



Finanziato  
dall'Unione europea  
NextGenerationEU



Ministero  
dell'Università  
e della Ricerca



Italiadomani

PIANO NAZIONALE  
DI RIPRESA E RESILIENZA



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



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

## Enabling the compute continuum at ICSC: A solution based on interLink.

*L. Anderlini, D. Ciangottini, D. Spiga, T. Boccali, D. Gasparrini, P. Lubrano, A. Troja, F. Fanzago, M. Sgaravatto, T. Tedeschi, G. Bianchini*

**110° Congresso della Società Italiana di Fisica (SIF 2024)**

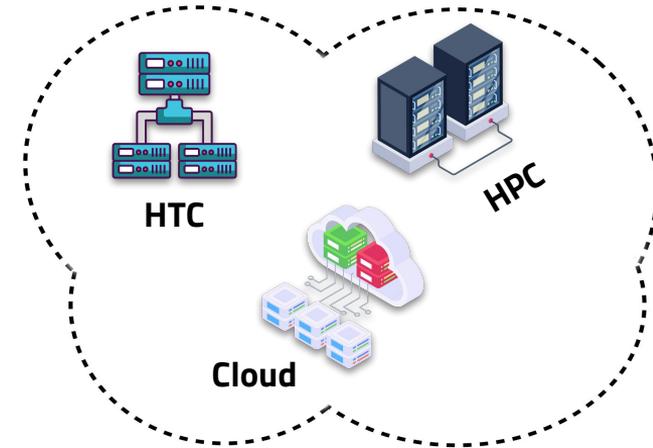
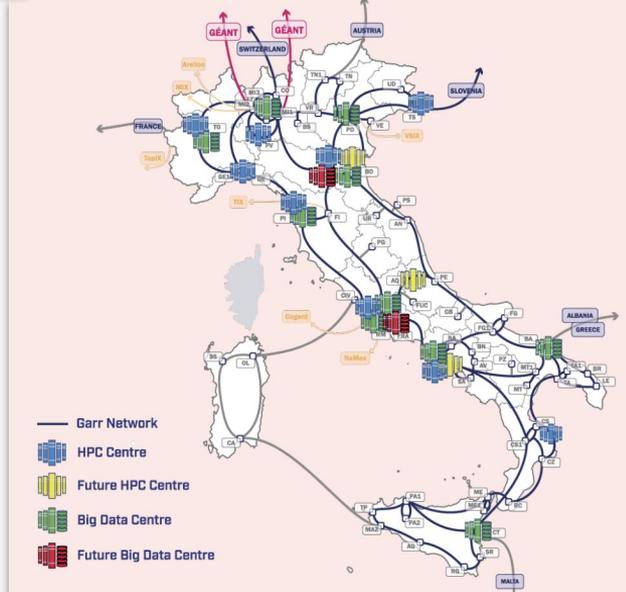
*"Emerging data and computing technologies for frontier physics" Symposium*

**September 10th, 2024**

## ICSC and Terabit: *towards a highly heterogeneous, geographically distributed computing infrastructure*



**Cutting-edge computing infrastructure connecting highly specialized centers (HPC, BigData and Cloud)** providing a comprehensive overview of the existing resources and integrating emerging technologies.



High-Performance Computing (HPC), High-Throughput Computing (HTC), and **Cloud computing** ... different resource providers, different backends, different authN/Z workflows, different network policies...



**Scope:** *Support complex workflows by distributing payloads throughout the network of computing centers, effectively mapping resource requests to compute backends while hiding the underlying heterogeneity and complexity from the user.*

## Kubernetes as the emerged standard for distributed applications (cloud-native)

As Python+NumPy emerged as *de facto* standard for Data Science applications in early 2000s, the **Kubernetes APIs have been emerging as the standard for distributed applications.**



**Kubernetes** is an open-source system for automating deployment, scaling and management of **containerized** applications on multiple nodes.

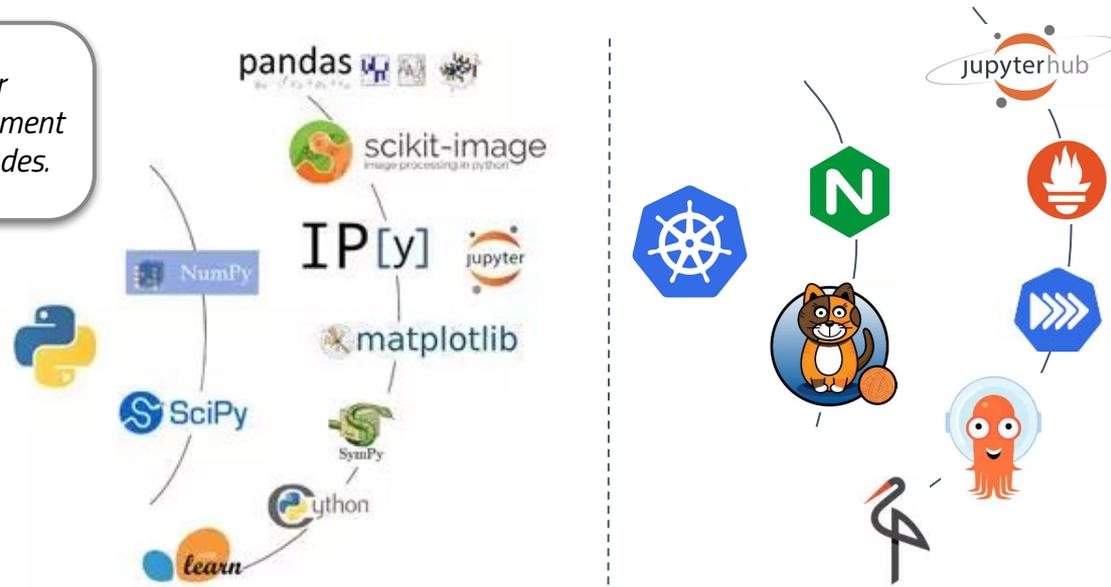
The **Cloud Native Computing Foundation** (CNCF) supports and coordinates the development of 200 k8s applications in more than 850 git repositories.

To benefit of K8s for compute-intensive applications, tasks can be organized in job queues and submitted to HPC/HTC centers.

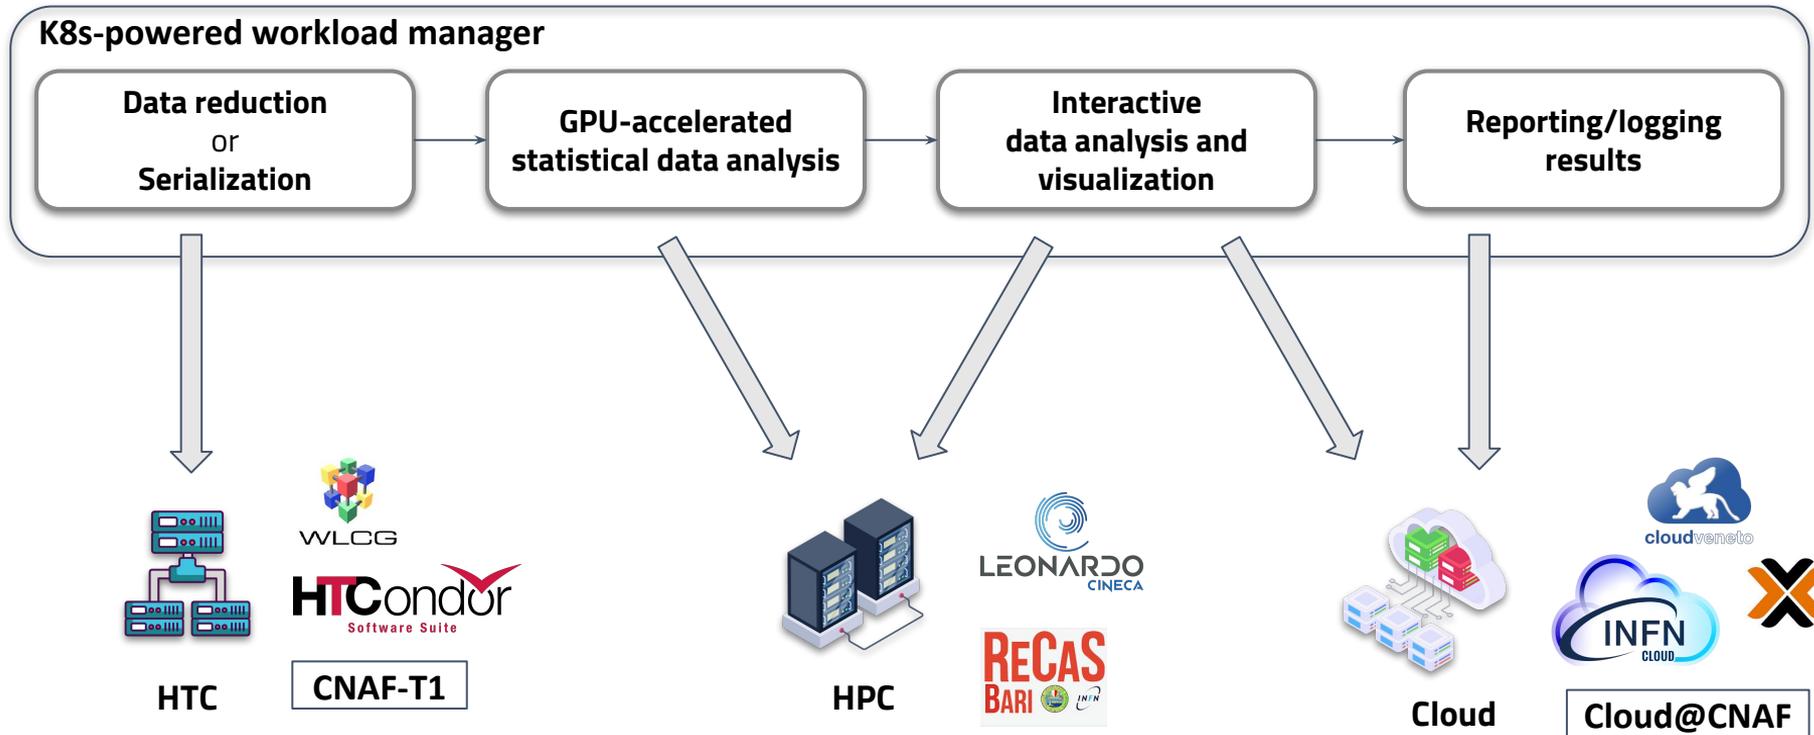
pip install



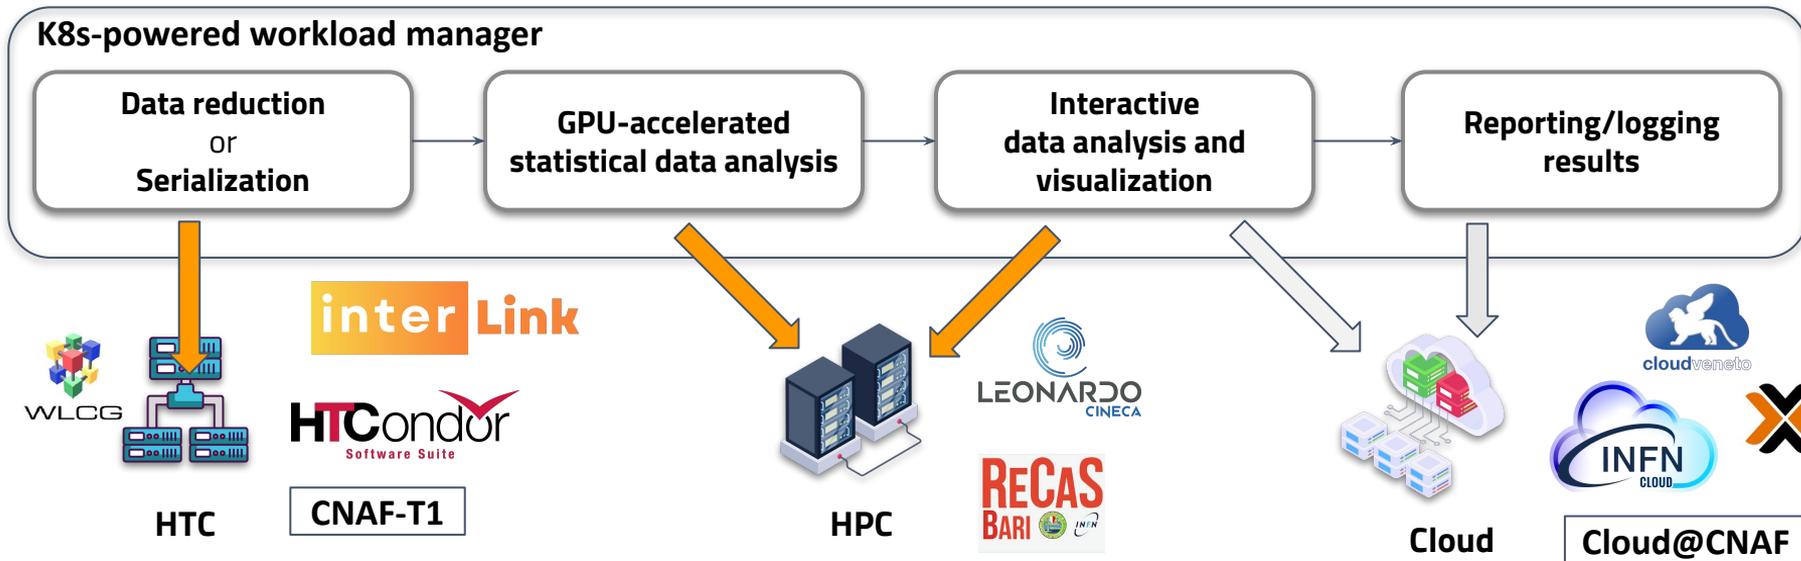
helm install



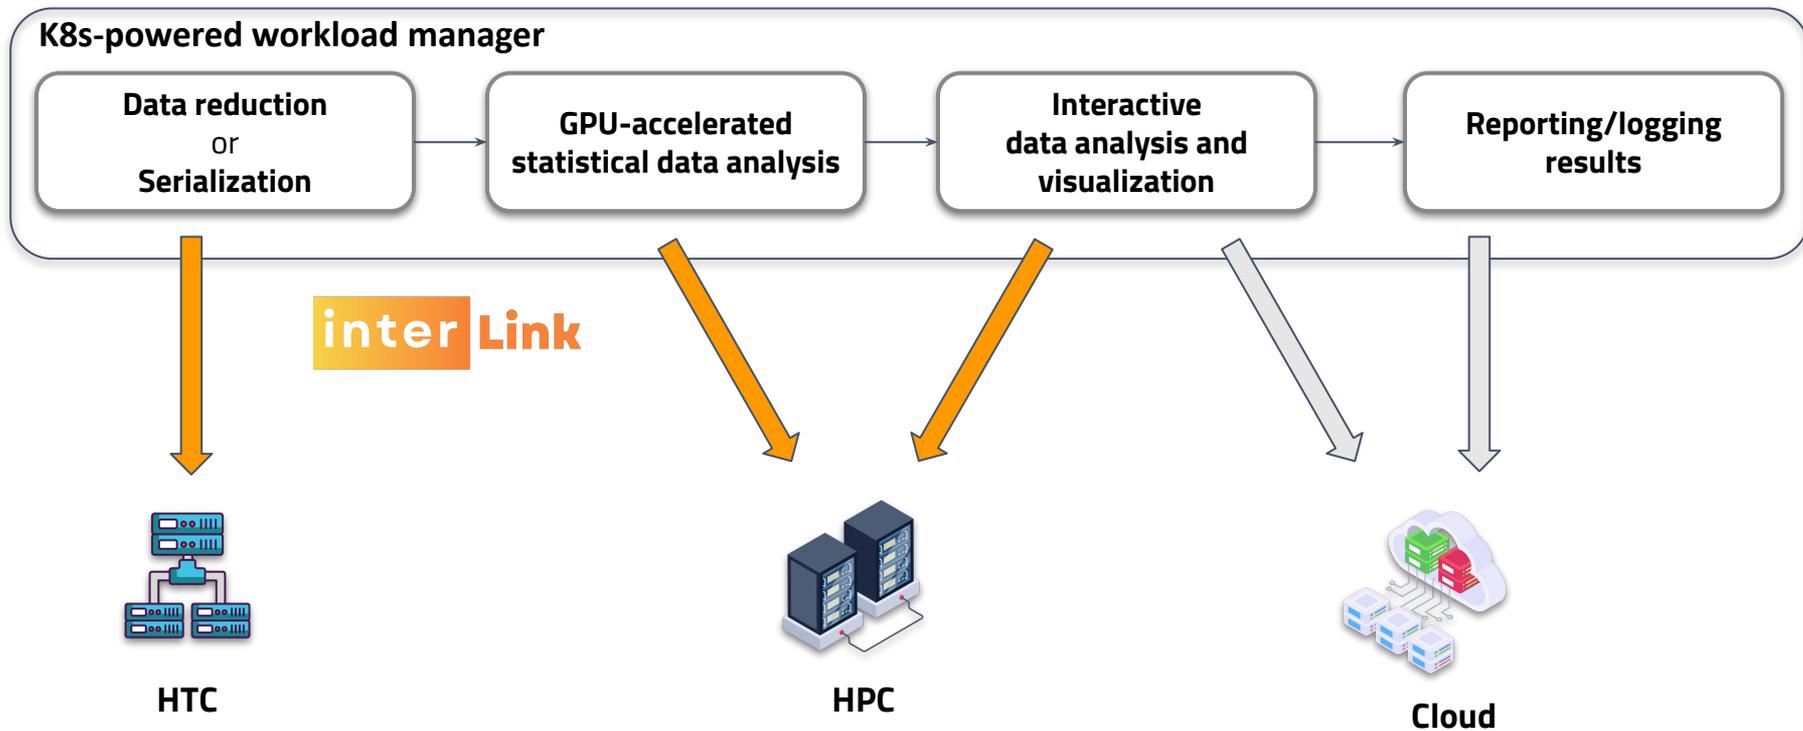
## A typical workflow, offloaded



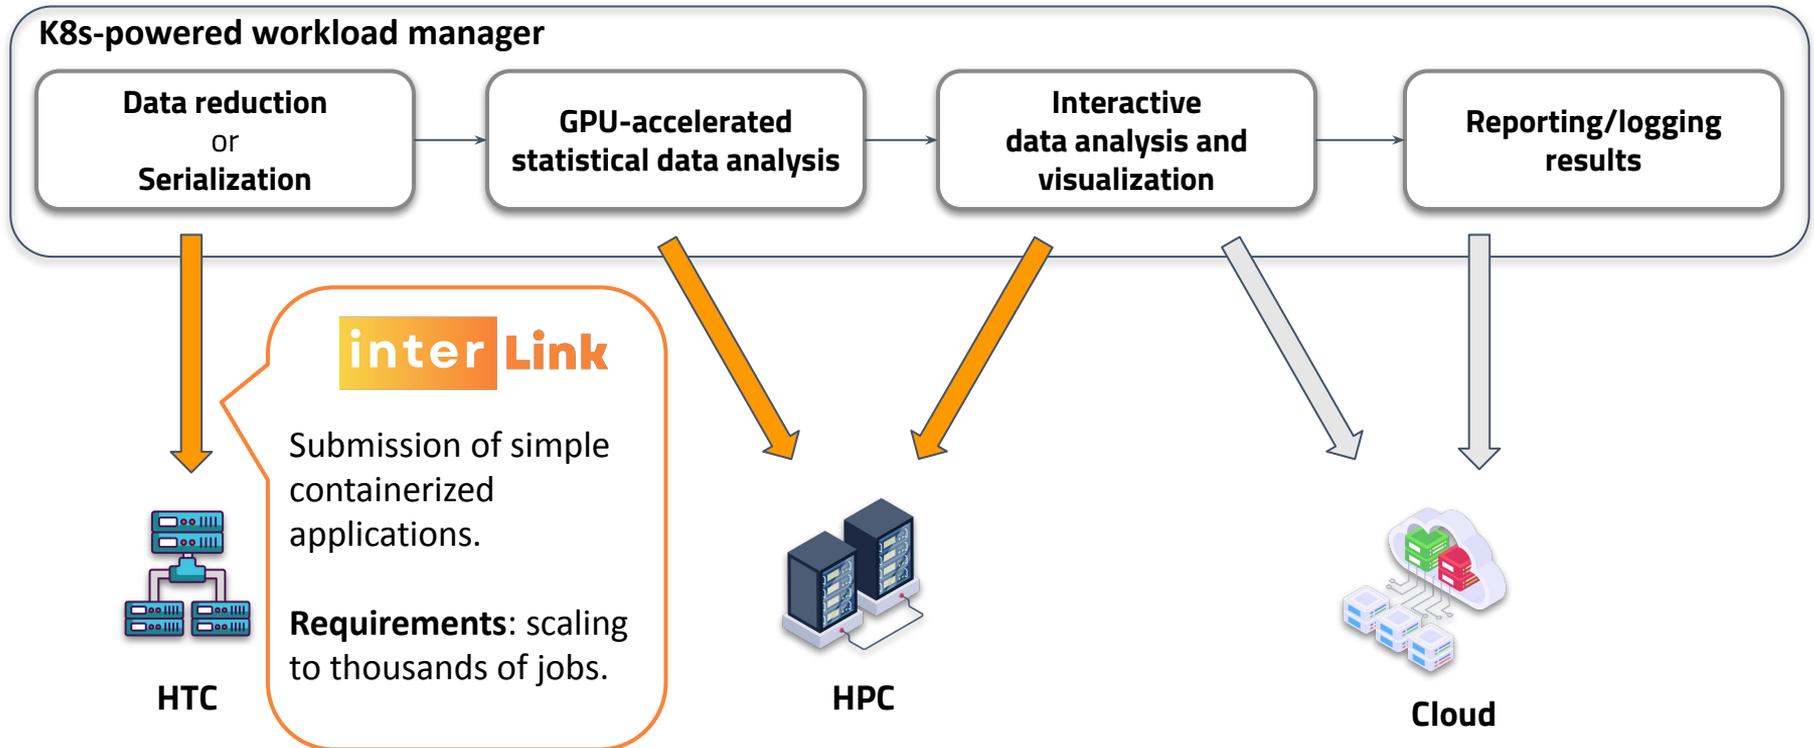
## A typical workflow, offloaded



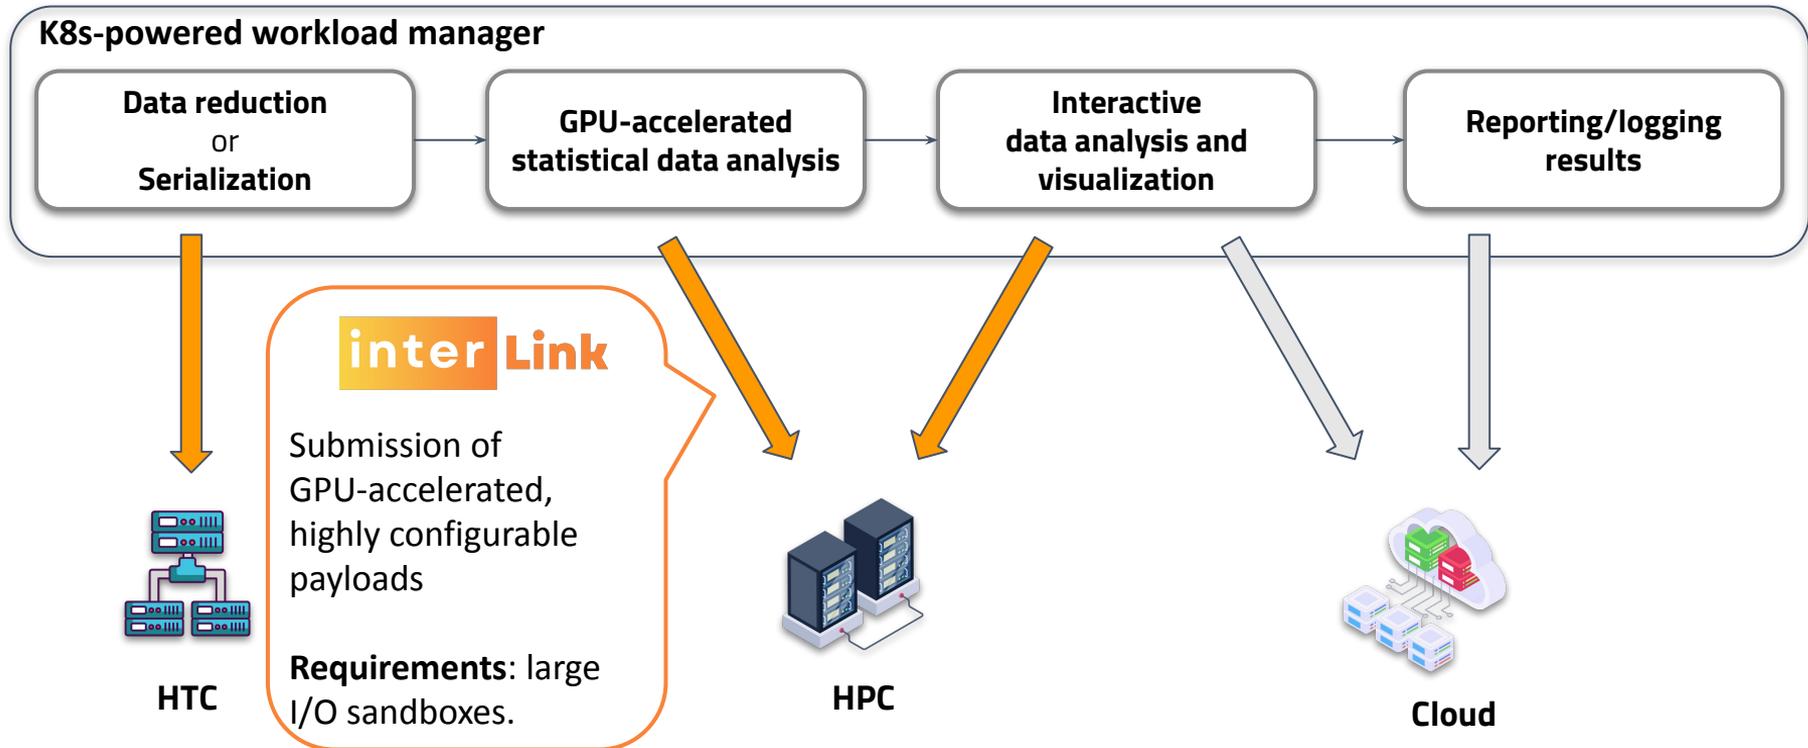
## A typical workflow, offloaded



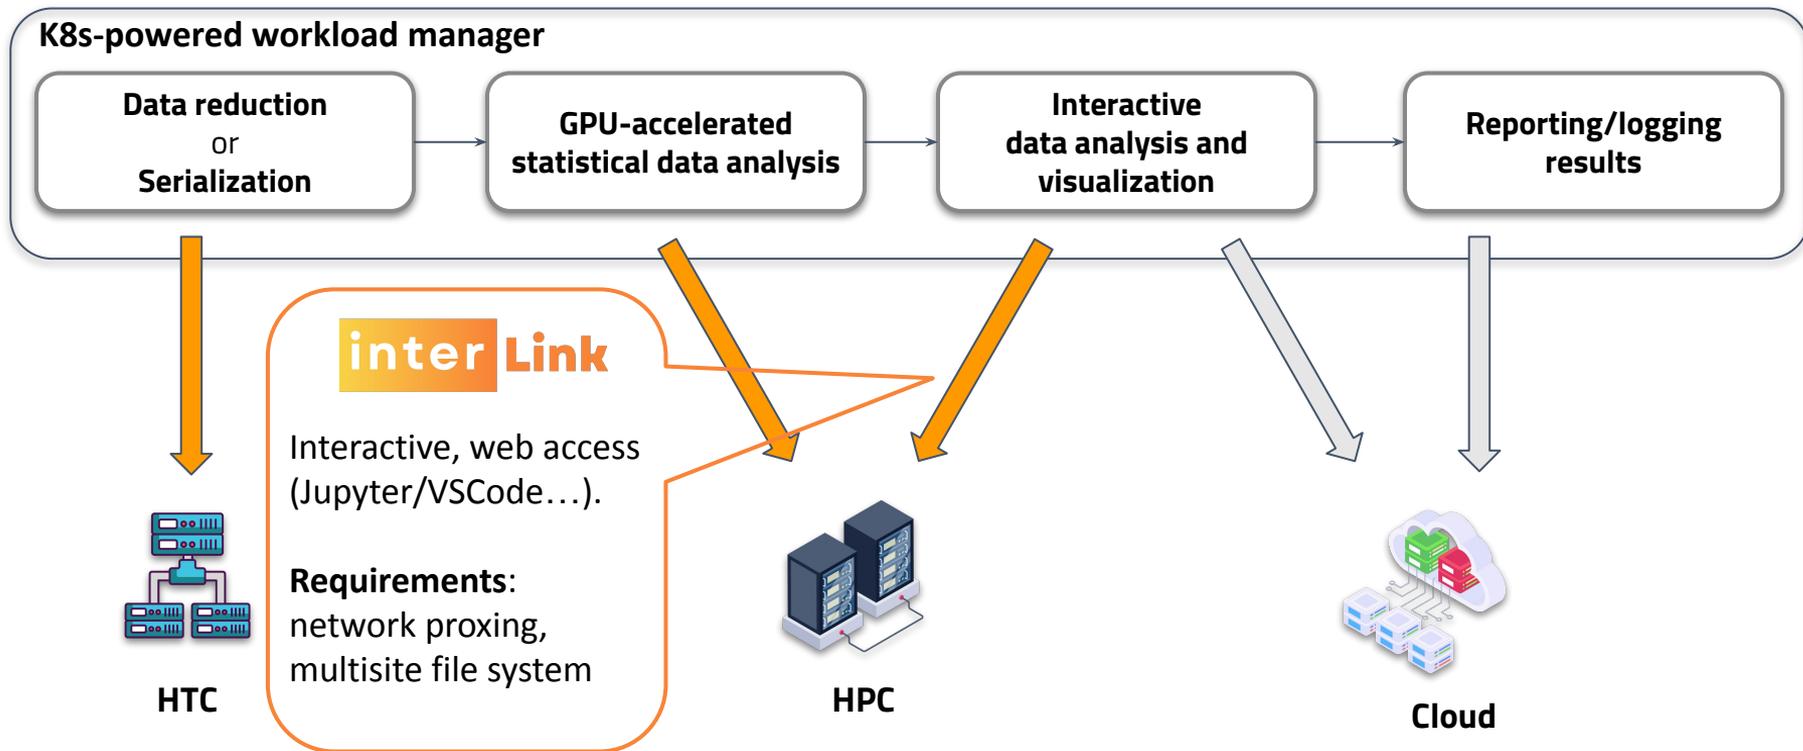
## A typical workflow, offloaded



## A typical workflow, offloaded



## A typical workflow, offloaded



## How does it work?

**Virtual Kubelet** is an open-source project providing an interface between Kubernetes and serverless container platforms.

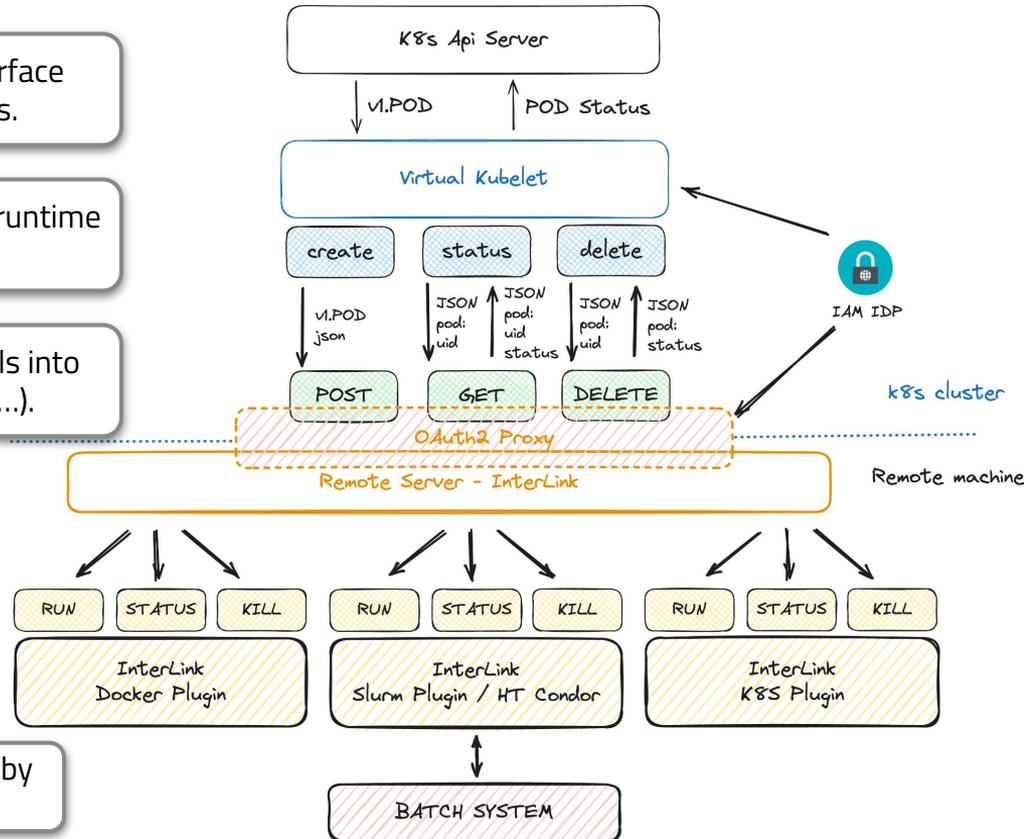
**Interlink** abstracts the interaction with the remote container runtime to few REST APIs: **create**, **status**, **delete**.

**Backend-specific plugins** are developed to translate API calls into the "proper language" for each backend (CLI, REST, gRPC...).

**OAuth2** and **OpenID Connect** are used to authenticate the Virtual Kubelet towards the remote backend.

The user submitting the job to the Kubernetes Workload manager is propagated as meta-data.

**Data**, accessed with remote protocols, can be **cached locally** by the compute backend (e.g. cvmfis)



## Status of the development

	 HTC	 HPC	 Cloud
Submission of self-contained applications	<b>Supported</b>	<b>Supported</b>	<b>Supported</b>
Remote volumes for I/O sandboxes	<b>PoC</b>	<b>PoC</b>	<b>Supported</b>
Network proxying for Jupyter access	<b>Supported</b>	<b>Supported</b>	<b>Supported</b>
Generic network proxying	<b>Ongoing</b>		
Multisite filesystem ( <i>e.g. for scripts and notebooks</i> )	<b>Tech. Tracking</b>		
Monitoring and observability	<b>Supported</b>		

## Scientific use-cases and their workflows

**Flash Simulation** for High Energy Physics experiments.

**Workflow:** Train models (HPC), Deploy models in grid-like simulation (HTC), Statistical validation (Cloud)

**Quasi-interactive HEP data analysis** for Analysis Facilities.

**Workflow:** An interactive application (Cloud) controls parallel, distributed data processing (HTC) and statistical interpretation (HPC)

Interactive **GPU algorithm optimization for astrophysical data analysis.**

**Workflow:** optimize code for astrophysical tasks on GPUs directly through interactive notebook sessions hosted on HPC.

## Summary and next steps

- Kubernetes is the standard for distributed applications, including workload managers.
- Offloading compute-intensive containerized workloads to remote dedicated compute backends with a serverless-computing approach enables complex workflows combining:
  - custom services (cloud)
  - High Throughput Computing (data reduction and grid-like jobs)
  - High Performance Computing (statistical data analysis and machine learning)
- Interlink is being deployed as part of the **Integration PoC** of ICSC and TERABit to integrate:
  - INFN Cloud instances (docker, Kubernetes)
  - CINECA Leonardo (slurm)
  - HTC resources at CNAF (HTCondor)
- Development is progressing fast and use-cases have been identified in the context of ICSC-Spoke 2 and ICSC-Spoke 3 as demonstrators