



Finanziato  
dall'Unione europea  
NextGenerationEU



Ministero  
dell'Università  
e della Ricerca



Italiadomani  
PIANO NAZIONALE  
DI RIPRESA E RESILIENZA



# Offloading of computing payloads to HPC resources: concepts and synergies

[D. Ciangottini](#) for WP5



Spoke3 General meeting, Biodola 07/05/2024

# Outline

- Recap and use cases
- Sinergies
- Outlook and plans

Finanziato dall'Unione europea NextGenerationEU

Ministero dell'Università e della Ricerca

Italiadomani

ICSC

## User stories and call for volunteer

**Container based use cases are the primary targets**

**"I have my analysis working in a container, just bring that to N GPU node for me"**

My container → DataCloud Offload technology (INFN) → HPC GPU Node

**"I want to develop an algorithm on GPU/FPGA, please spawn a containerised environment where I can code"**

**We can start playing with the system NOW-ish (at least as soon as we have a seed of resources to use). We are collecting interested use cases and get started with the experimentation.**

My notebooks → DataCloud Offload technology (INFN) → Jupyter → HPC GPU Node

ICSC Italian Research Center on High-Performance Computing, Big Data and Quantum Computing

Missione 4 • Istruzione e Ricerca

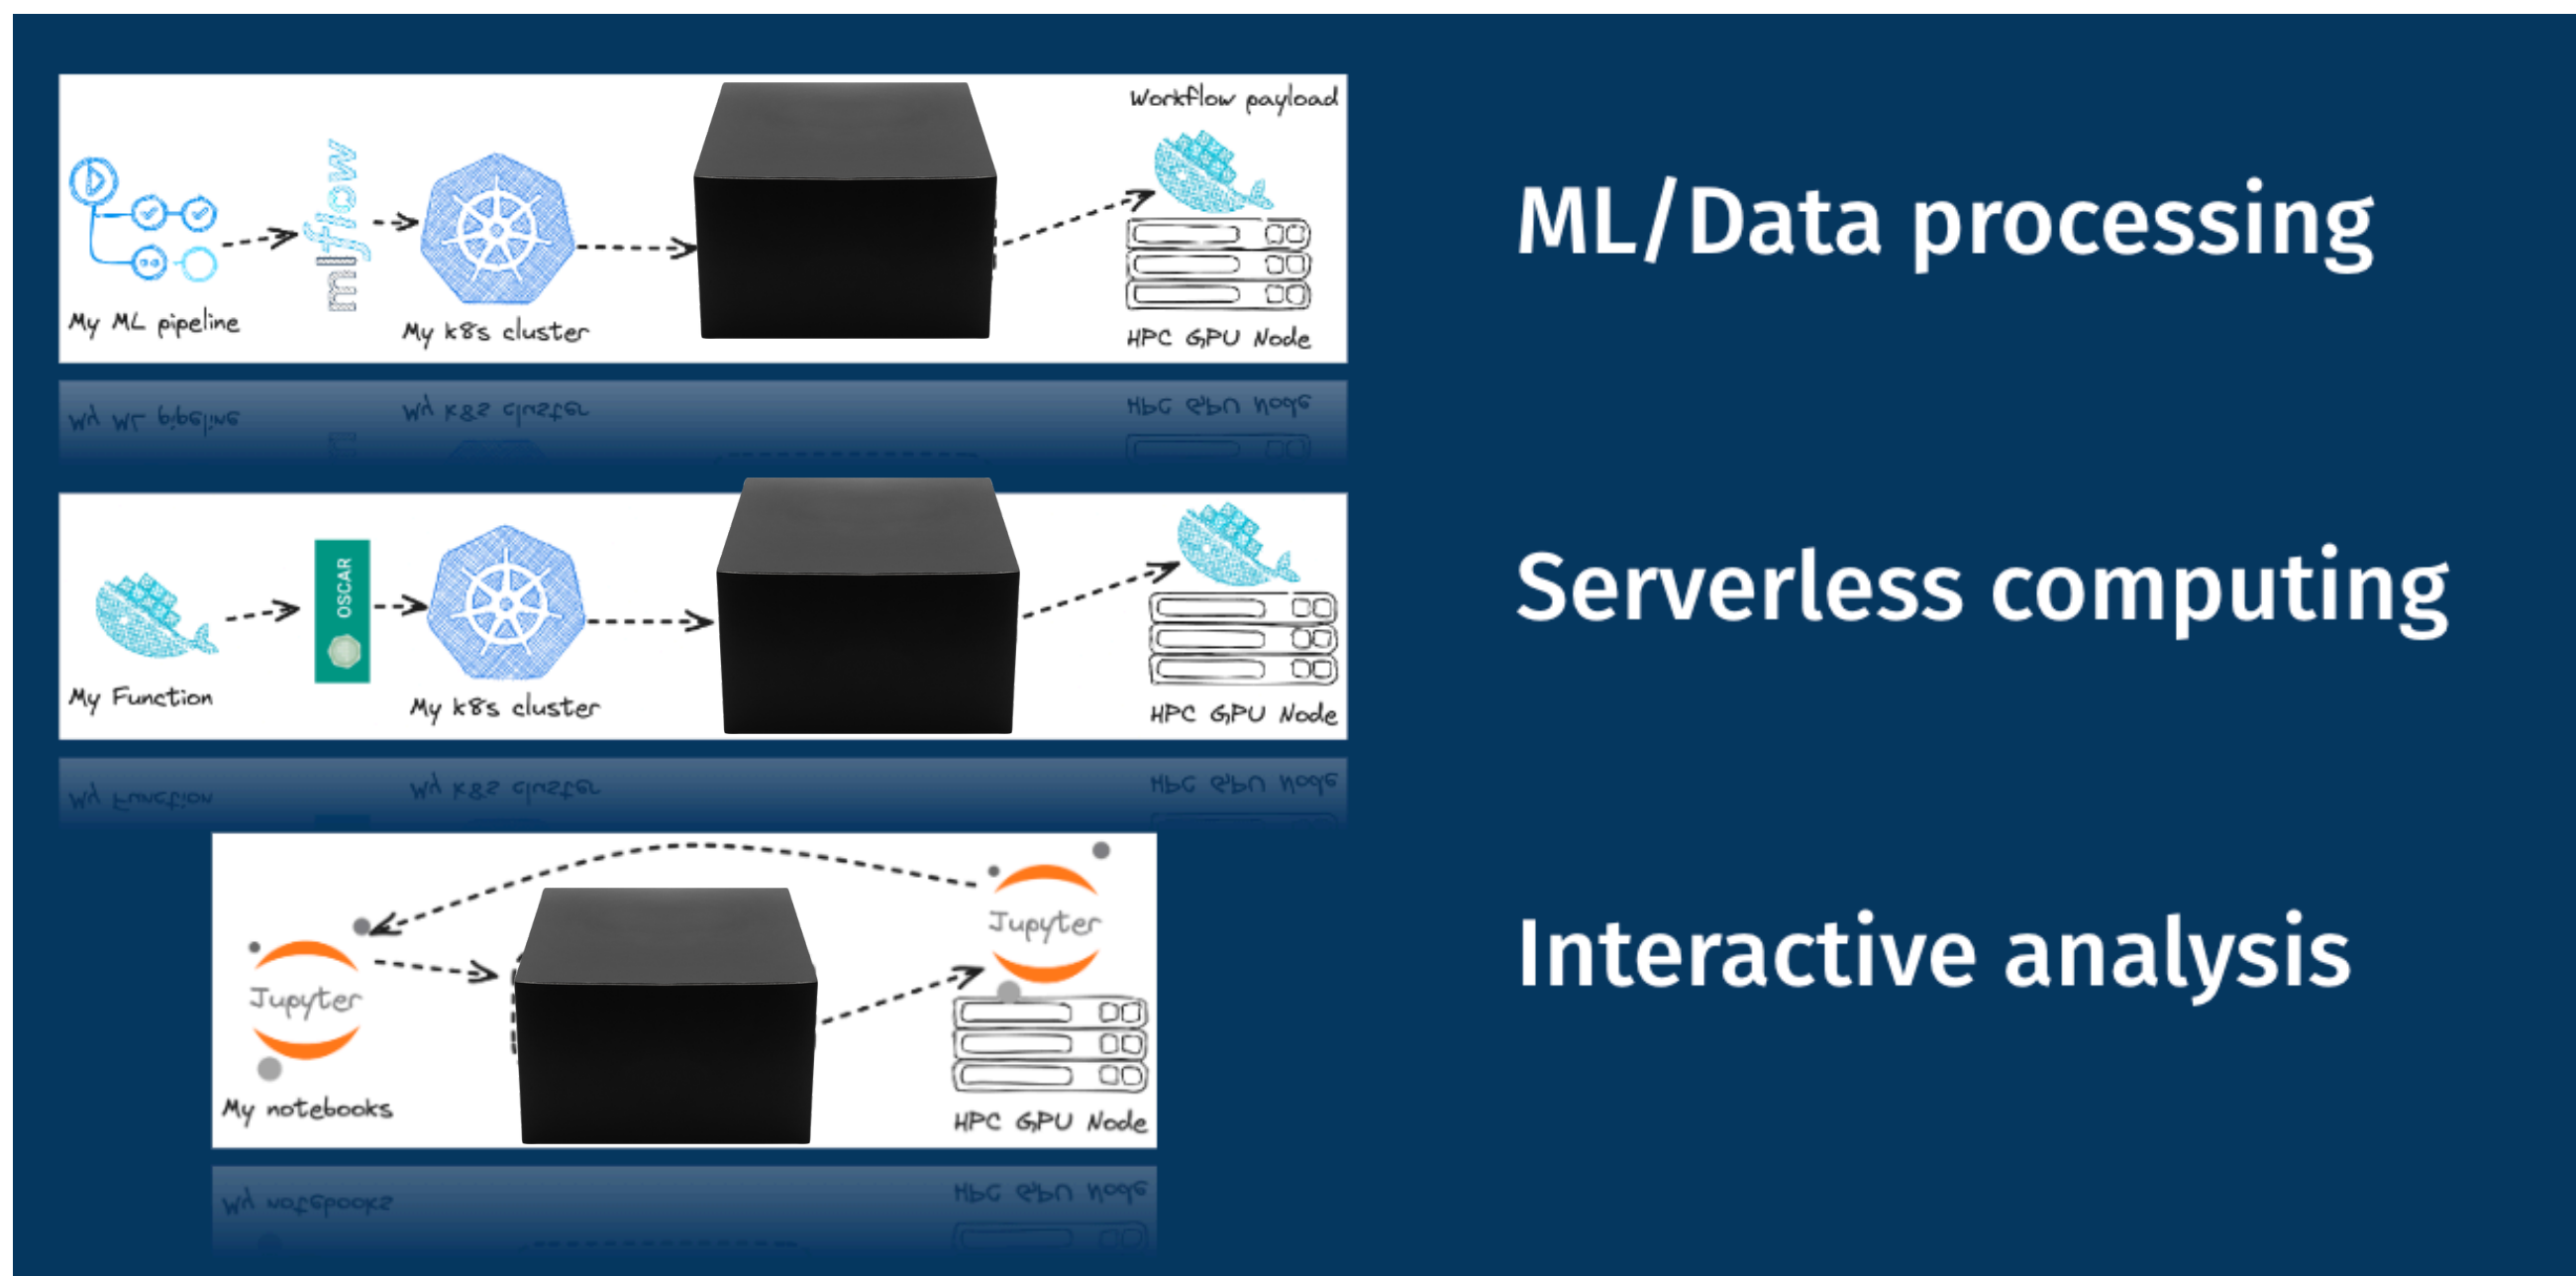
**We were here**  
**(last tech workshop...)**

The background of the slide is a deep blue color. On the left side, there is a vertical strip of abstract light trails and dots in various shades of blue and white, creating a sense of depth and movement. The text is positioned on the right side of the slide.

# **Use cases: a recap**

# A black box to manage to manage our payloads everywhere

e.g. accessing GPUs on-demand "seamlessly", no matter their location



ML/Data processing

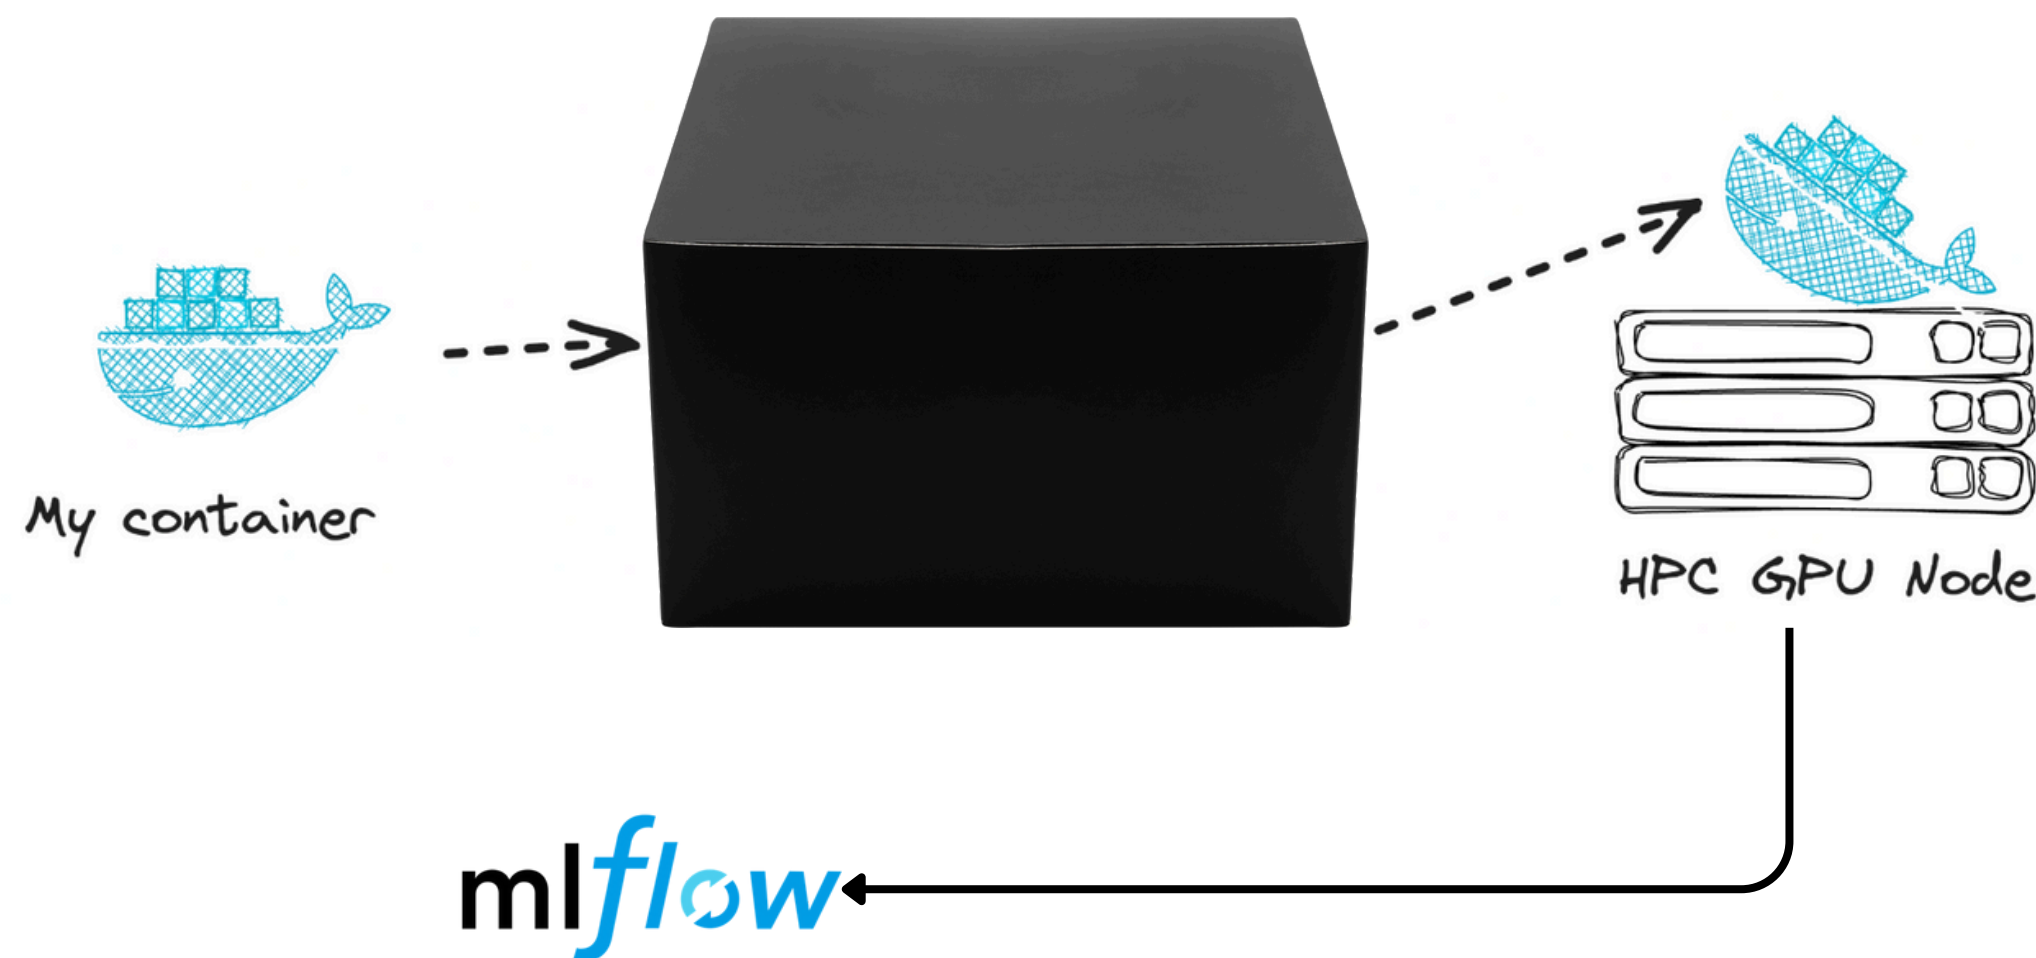
Serverless computing

Interactive analysis



## User story 1:

I have a container and I want to execute it on a node with N GPUs




### Example: Tracking ML training

- Submitting one or more training payloads to a remote machine with GPUs
- Track the performances of the training
- Send tracking information to a central service (along with ML final model as well)
- Retrieve the information from either a CLI or a web browser

## User story 2:

Interactively develop in an notebook that lives "somewhere else"




### An HUB for tinkering your code

- Spawn a JupyterLAB session on a remote machine (even in a HPC center)
- Use a default container image for that, or a custom one of yours
- Get access either through browser or terminal to that session
- Develop and test interactively your code in a on-demand environment

## User story 3:

I have my own framework, I want it to delegate heavy lifting "somewhere else"




### Analysis pipelines

- Framework evolution toward cloud-based resource orchestration is a common pattern lately
- We must be able to leverage this to hide the complexity of the infrastructure
- In other words, we should get something compatible with the most adopted container orchestrator out there.....

## User story 3:

I have my own framework, I want it to delegate heavy lifting "somewhere else"



### Analysis pipelines

- Framework evolution toward cloud-based resource orchestration is a common pattern lately
- We must be able to leverage this to hide the complexity of the infrastructure
- In other words, we should get something compatible with the most adopted container orchestrator out there.....



The background is a deep blue gradient. On the left side, there is a complex pattern of light trails and dots. These elements appear to be digital or data-related, with lines of light converging towards a central point and then diverging outwards. The dots are small, bright blue spheres of varying sizes, some appearing as blurred streaks due to motion or depth. The overall effect is one of dynamic energy and connectivity.

**Sinergies**

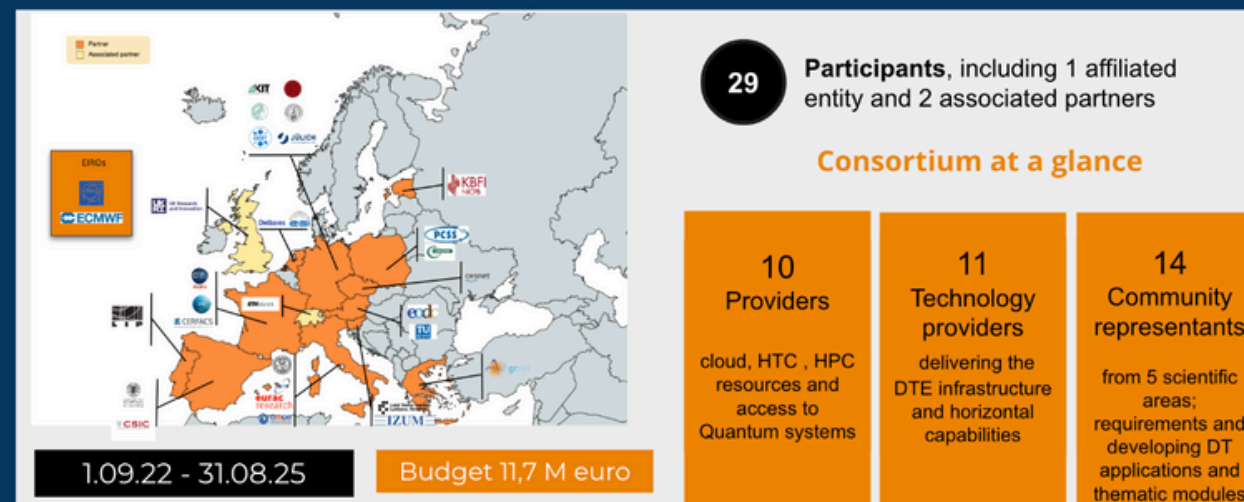
## interTwin and INFN:

striving to get an efficient interface to match the evolution of analysis frameworks with a seamless resource brokering.

### interTwin: designing a Digital Twin Engine

Very similar set of requirements with an even stronger need for cloud native framework integration with HPC centers (kind of a cloud-prohibitive environment)


**A Digital Twin Engine should be able to provide a platform capable of distributing DT applications to remote sites, maintaining an interoperable (cloud) interface.**



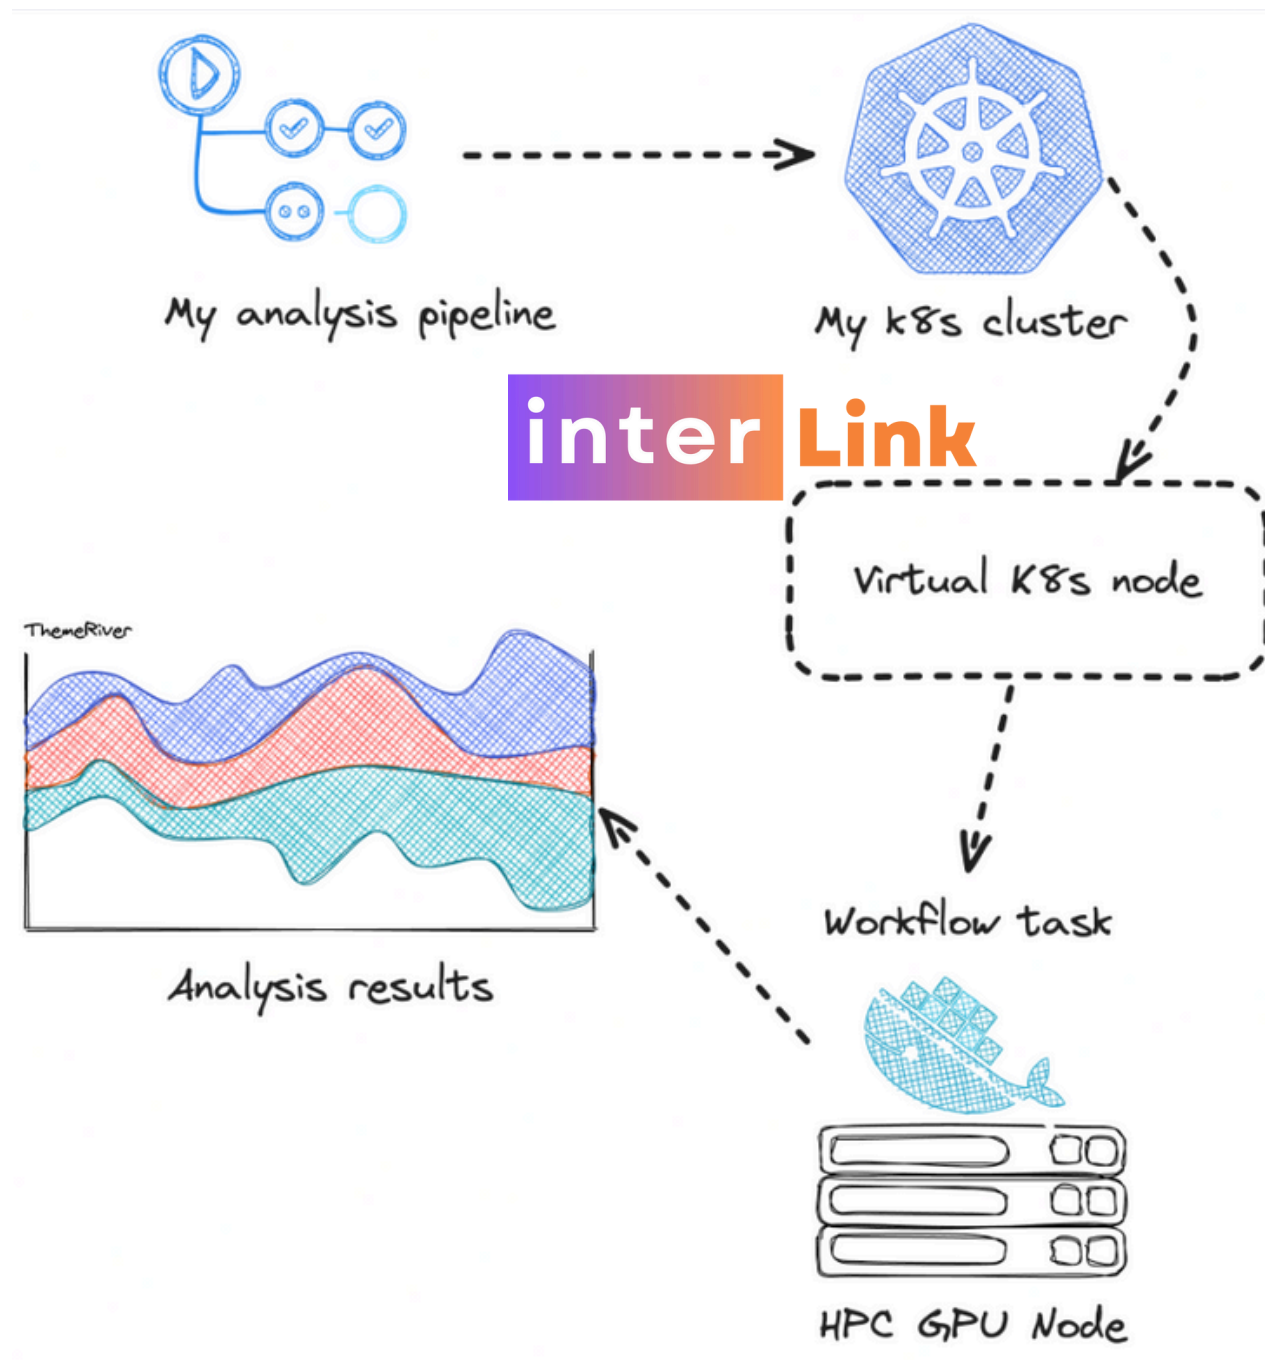
### Developed a first prototype for:

- Getting a seamless distribution of data analysis payloads over EuroHPC centers
- Validating Digital Twin Engine use cases against the prototype
- Extending the volunteer base of both the use cases and the resource providers (e.g. additional HPC and computing centers in general)

What's the tool that has been developed?



# A tool to abstract the execution of a container: InterLink translates a container into a job submitted to the HPC SLURM queue



"in a nutshell"

**A cloud-native friendly initiative:** based on known open technologies

**Not only HPC:** in fact, different initiative are giving it a stab for remote execution of containers over several remote services

**A recent talk/demo at the KubeCon EU 24 to get an idea**

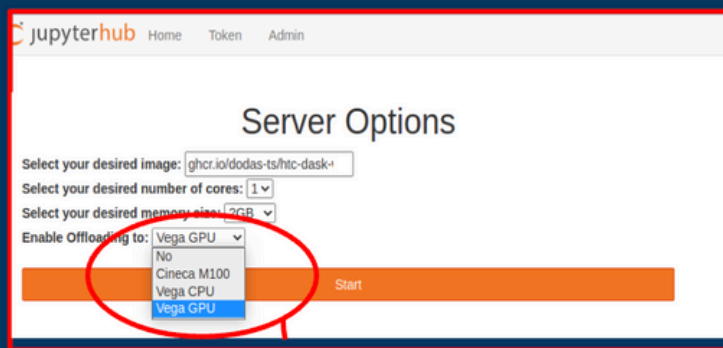
**Documentation**

The background is a deep blue gradient. On the left side, there are numerous thin, curved lines of light that appear to be receding into the distance, creating a sense of depth and motion. These lines are punctuated by small, bright blue dots, some of which are slightly out of focus, giving the impression of a digital or data-driven environment.

**Possible scenarios...**  
**~today.**

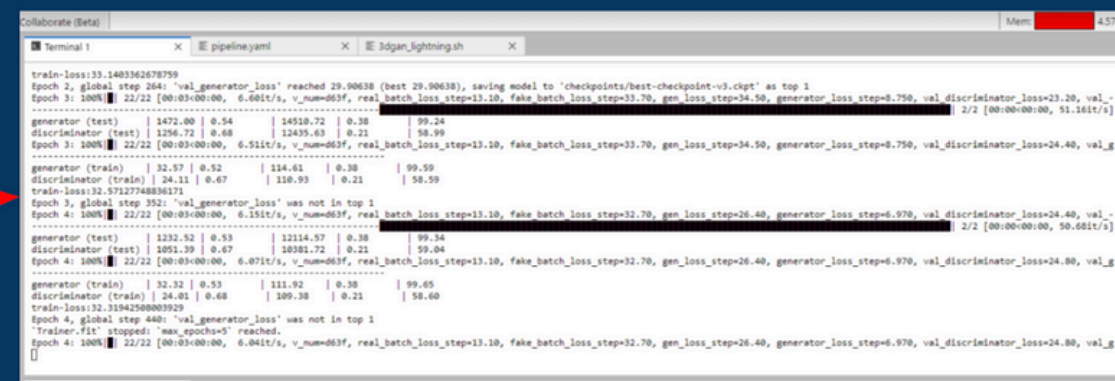
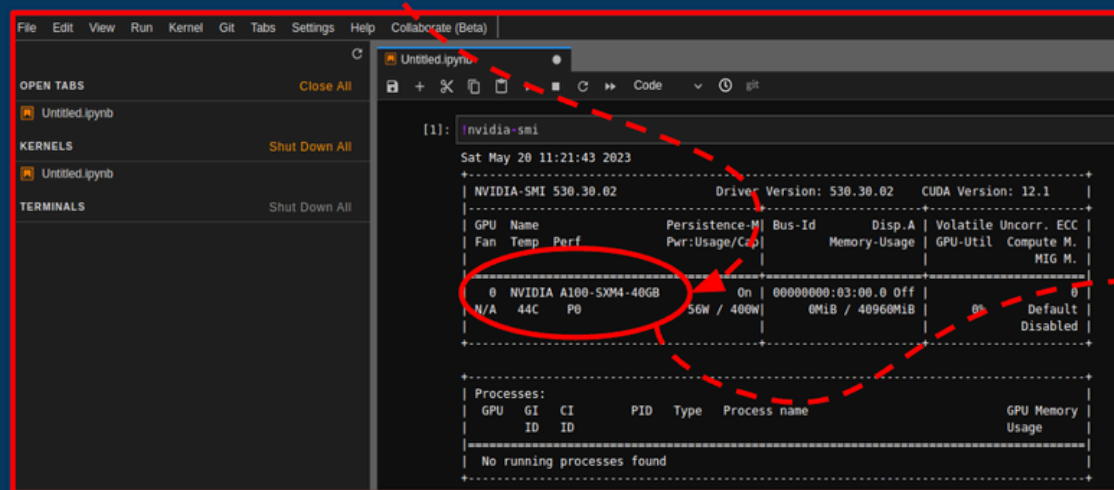
# It might be technically possible to reproduce the interTwin PoC meaning to "offload" from an hub hosted on the cloud to HPC resources

## Interactive prototyping



Successfully spawning on-demand jupyterLab instance on HPC along with other more "cloud-ish" instance on K8s.

Crucial for experimenting ML workflows interactively before scaling out!



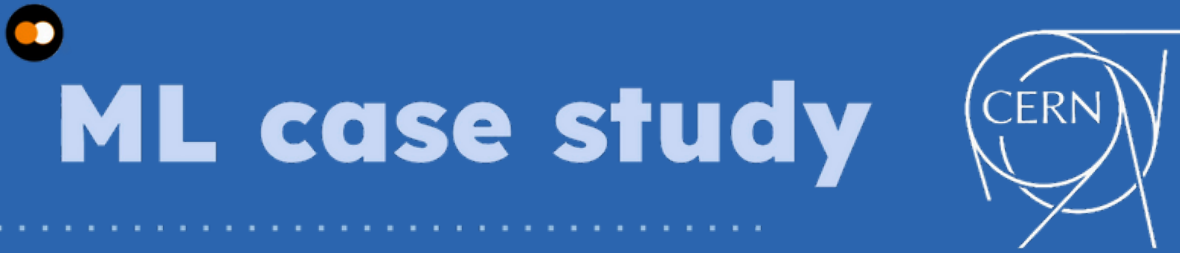
## Access to an extensible HUB

- Hosting core services on cloud resource
- Extend the resource pool for interactive sessions over HPC and, more generally, distributed resources

There are then cases where automation and scalability comes without (or after) interactive development

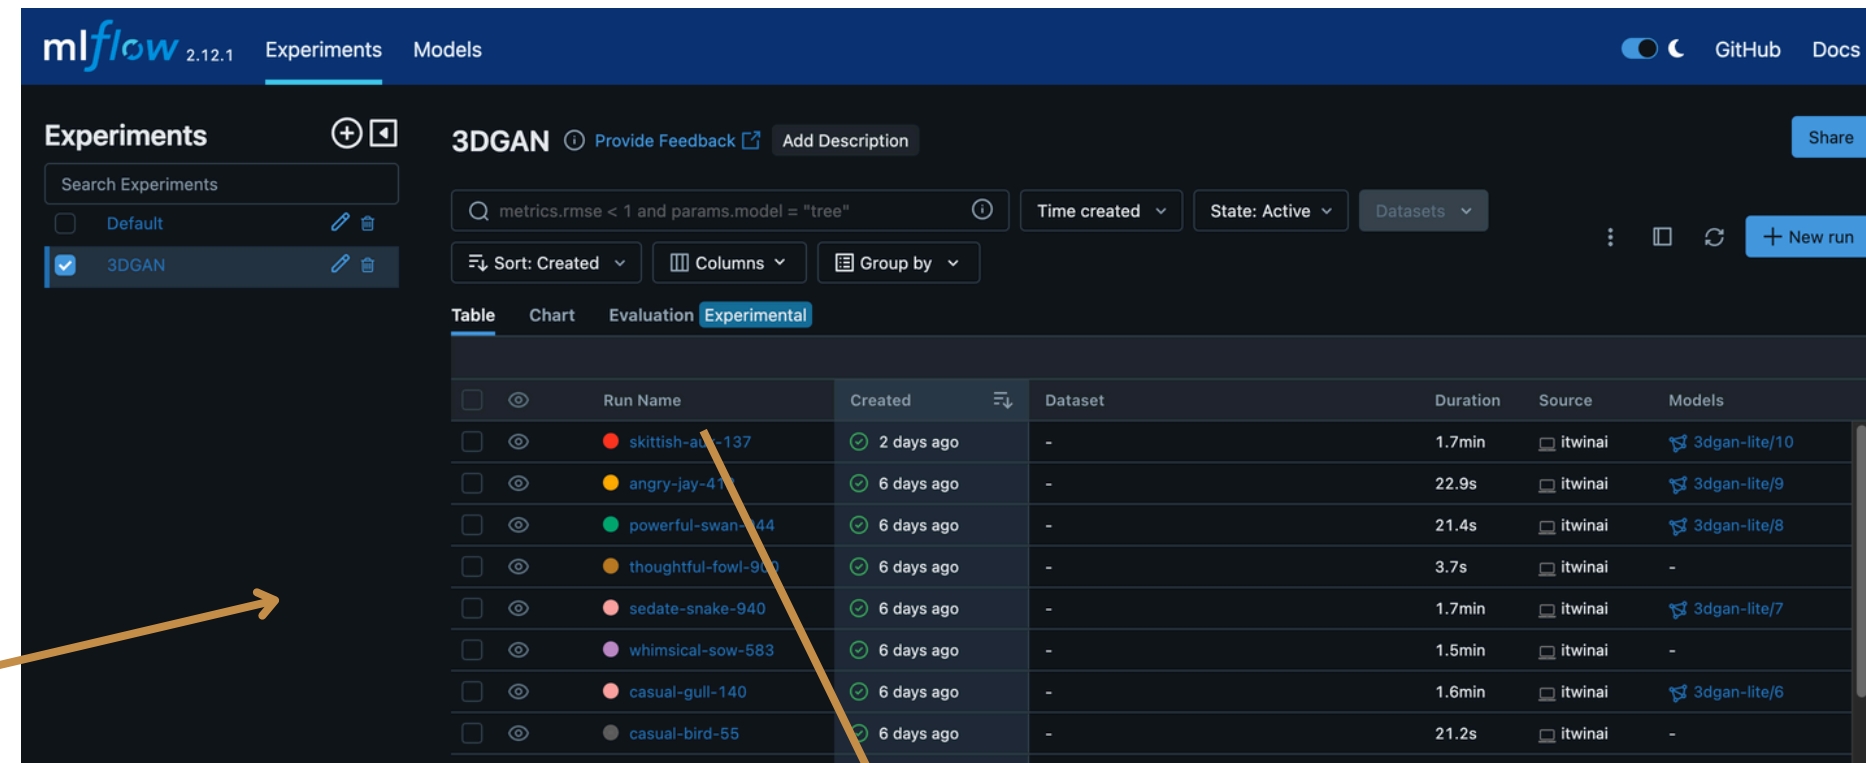


# "Sending" containers via any cloud-friendly framework might be also available



## ML case study

Use case driven by 3D GAN algorithm developed at CERN and integrated with MLFlow tracking.



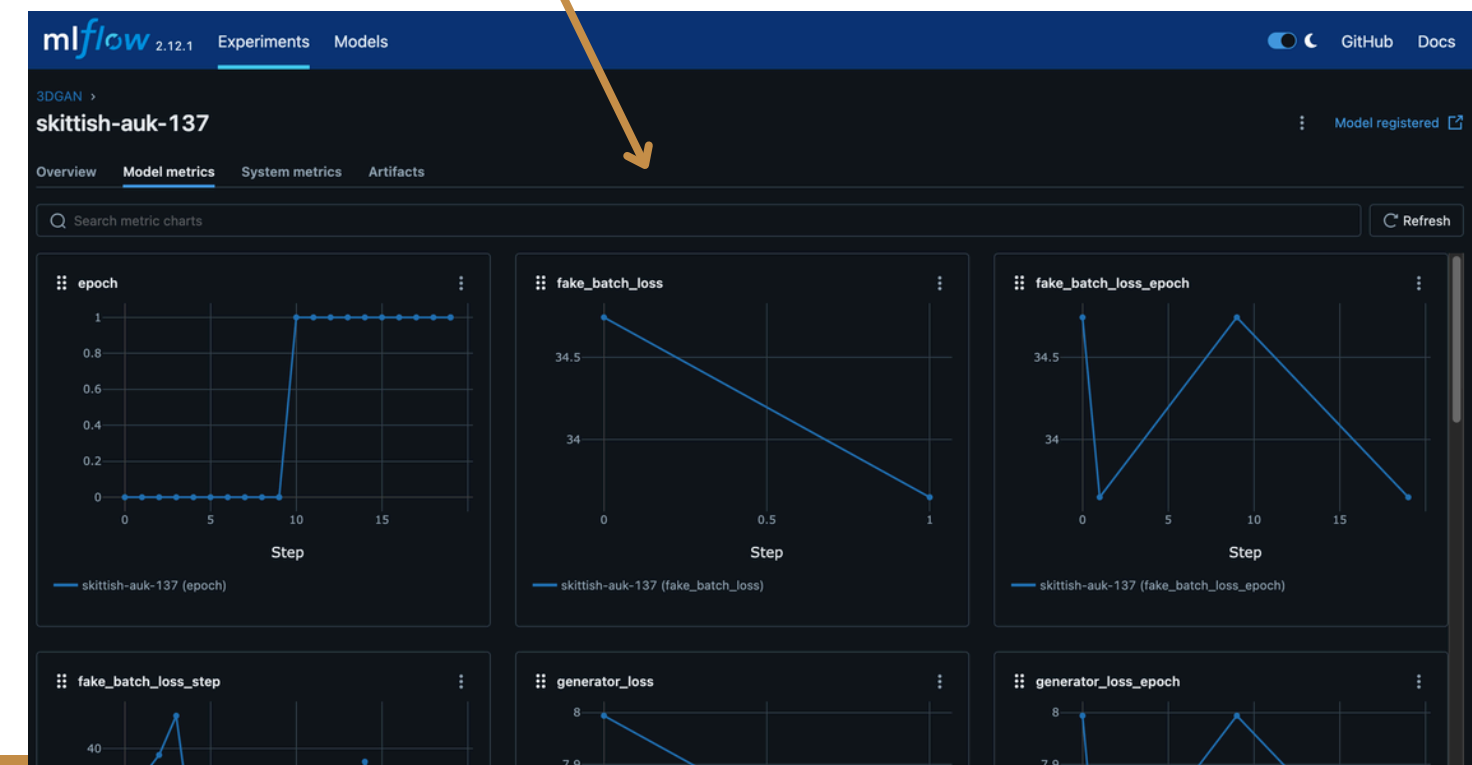
```

[memon1@hdfml01 ~]$ squeue --me
JOBID PARTITION NAME USER ST TIME NODES NODELIST(REASON)
60936 batch Prototyp memon1 CF 1:28 1 hdfmlc03

[ciangottinid@intertwin ~]$ squeue --me -i5
Tue Sep 26 09:41:12 2023
JOBID PARTITION NAME USER ST TIME NODES NODELIST(REASON)
4086045 cpu test-pod ciangott R 1-00:41:14 1 cn0130

#####
# Starting execution of 'Lightning3DGANPredictor'... #
#####
Searching in : /tmp/data/exp_data/*/*.h5
Found 1 files.

Predicting: | | 0/? [00:00<?, ?it/s]
Predicting: 0%| | 0/2 [00:00<?, ?it/s]
Predicting DataLoader 0: 0%| | 0/2 [00:00<?, ?it/s]
Predicting DataLoader 0: 50%| | 1/2 [00:02<00:02, 0.34it/s]
Predicting DataLoader 0: 100%| | 2/2 [00:03<00:00, 0.58it/s]
Predicting DataLoader 0: 100%| | 2/2 [00:04<00:00, 0.49it/s]
    
```



The background features a deep blue gradient. On the left side, there is a complex pattern of light trails and dots. These elements appear to be part of a digital or data visualization, with lines and points converging towards a central point on the left, creating a sense of depth and movement. The dots are small and bright, while the lines are thin and slightly blurred, giving the impression of light or data being processed or transmitted.

**Wrap Up**

## WP5 liaisons with infrastructure and middleware is a key here: we can and should act as the contact point

### Sneak peaking/early access

- As already mentioned in this presentation, being aware of what is going on, and getting prepared for it
- Take advantages of the point above to get access to any kind of experimental service --> access to resources
- Help the design/development of the solutions --> include Spoke3 critical features from the beginning

**I think we should get the best out of any upcoming Proof Of Concept instantiation.**

WP5 should act as a two-ways bridge:

- Bringing well defined use cases as a validation tool for the infrastructure PoC
- Getting access to early development decisions

