



Diego Ciangottini – INFN Sezione Perugia
Federica Cuna – INFN Sezione Bari
Maria Bossa – INFN Sezione Napoli
Dario Gasparrini – INFN Sezione Roma 2
Matteo Pasqui – INFN Sezione Roma 2

RaySTER

Ray-powered Scalable Training and Execution Runtime

1. RaySTER is an initiative to deploy and **manage distributed Ray clusters on HPC infrastructure** through seamless Kubernetes-to-SLURM integration using the interLink Virtual Kubelet provider.
2. The **target is using Terabit resources to demonstrate a sustainable model** to be replicated also on other scenarios with similar topologies.
3. The project demonstrates how **Ray's cloud-native distributed computing framework can be transparently offloaded to traditional HPC batch systems**, bridging the gap between modern AI/ML development workflows and high-performance computing resources.

Scientific Rationale: AI tracking algorithm

To test our algorithm on **more complex data**, we decided to analyze HERD simulated data. The HERD (High Energy cosmic-Radiation Detection) experiment is a space mission designed to directly detect cosmic rays, and it is set to be installed on the **Chinese Space Station (CSS)** in 2028. The main goals of the mission:


- enhance our understanding of high-energy cosmic rays,
- search for indirect signals of dark matter,
- probe sources of high-energy particles such as protons, electrons, and photons.

Fiber Tracker (FIT):

- Surrounds CALO on top and sides.
- Particle tracking and charge measurements
- 5 sectors, each with 7 X-Y scintillating fiber layers.
- Provides 7 precise position measurements.

The HERD simulation dataset
The full dataset, generated using the custom HerdSoftware simulation framework, consists of 4,300,000 events, equally divided into 2,150,000 electron events and 2,150,000 proton events. Both event sets are simulated within a power-law energy spectrum E^{-1} , spanning an energy range from 100 GeV to 1 TeV, and are distributed within a spherical region surrounding the HERD detector.

Electrons set was used for the tracking algorithm



ICSC Spoke3 use case

A GNN-driven tracking to enhance tracking performances on tracking purposes over the analytical/traditional approach.

The GNN algorithm

- SageConv architecture
- 18 layers
- Mean aggregation function
- 128 hidden size
- Adam optimizer
- Binary cross entropy loss function

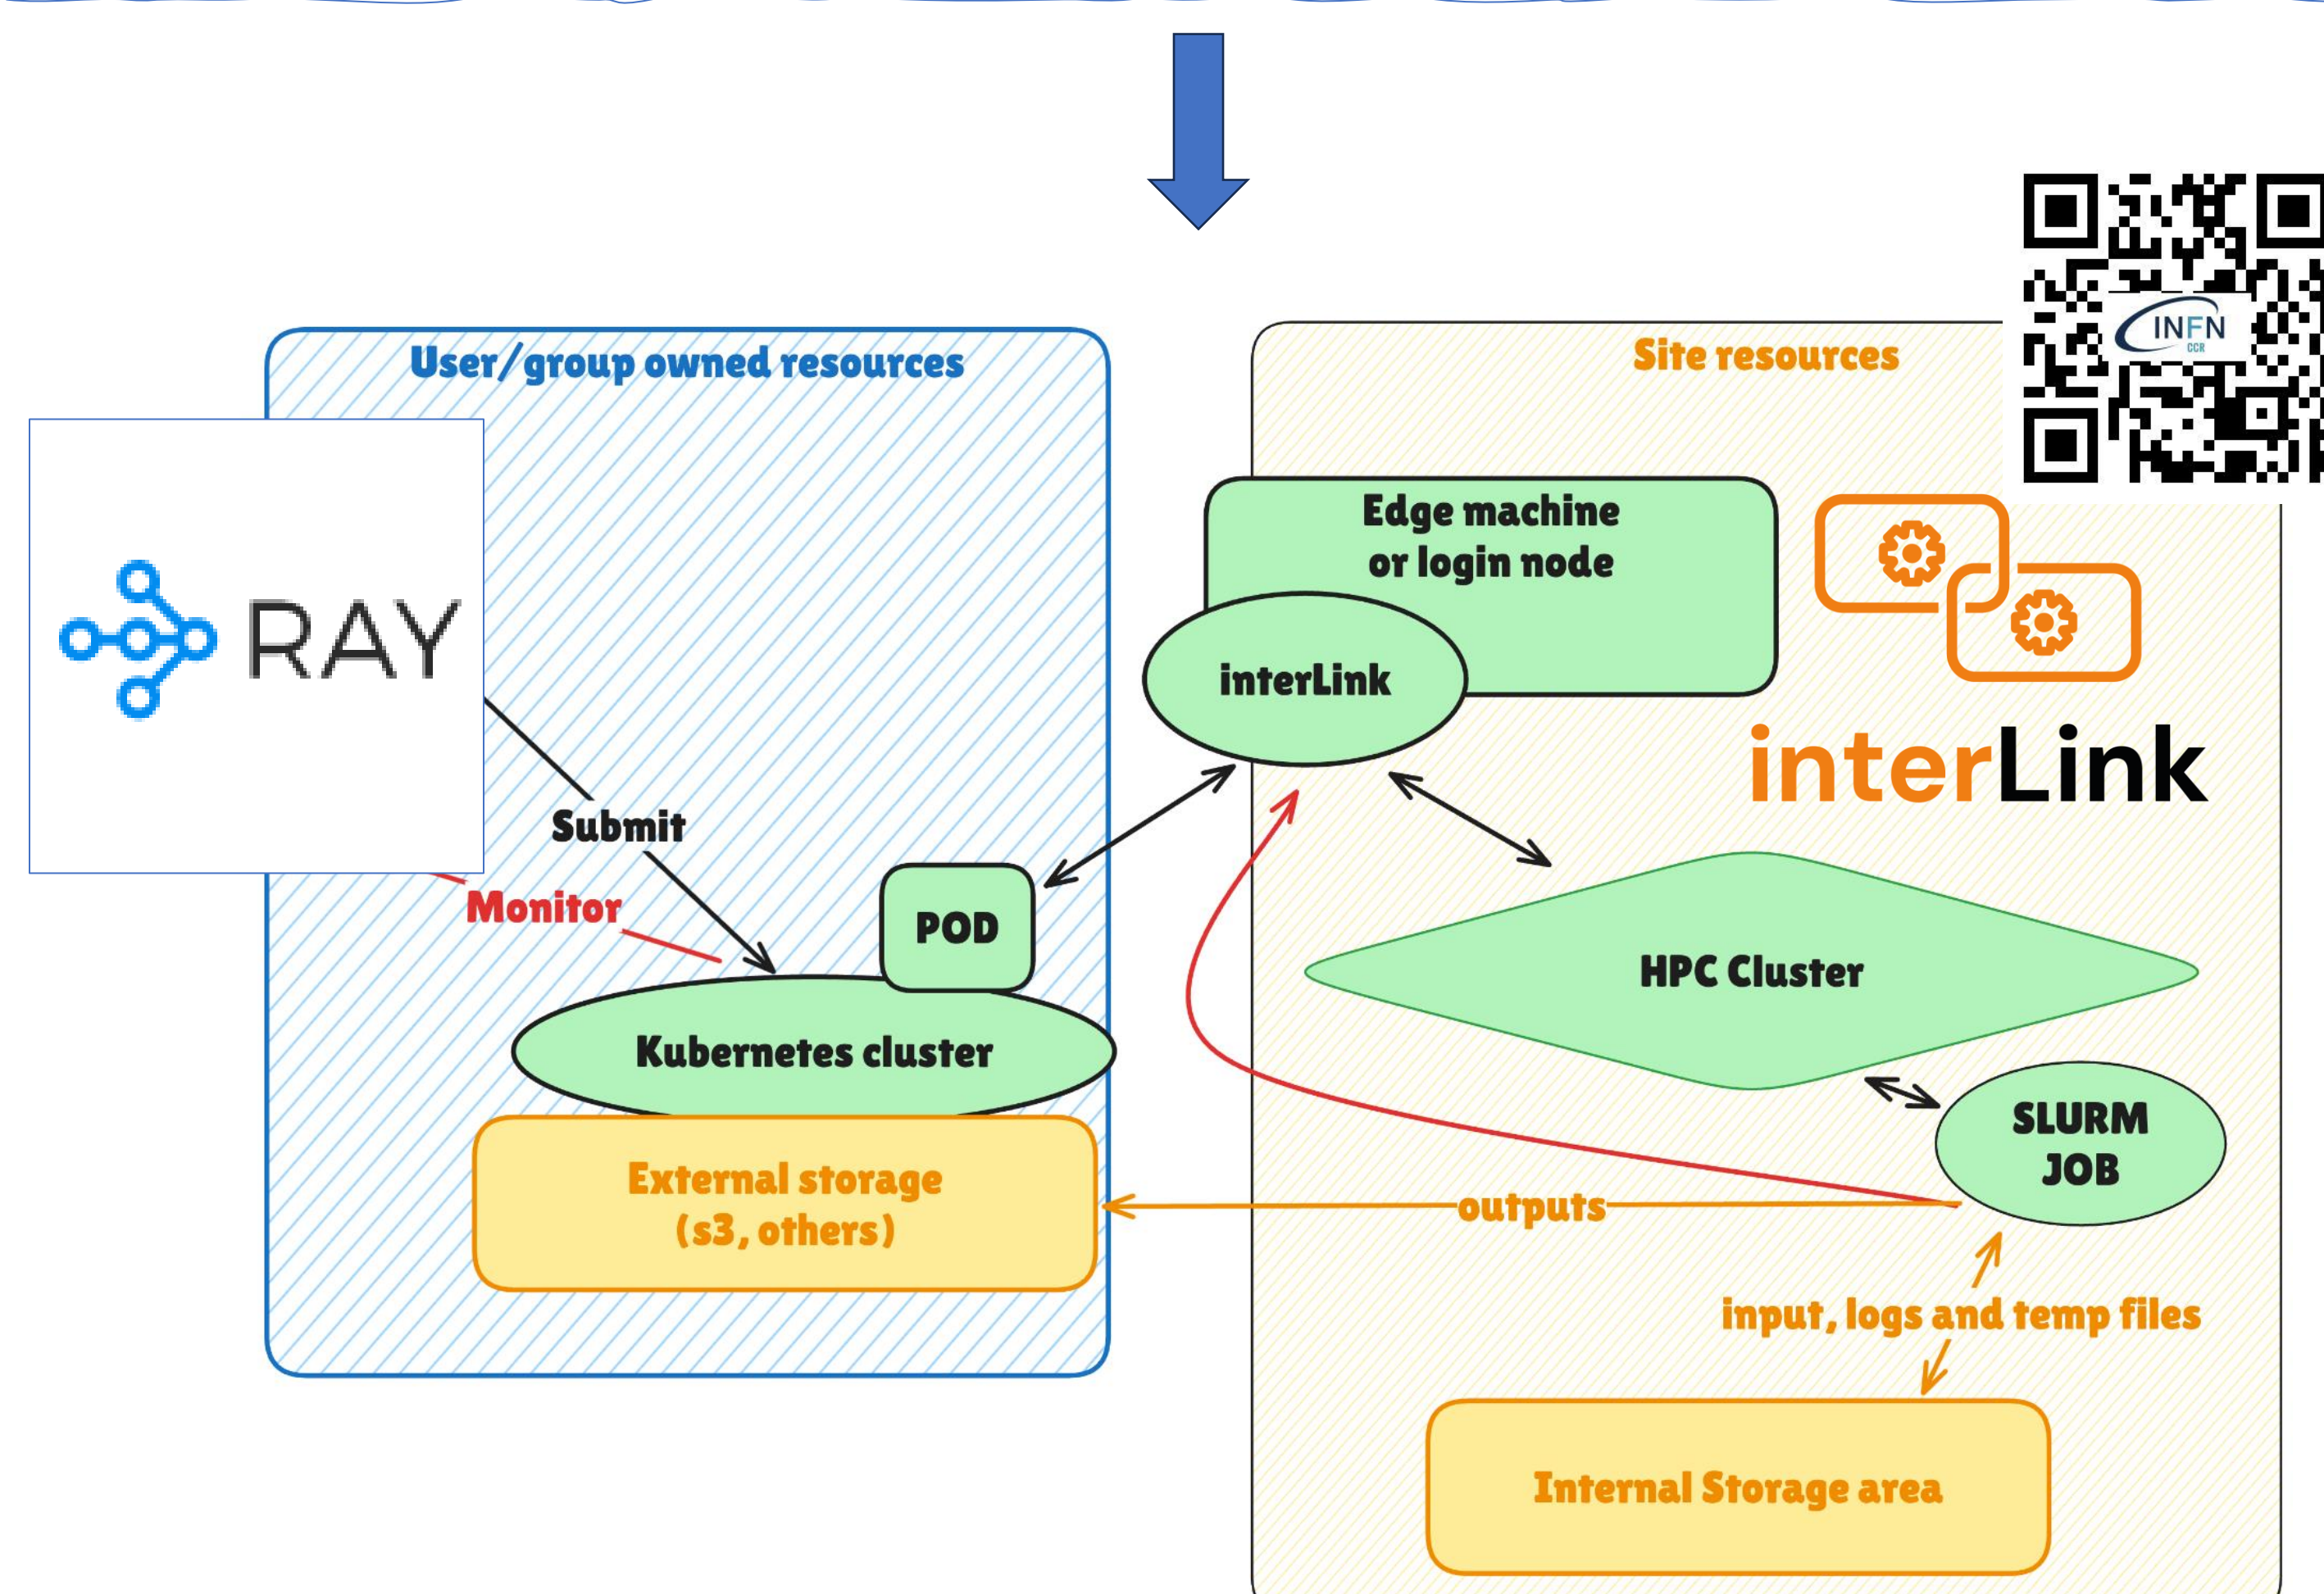
4,5 million simulated track data
75% train-15% validation-10% test

Time and resource consuming

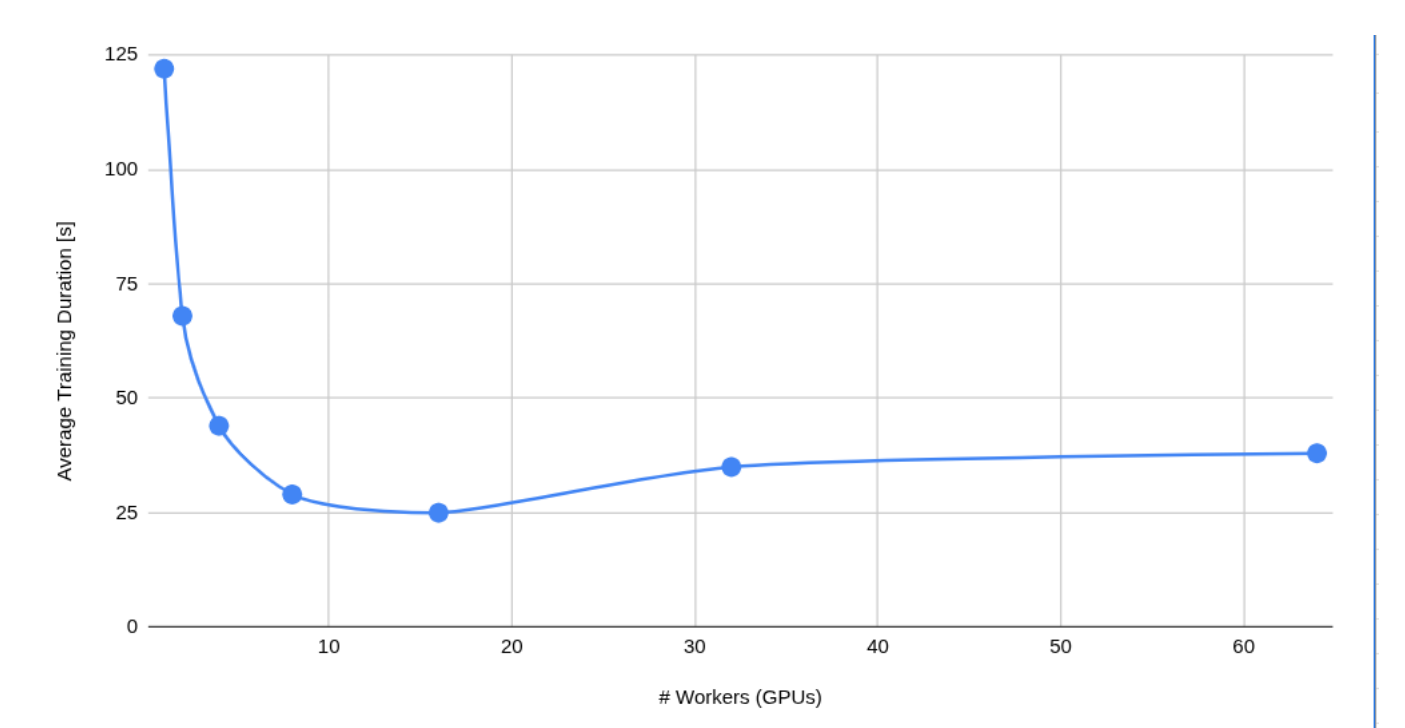


A distributed training approach over multiple GPUs (and nodes) has been implemented over a widely adopted cloud-native framework: KubeRAY clusters

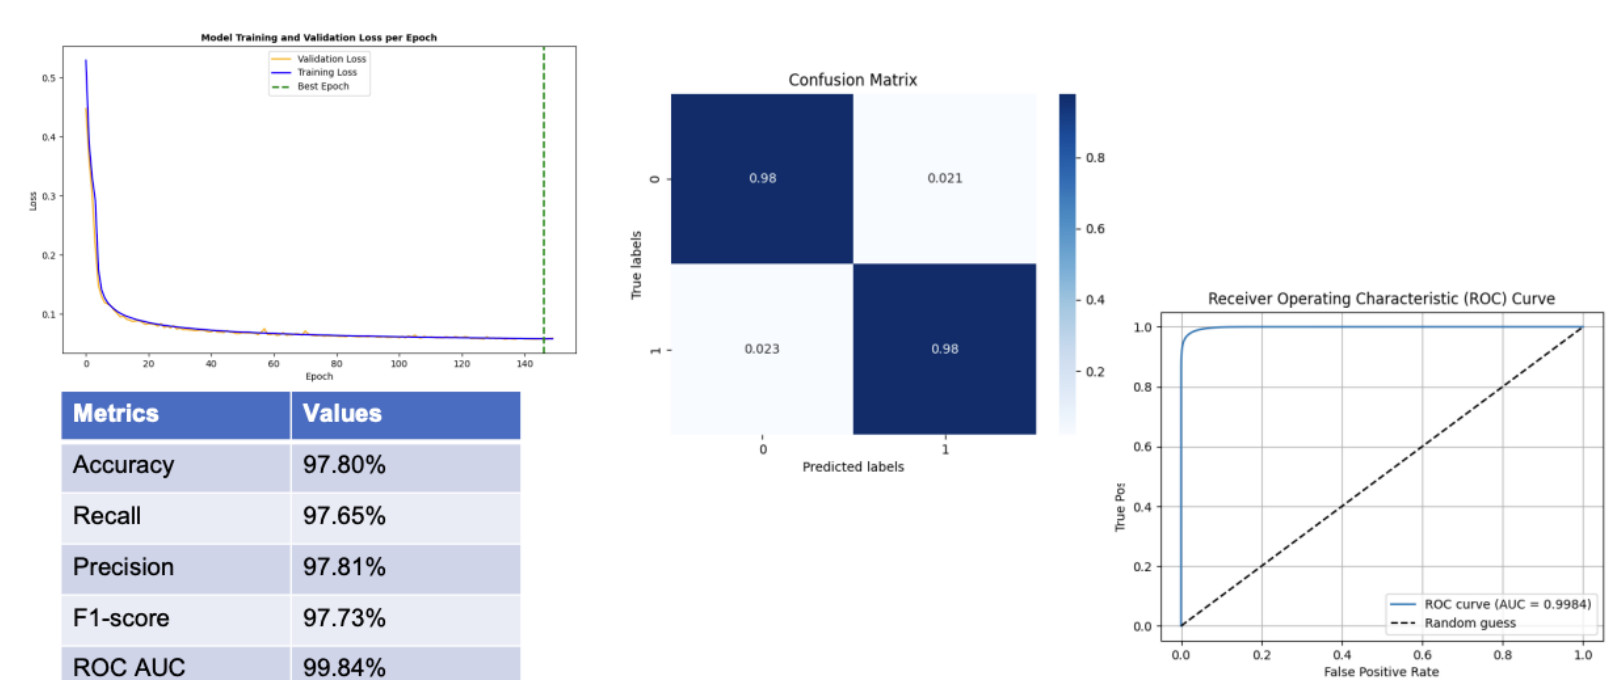
Through TERABIT resources we deployed a model where we can offload such trainings from a K8s cluster hosted on cloud resources to HPC SLURM queue in Padova, without any change to the framework



Scalability validation
We validate the scalability of the offloaded cluster when increasing its size



Scalability validation
We validate the scalability of the offloaded cluster when increasing its size



Conclusions

We demonstrated the viability of hybrid cloud-HPC architectures for high-energy astrophysics detector simulation workloads (developed in the context of the ICSC Spoke3 activity), enabling researchers to exploit both the flexibility of Ray's unified distributed computing platform and the specialized hardware capabilities of HPC systems without architectural compromises or workflow modifications.

Acknowledgements

This work was possible thanks to the invaluable support of the Padova HPC administrators, ICSC Spoke0 and to the use case effort in the context of "Fondazione ICSC, Spoke 3 Astrophysics and Cosmos Observations."

