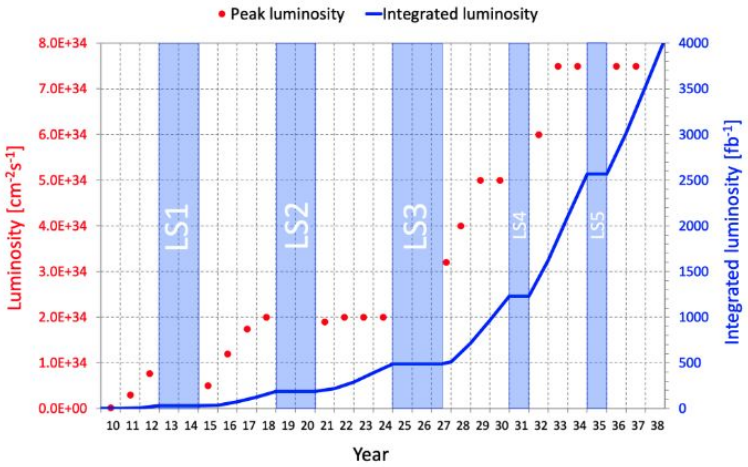

link

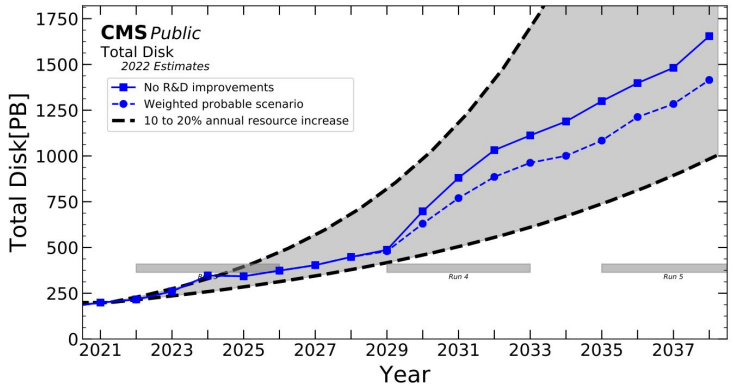

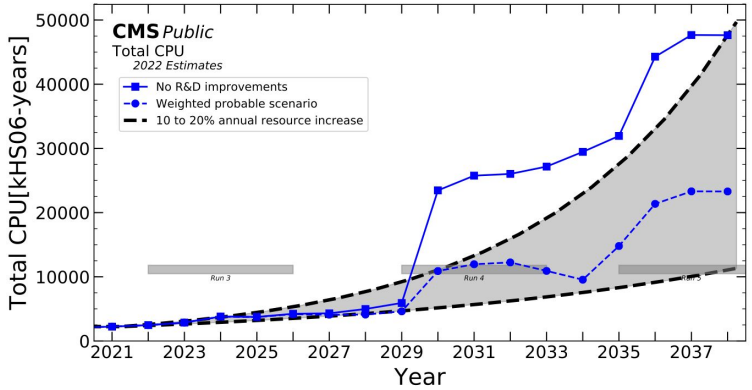

Higher rates of collision events → Higher demand for computing and storage resources

To better analyse this increasing amount of Big Data:

- Optimize the usage of CPU and storage;
- Promote the usage of better data formats;
- **Develop new analysis paradigms**!

→ "Quasi interactive analysis of big data with high throughput"

Common situation nowadays, extendable also to other Physics contexts…

Data processing stream in a HEP experiment

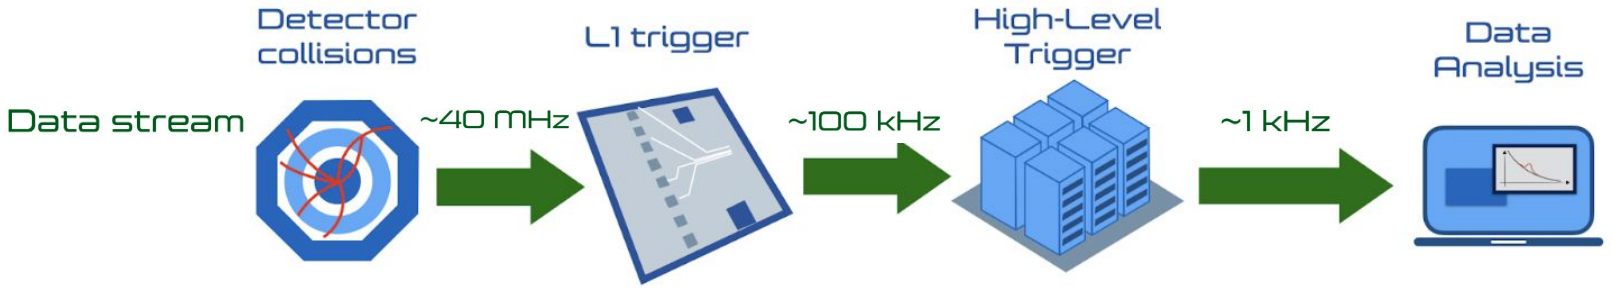

LHC example image taken from: link

High-throughput analysis within ICSC

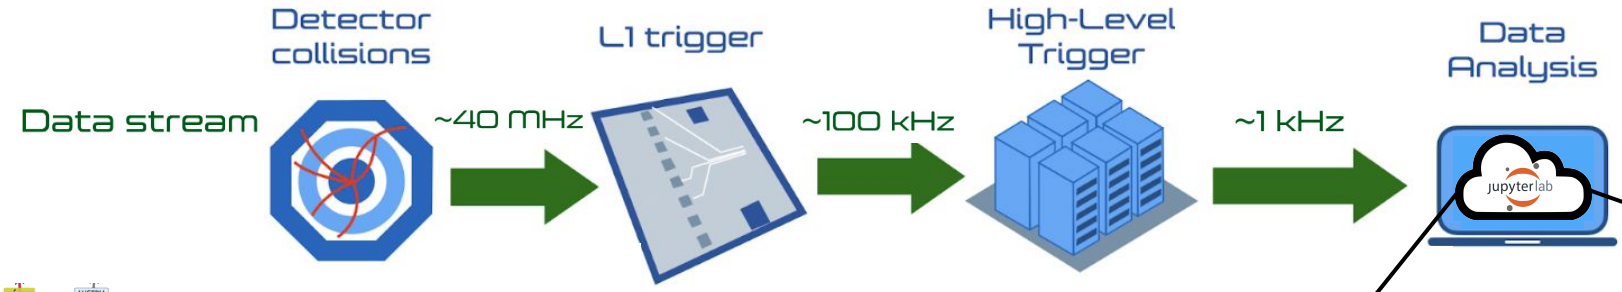

LHC example image taken from: link

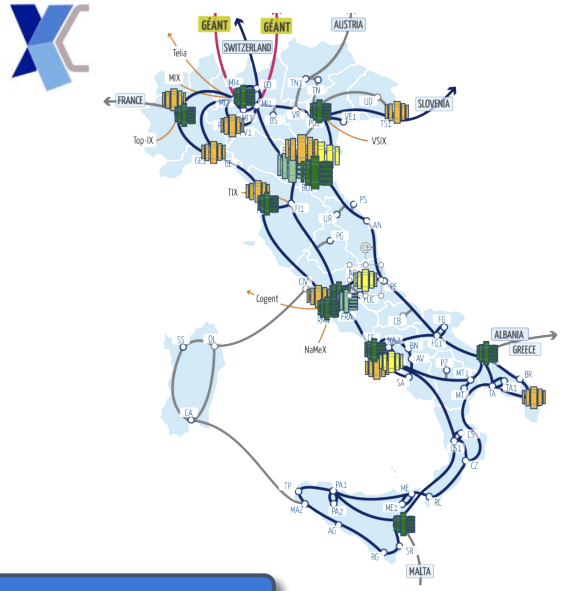

Resource monitoring dashboard

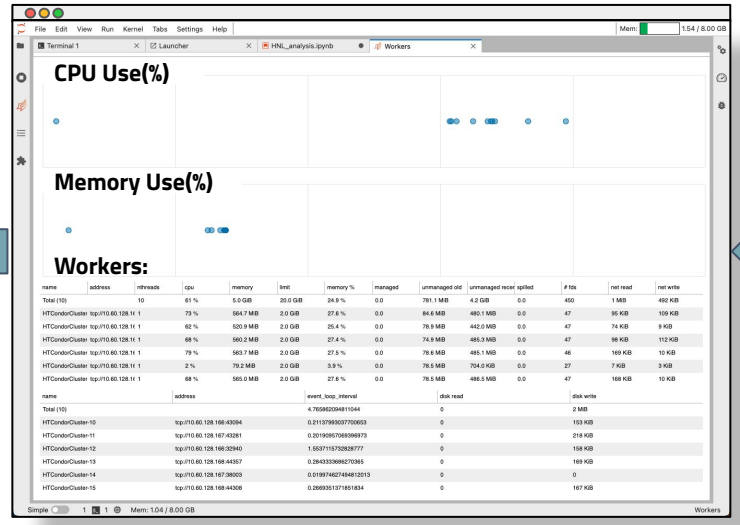

Analysis code

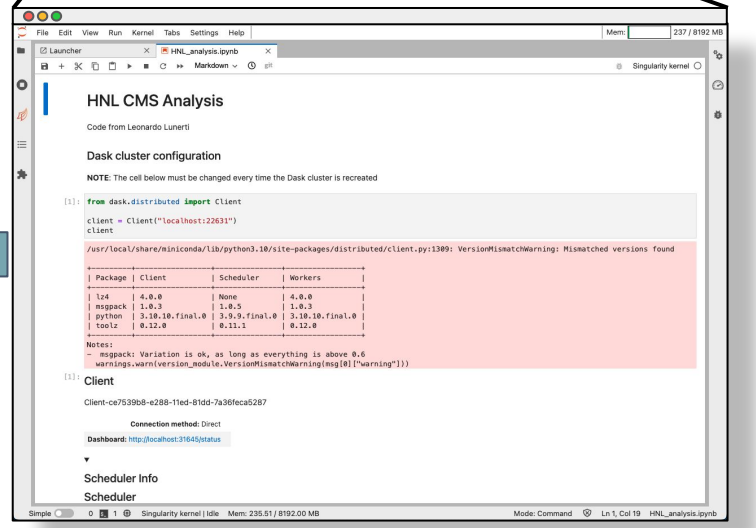

Want to know more?

More details of such platform in the lightning talks by G.Sabella and T.Tedeschi

Current activities in usecase UC2.2.2

Vector Boson Scattering ssWW analysis in hadronic tau and light lepton

Heavy Neutral Lepton search on heavy neutrinos in the D_S decays

Continuous Integration pipeline, triggering analysis execution on AF

Talk by M.Bartolini

top quark+MET analysis

Search of rare events in tau to 3 muons decay

Differential cross section measurement for ttbar inclusive production

Search for new phenomena in events with two opposite-charge leptons, jets and missing transverse momentum

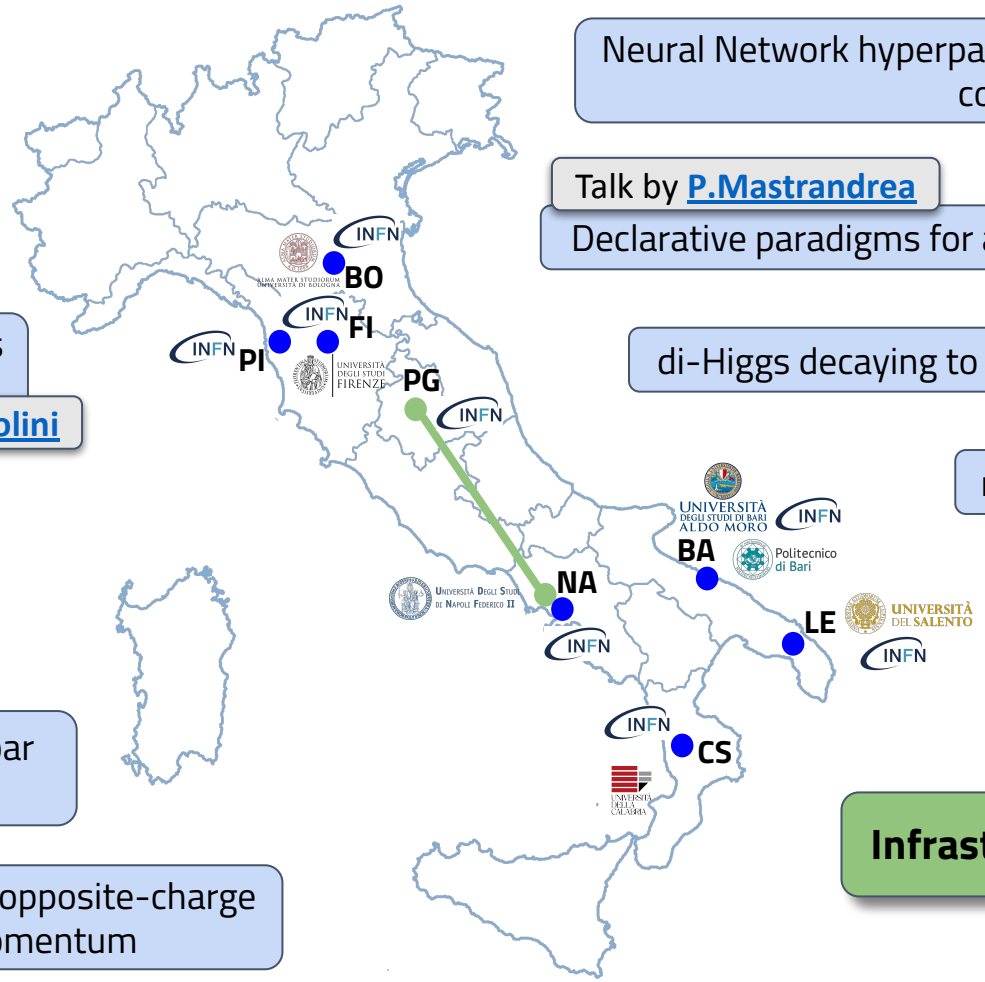

Neural Network hyperparameter optimisation applied to future colliders (FCC-ee)

Talk by P.Mastrandrea

Declarative paradigms for analysis description and implementation

di-Higgs decaying to two b quarks and two muons

muon detector performance analysis

Talk by A.D'Onofrio

Benchmark interactive analysis for future colliders (FCC-ee)

Infrastructural support WP5 backbone

and more to come...

Current new-hired personnel under UC2.2.2

INFN Pisa (PI):
- Paolo Mastrandrea, INFN technologist

University of Bologna (BO):
- Tommaso Diotalevi, Researcher (RTDa)

INFN Perugia (PG):
- Tommaso Tedeschi, INFN PostDoc

University of Firenze (FI):
- Matteo Bartolini, PostDoc

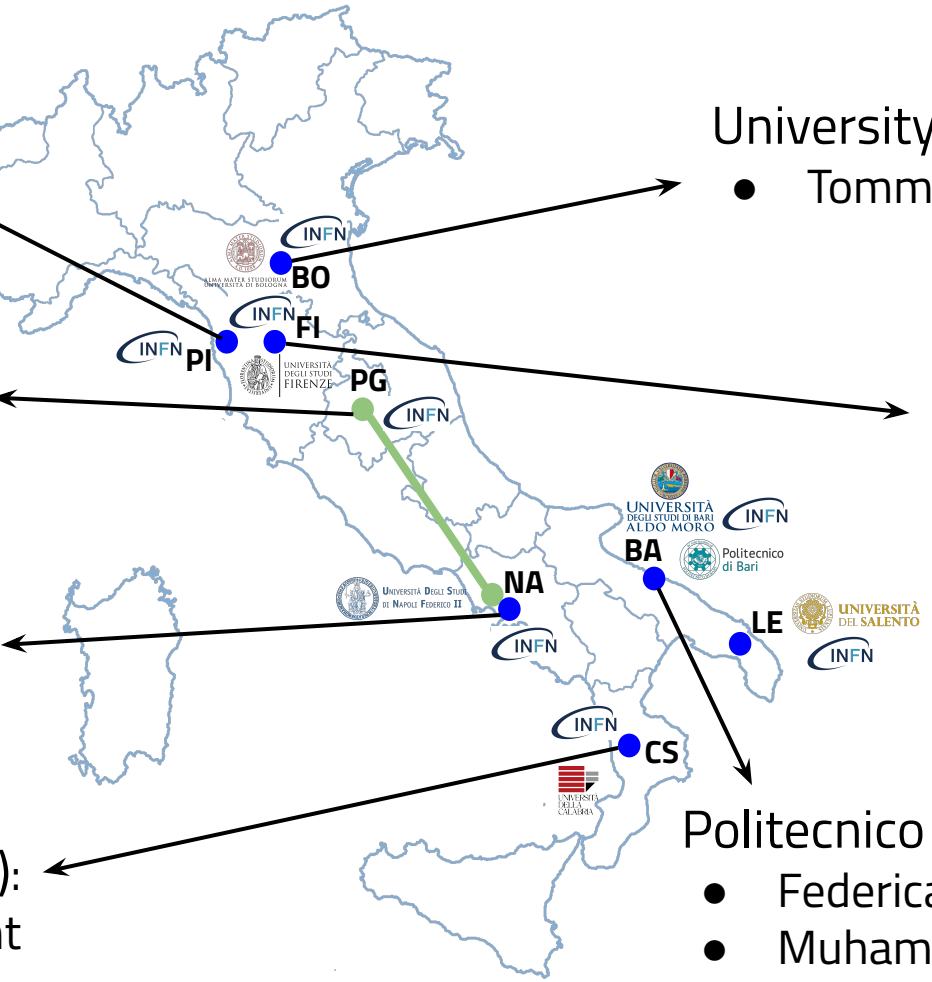

INFN Napoli (NA):
- Adelina D'Onofrio, INFN technologist

University of Calabria (CS):
- Alessandro Tarasio, PhD student

Politecnico of Bari (BA):
- Federica Maria Simone, Researcher (RTDa)
- Muhammad Anwar, PhD Student

Conclusions

What happened this year:

- Definition of the Flagship activity UC2.2.2 inside WP2, with definition of KPIs and milestones;
- Surveying all the participating institutions and gathering all the physical applications;
- Exploiting the existing infrastructure test-beds (in synergy with WP5), with some preliminary study for performance comparison speed-up;
- Quantify and formalize the amount of resources required by the usecase, via RAC.

What will happen next:

- Working on the identification of metrics and benchmarks for performance evaluation among the different applications;
- Using the central deployment of the high-throughput infrastructure (in synergy with the WP5 goals) with ICSC resources;
- Fulfilling the established deliverables.

Finanziato dall'Unione europea NextGenerationEU
Ministero dell'Università e della Ricerca
Italiadomani PIANO NAZIONALE DI RIPRESA E RESILIENZA
ICSC Centro Nazionale di Ricerca in HPC, Big Data and Quantum Computing

BACKUP

KPI - Key Performance Indicator

KPI ID	Description	Acceptance threshold
KPI2.2.2.1	Implementation of *N* data analyses in the AF	$N \geq 2$
KPI2.2.2.2	Reference documentation of the AF	≥ 1 dedicated web site
KPI2.2.2.3	Hands-on workshops for AF users	≥ 1 workshops
KPI2.2.2.4	Scaling up the testbed AF infrastructure, serving *k* tenants, for a total of *N* data analyses	≥ (200·*M*) cores
KPI2.2.2.5	Talks at conferences/workshops about AF activities	≥ 1 talk

RAC resource request

7.1 Resources granted by INFN-CLOUD (PaaS)

vCPU (number of vCores and requested allocation time)

Number	Time	vCORE	Notes
*	*	*	time in hours
4	17520	32	

Number of requested GPU and allocation time

Number	Time	GPU	Notes
*	*	*	NO GPU REQUIRED. The numbers were entered due to a problem with the form
1	1	1	

vCPU (number of vCores and requested allocation time)

RAM per VCORE	Notes
*	Memory in GB
4	

☐ Software used or required, including preferred Cloud services

I.e.: Kubernetes-as-a-service, Jupyter Notebook as a Service, Private Container Image Registry, Spark and Grafana as a service, Dropbox-like sync-and-share service

Software used or required	Notes
	Total core hours: 4*17520*32 = 2242560. This kind of resource does not follow time evolution, to be considered constant for the entire project duration.

7.2 Resources granted by INFN-GRID (Batch processing)

CPU (number of Cores and requested allocation time)

Number	Time	vCORE	Notes
*	*	*	time in hours
21	**17520**	**32**	

RAM Requirements

RAM per VCORE	Notes
*	Memory in GB
4	

Software used or required	Notes
	Total core hours: 21*17520*32 = 11773440. This kind of resource follows a time evolution, where: **- first 6 months: 876000 core hours, to be considered as 1000 cores * 20% of time (1000 cores distributed in *5* data analyses);** **- following 4 months: 730000 core hours, to be considered as 1000 cores * 25% of time (1000 cores distributed in *5* data analyses)** **- following 4 months: 730000 core hours, to be considered as 1000 cores * 25% of time (1000 cores distributed in *5* data analyses)** **- last 10 months: 9437440 core hours, to be considered as 5000 cores * ~26% of time (5000 cores distributed in *5* data analyses)** **These nodes, on N CPU systems without the requirement of fast inter-node infiniband communication (HTC-like), must have access to permanent storage.**

High throughput data analysis platform

- Access to a single Hub and authentication via token (using IAM instance)
- Based on standard industry technologies (Jupyter, Docker, Kubernetes)
- Configurable python Kernel (with user-specific software)
- Containerized workspace (loaded from public registries)
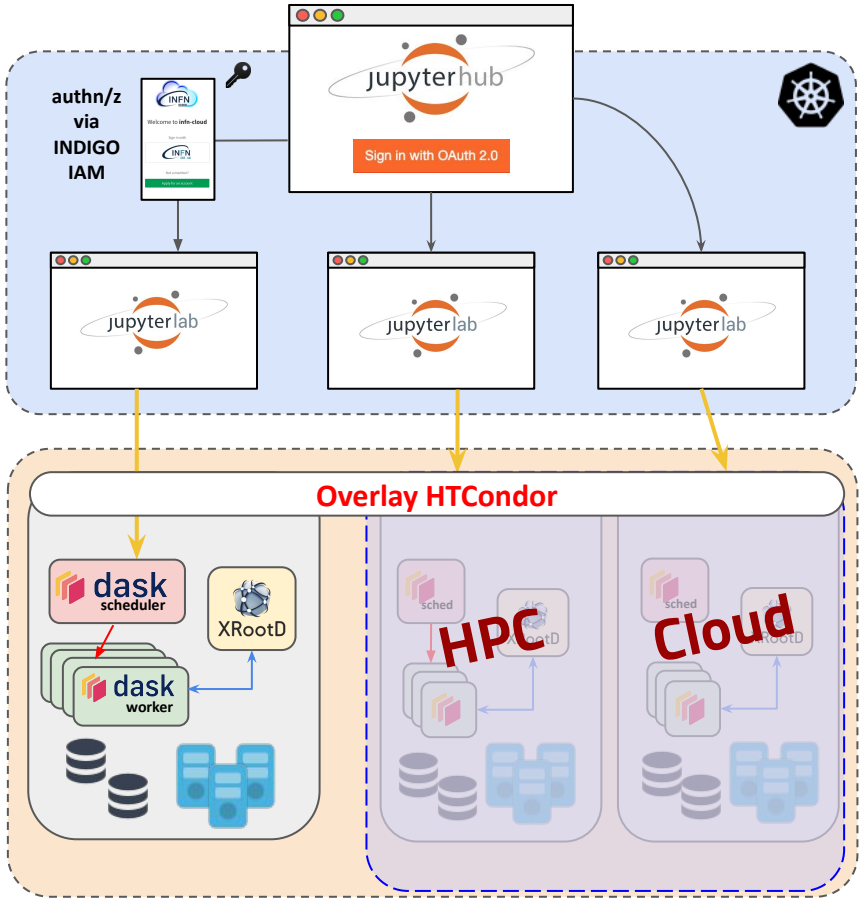
- HTCondor based overlay (also available standalone)
- DASK library (python) for distributed code execution (among the available nodes)
- Possible usage on heterogeneous resources (HTC/HPC/Cloud)
- Interface with Worldwide LHC Computing Grid (xrootd, WebDAV, ...)