The Platform-as-a-Service paradigm meets ATLAS: developing an automated analysis workflow on the newly established INFN Cloud

International Conference on Computing in High Energy and Nuclear Physics

May 8 – 12, 2023 - Norfolk Waterside Marriott

Caterina Marcon, Leonardo Carminati, David Rebatto

On behalf of the ATLAS Computing Activity

INFN MILANO
Istituto Nazionale di Fisica Nucleare
Sezione di Milano

ATLAS EXPERIMENT

Motivation

- ATLAS has been using a complex and distributed computing infrastructure: the **Worldwide LHC Computing Grid (WLCG)** characterized by almost a million computing cores and an exabyte of storage deployed in different sites worldwide;

- The computing needs (power and storage) of ATLAS in the **HL-LHC** era will represent an **unprecedented challenge** for the existing infrastructure:

  - New software and hardware technologies are being explored;
  - The experiment is considering integrating various alternative computing resources into the distributed computing system, including **cloud computing technologies**.

- Cloud technology allows **dynamic, flexible and cost-effective resource provisioning**.

INFN Cloud infrastructure

- INFN CLOUD infrastructure in production since March 2021;
- Backbone connecting the large data centers of CNAF and Bari;
- Smaller federated sites offer opportunistic resources;
- Resources orchestrated by OpenStack;
- Active INFN users can access all the federated resources;
- Appointed “administrators” can provide sub-services;
- Two operation models:
  - Platform-as-a-Service (PaaS)
  - Software-as-a-Service (SaaS)

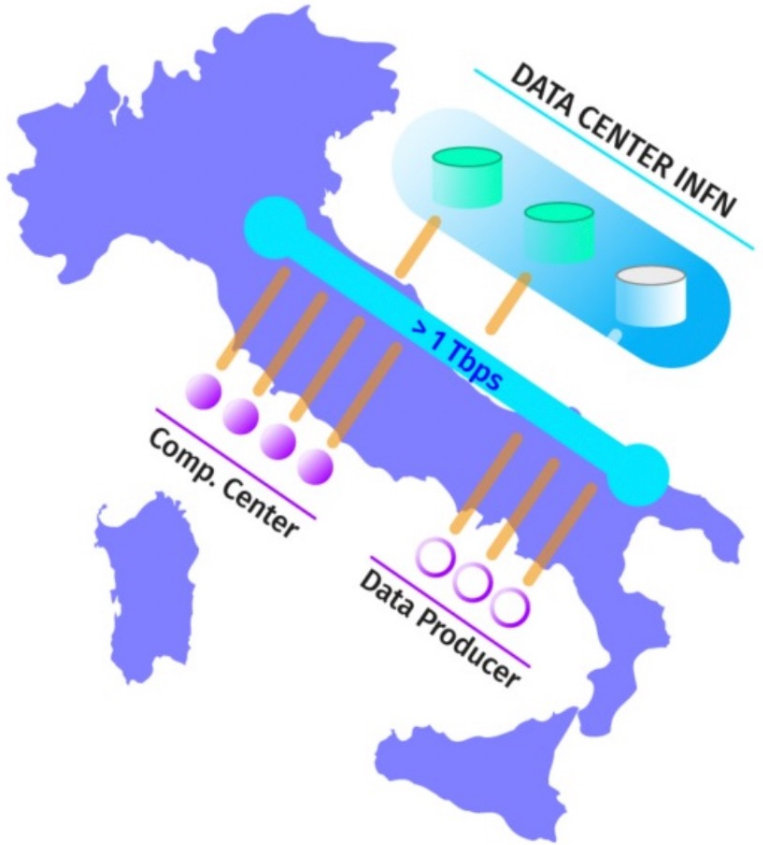

Objectives and Tools

- Investigate the possibility of implementing two distinct (but not orthogonal) analysis workflows by exploiting the computational resources of **INFN Cloud**:
  - create a batch-like system capable of obtaining flat n-tuples (compatible with analysis flows for result extraction) from structured and complex data;
  - develop **interactive analysis flows** (similar to Jupyter Notebook-as-a-Service).

- High level building blocks:
  - Different **Docker images** to create an alternate **ATLAS software stack provisioning** architecture;
  - Using **Kubernetes** for **resource orchestration**;
  - Using **HTCondor** as the **job scheduling system**.

The Platform-as-a-Service paradigm

Users maintain Applications and Data

+

Providers maintain the underlying infrastructure

=

PaaS paradigm

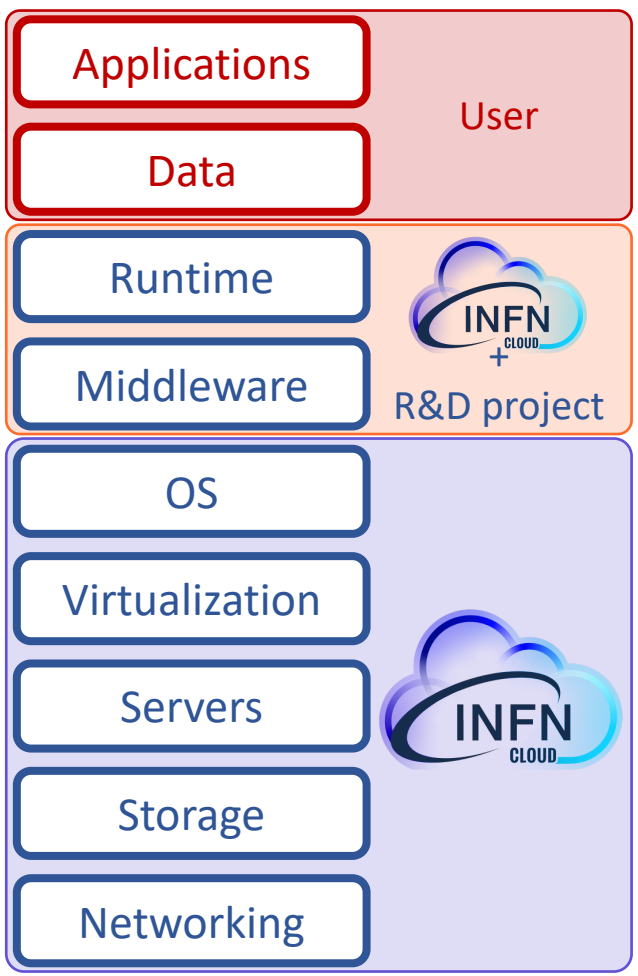

CERN - ATLAS middleware and runtime added on top of existing PaaS services from INFN Cloud

The R&D resource pool

- The following resource pool has been provisioned:

CPU	92
RAM [GB]	168
Volumes [GB]	1000
External storage (compatible with S3) [GB]	2048

- A pre-defined set of cloud applications is available:
  - pure Kubernetes clusters;
  - HTCondor clusters deployed on Kubernetes;
  - General purpose Virtual Machines (with Ubuntu 18.04, Ubuntu 20.04 or CentOS 7);
  - S3 storage.

- The scale of these applications is configurable and resources are drawn from the reserved pool.

HTCondor on Kubernetes (I)

- **Kubernetes** (K8s) cluster with **HTCondor batch system** on top created via INFN Cloud Dashboard;
- Resources drawn from R&D pool and orchestrated by OpenStack
- Basic monitoring services configured by default (e.g. Grafana dashboard with Prometheus);
- Limited user configurability:
  - number of worker nodes;
  - Docker image of the worker nodes;
  - master and worker node VM size (RAM and CPU).

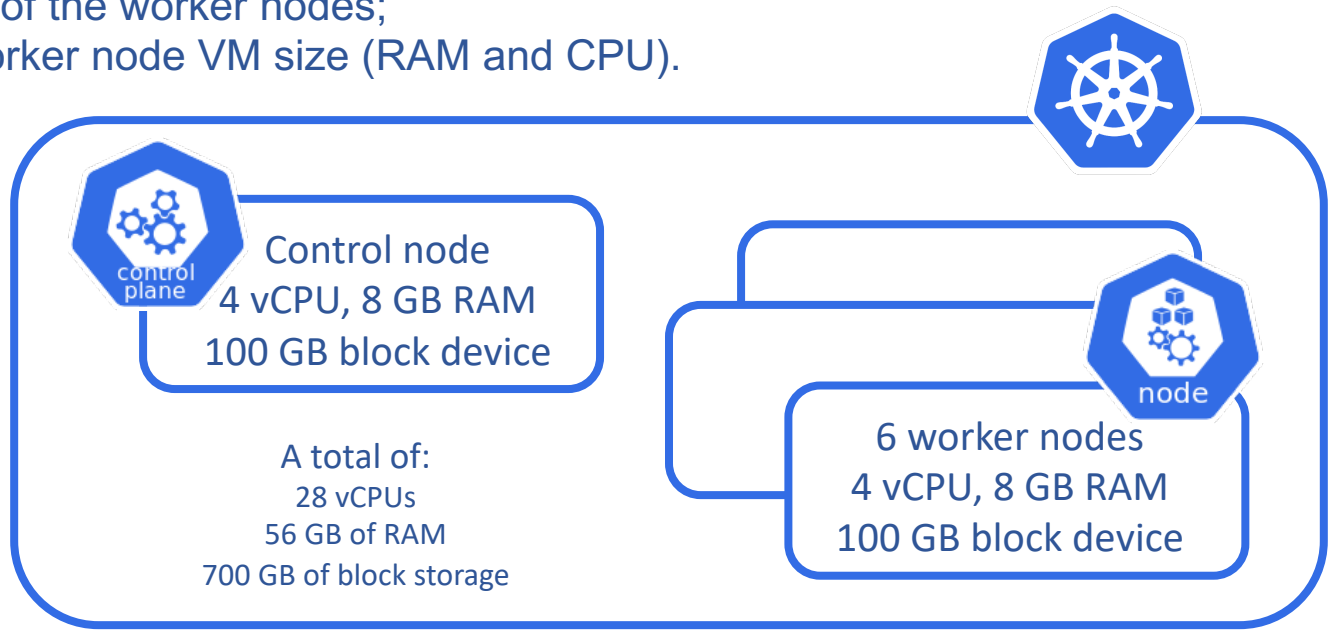

HTCondor on Kubernetes (II)

- HTCondor components configured as **K8s deployments**;
- Deployments can be easily **scaled** by the cluster administrator;
- **No HTCondor submit node** on cluster by design to allow remote job submission.

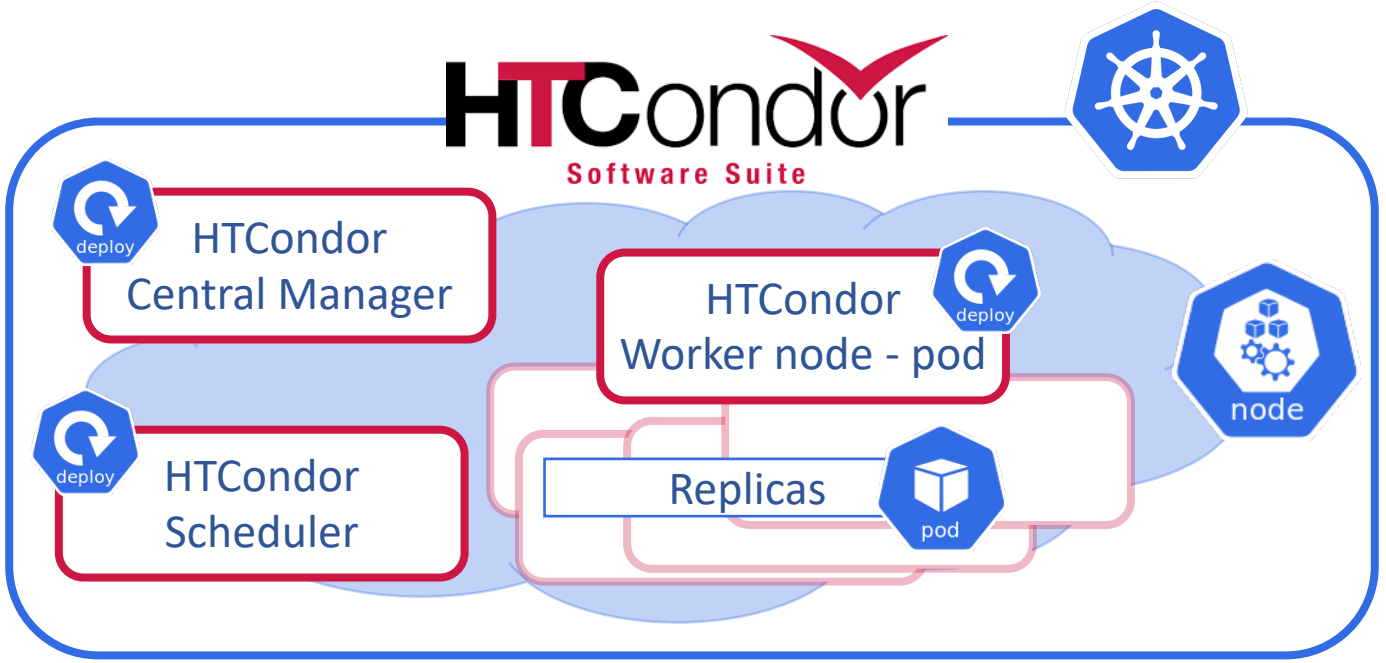

HTCondor on Kubernetes (III)

- Submit node designed as a **satellite Docker container**;
- Jobs can be submitted to the cluster from any **remote location**;
- Authentication to the HTCondor cluster via the **INFN IAM infrastructure**.

Merging CERN and INFN resources (I)

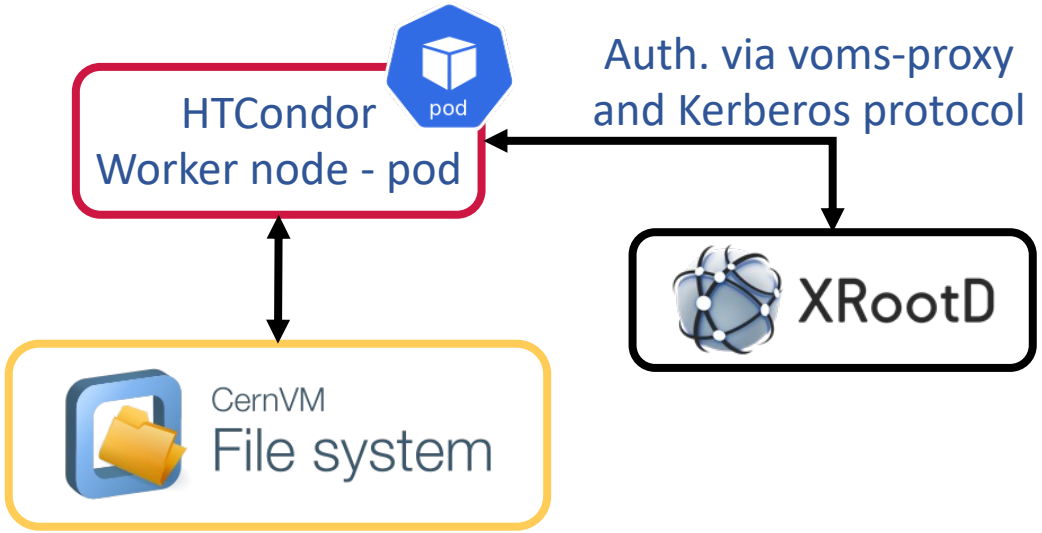

- ATLAS resources must be linked:
  - CVMFS to retrieve the required software;
  - XRootD for data file transfer.
- HTCondor worker pod images updated to include CVMFS and support for X509 and Kerberos authentication.

Merging CERN and INFN resources (II)

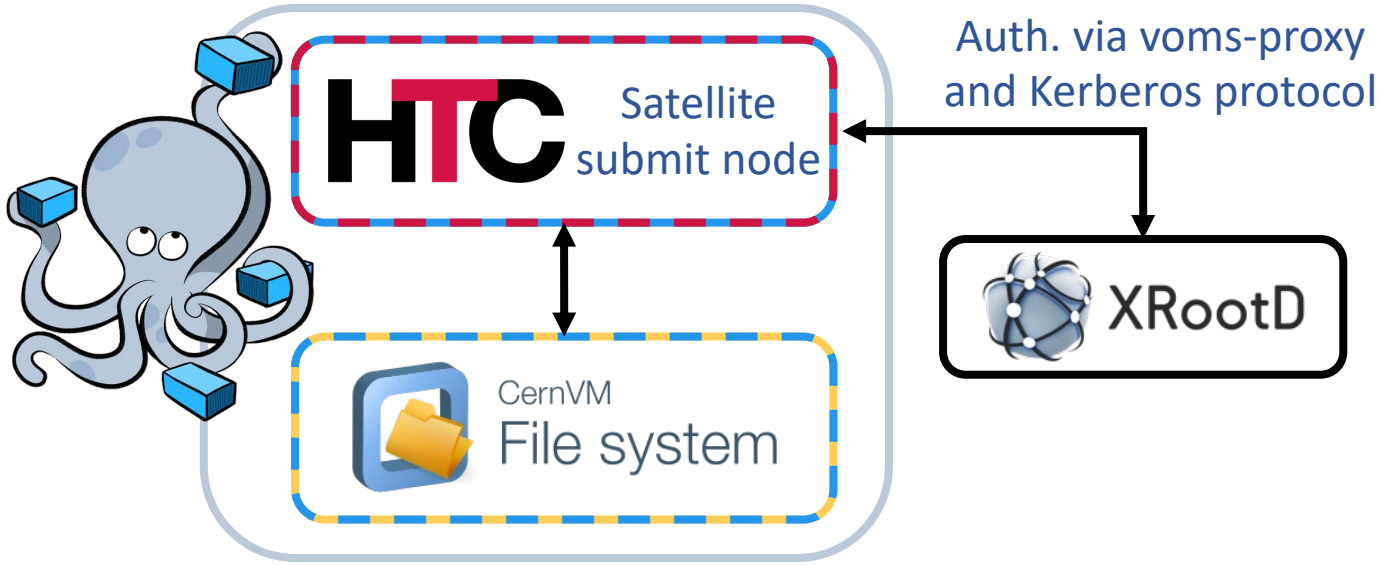

- CERN authentication for XRoot access added to the submit node;
- CVMFS running in a separate container using the official cvmfs/service:2.10.1-1 image [1], [2];
- Container integration via Docker Compose;
- Host machine and containerized ecosystem are isolated (except for the shared kernel).

[1] https://hub.docker.com/layers/cvmfs/service/2.10.1-1/images/sha256-511c85a96c50f89871dbfc1ebd9ab1d7df6b54310cc3745b9043e08fbfabea89
[2] https://cernvm.cern.ch/fs/

Merging CERN and INFN resources (III)

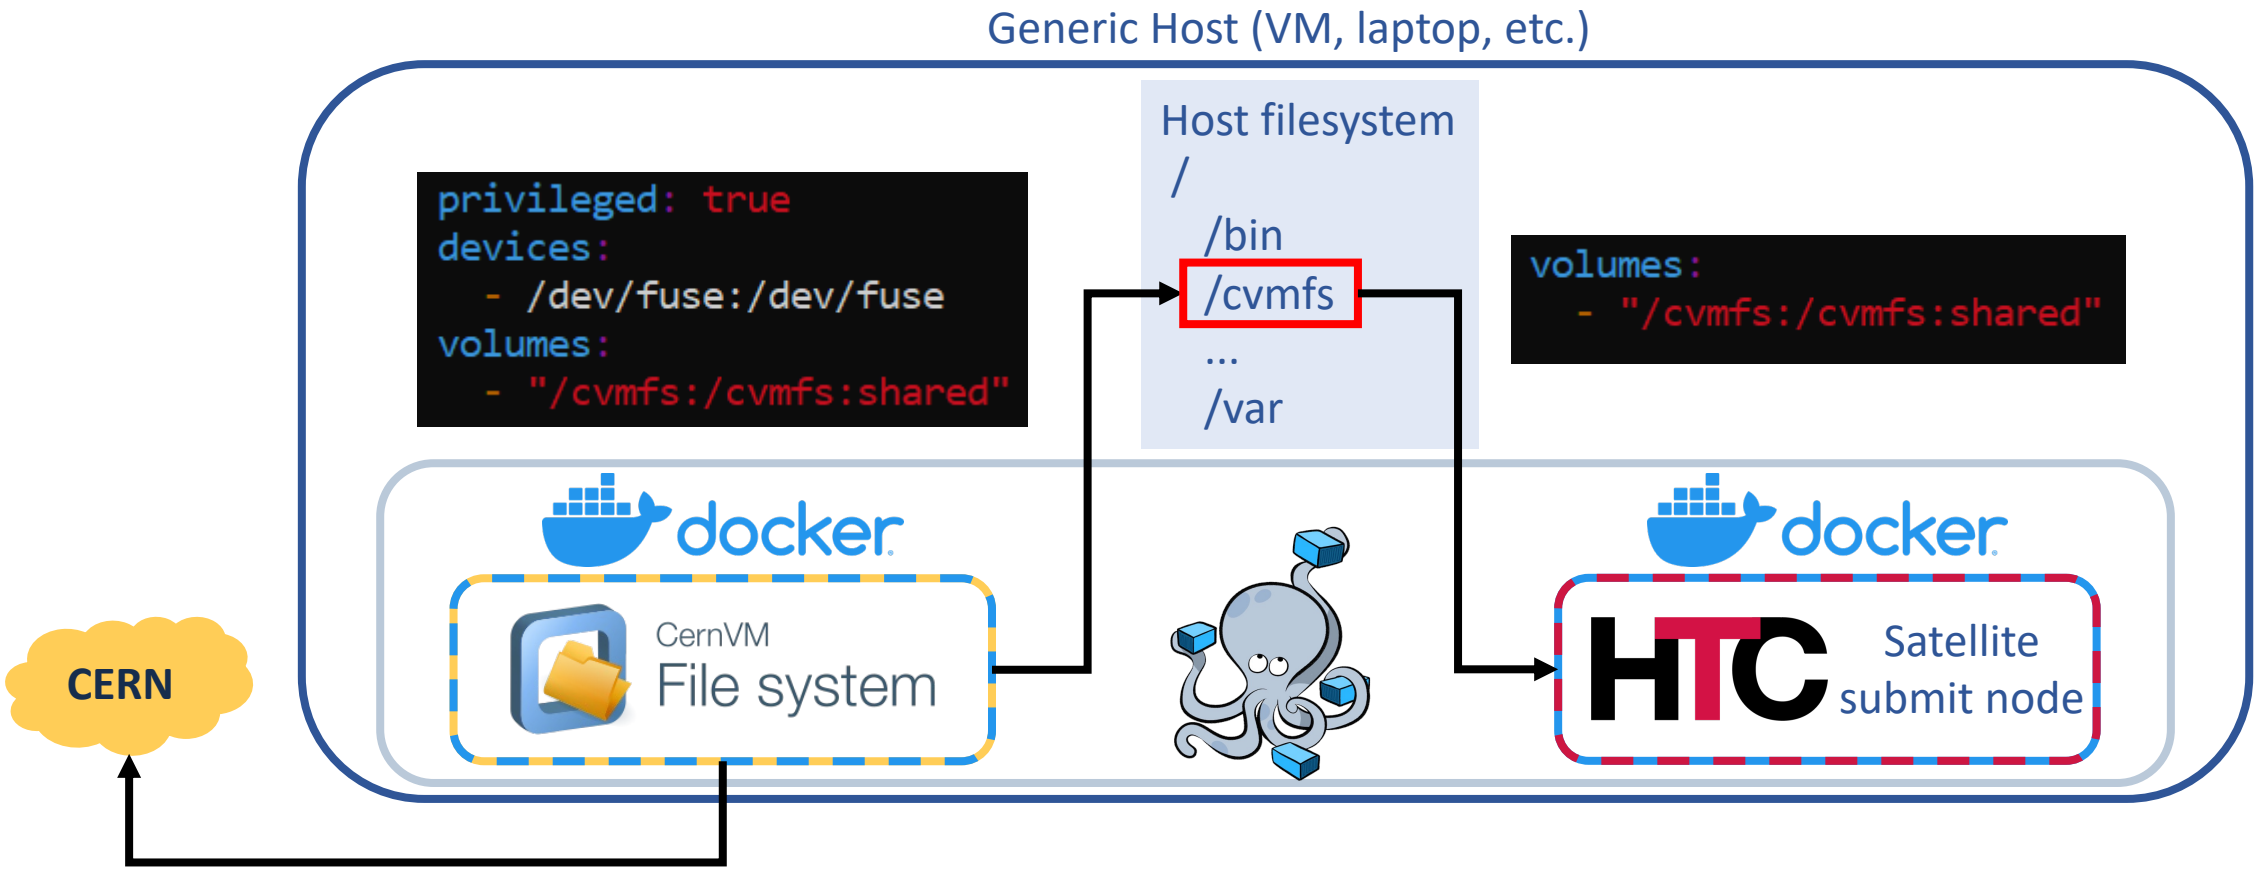

External storage integration

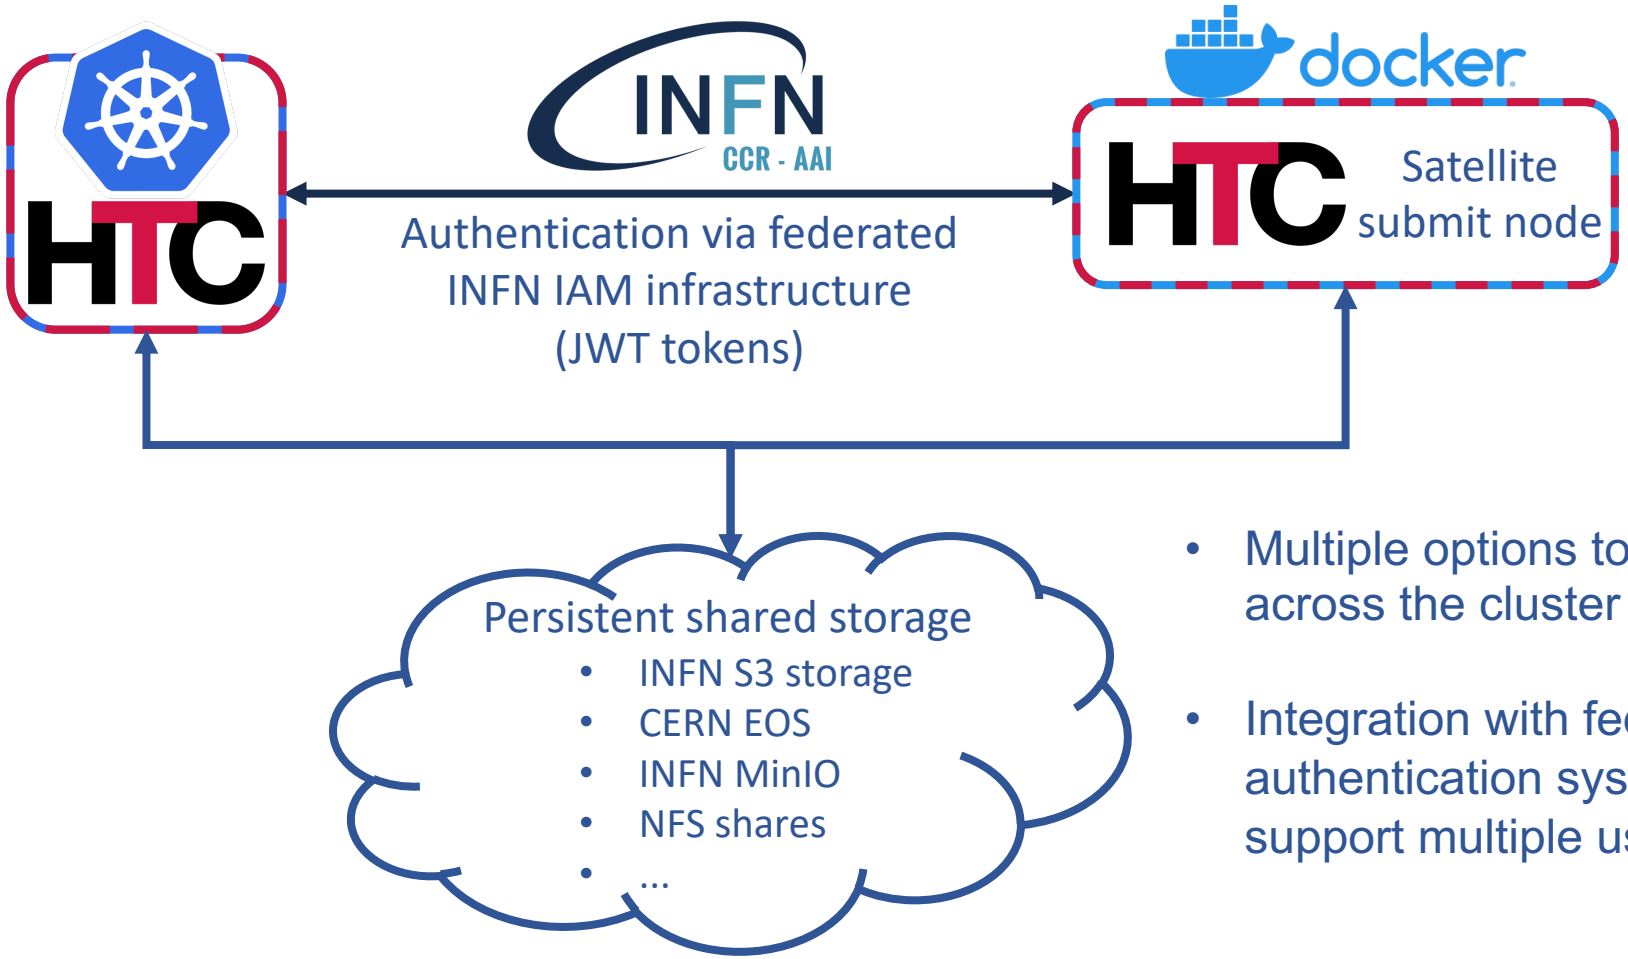

- Multiple options to share data across the cluster nodes;
- Integration with federated/SSO authentication systems crucial to support multiple users.

Future steps

- Define a strategy to **extend the availability** to a larger user base;

- Start a **thorough testing phase**:
  - Computing performance;
  - Scalability;
  - Behaviour under different type of analysis workloads;
  - Flexibility to support different workflows.

- Create a **pre-configured environment** easily deployable on INFN Cloud;

- Add **interactive workflow capabilities** to the batch-like system:
  - Batch system for data pre-processing and simplification;
  - Interactive platform for data manipulation and result extraction (plots, charts, etc.).

Conclusions

- The increasing need for computing resources foreseen for the HL-LHC era is accelerating the development of new analysis strategies based on distributed computing systems, including **cloud computing technologies**;
- In 2021, INFN joined this effort and the new **INFN CLOUD infrastructure** started operations;
- The core of the infrastructure is fully managed by INFN and users/administrators have access to **PaaS** and **SaaS** solutions, such as K8s/HTC clusters, VM or Jupyter Notebook-as-a-Service applications;
- A new R&D project has started to extend the existing PaaS paradigm with **ATLAS data analysis capabilities**;
- The objective of this project is to support **batch-like and interactive analysis workflows**, ensuring an optimized use of computing resources;
- So far, a working prototype has been produced and **testing is ongoing**.

INFN MILANO

Acknowledgments

This R&D project is funded by INFN Milano, INFN CLOUD and supported by ATLAS Italia Calcolo.