

# Notebooks and Computing Workflows

Ricardo Rocha - IT-CM-RPS

ABP Computing Forum

<https://indico.cern.ch/event/976158/>

# About

Computing Engineer in CERN IT - Cloud Team

Containers, Docker, Kubernetes, Networking

Machine Learning Infrastructure

Cloud Native Computing Foundation (CNCF) TOC

Reach out for any requests on Kubernetes, Prometheus, Helm, Fluentd, ...

[ricardo.rocha@cern.ch](mailto:ricardo.rocha@cern.ch)

# Today

Workflows

Docker Containers and Reproducibility

Container Orchestration

Notebooks

Pipelines and Distributed Workflows

Some Demos!

Not covered today... but do check them out!

SWAN: <https://swan.cern.ch>

REANA: <https://reana.cern.ch/>

Batch / HTCondor: <https://batchdocs.web.cern.ch/>

Spark

# Workflows



It often starts like this... compile, run, analyse, ...

Hard to keep a clean environment

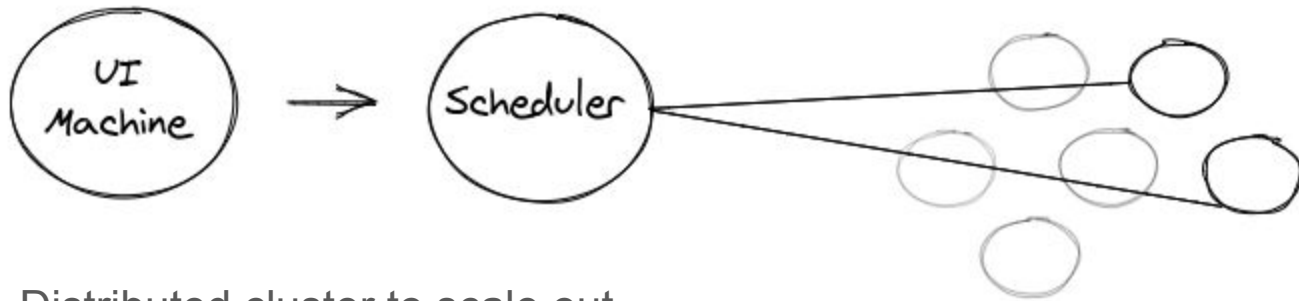
Often hard to find packages for dependencies

Unless everyone agrees on the same system / distribution

Limited resources, limited scalability

Hard to ensure reproducibility, sharing with colleagues

*"It worked in my machine..."*



Distributed cluster to scale out

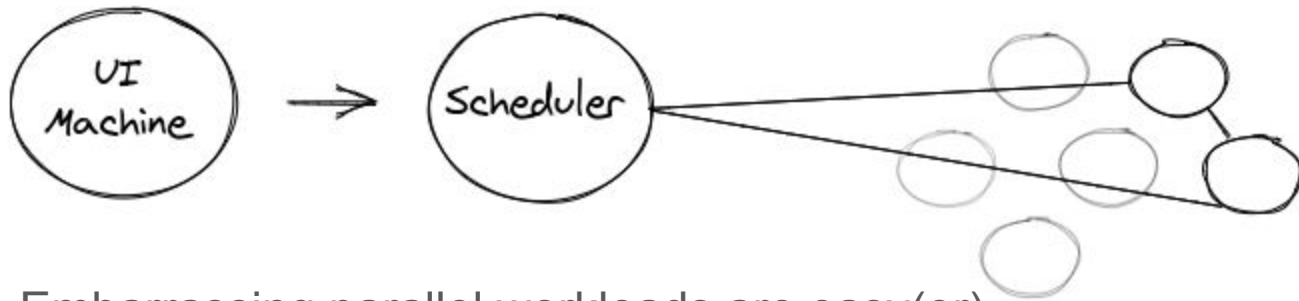
Typical mode for batch systems: HTCondor, SLURM, ...

Worker nodes run one (or a few) sets of systems / distributions

Not trivial to comply with everyone's requirements

How to push specific job dependencies

Shared file-system is the most common solution, but manage / update has issues



Embarrassing parallel workloads are easy(er)

Things get trickier when jobs need to talk to each other

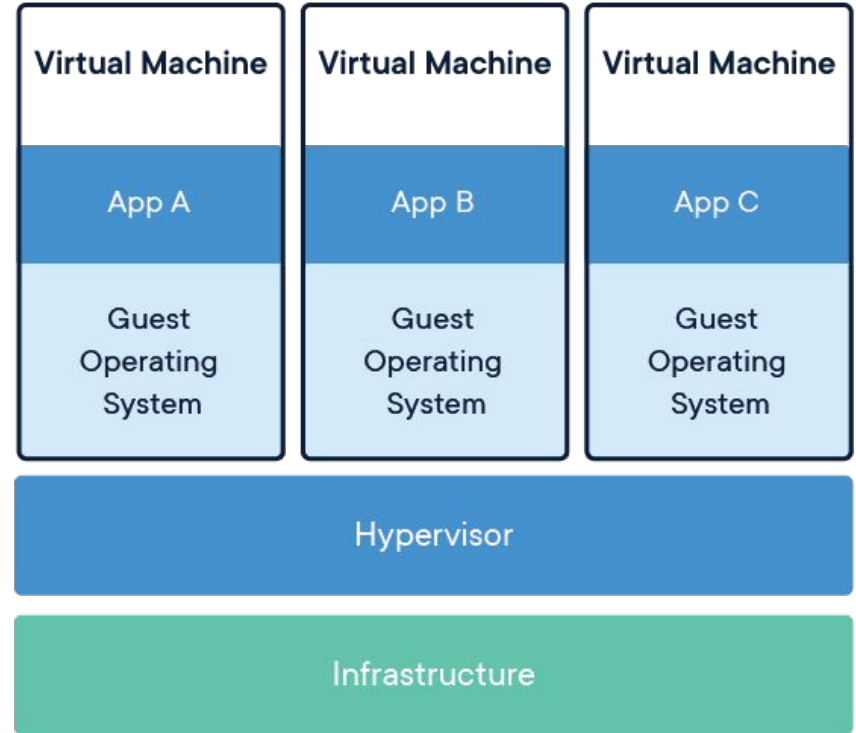
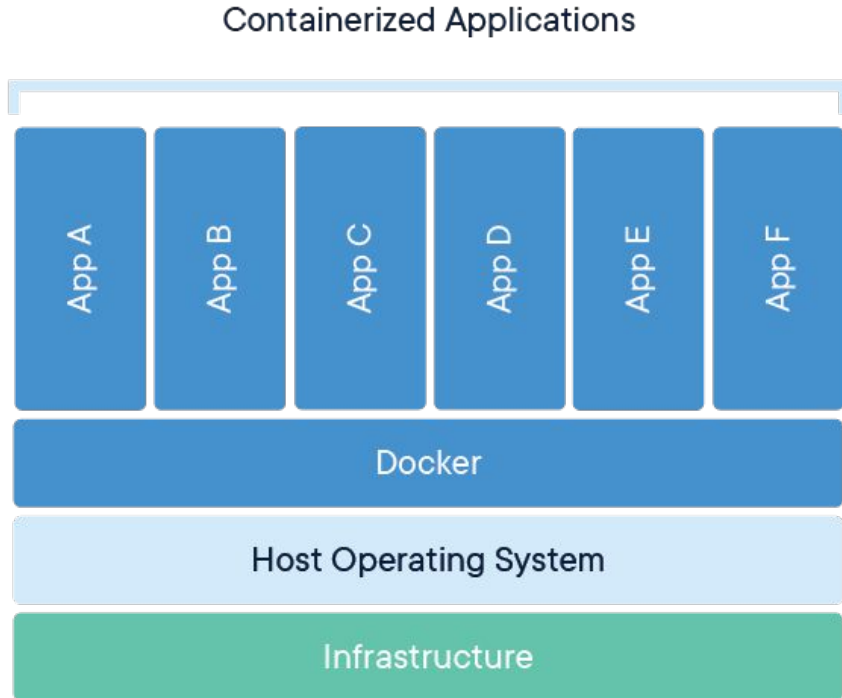
MPI, DAGs, gang scheduling, ...

Systems have their own ways to express these dependencies



# Docker Containers and Reproducibility

# Basic Unit: Docker Container



# Basic Unit: Docker Container

Builds are done locally, or using a CI process, or...

Resulting images are uniquely identified by a sha (immutable) and/or tag

This ensures reproducibility

An image *Push* makes it available to others in a registry

Pulled images are cached locally (pulled only the first time)



# Quick Demo

# Container Orchestration with Kubernetes

Manage containerized deployments in multi-node clusters

- Support for heterogeneous setups by grouping nodes

Many low level concepts covering compute, networking and storage

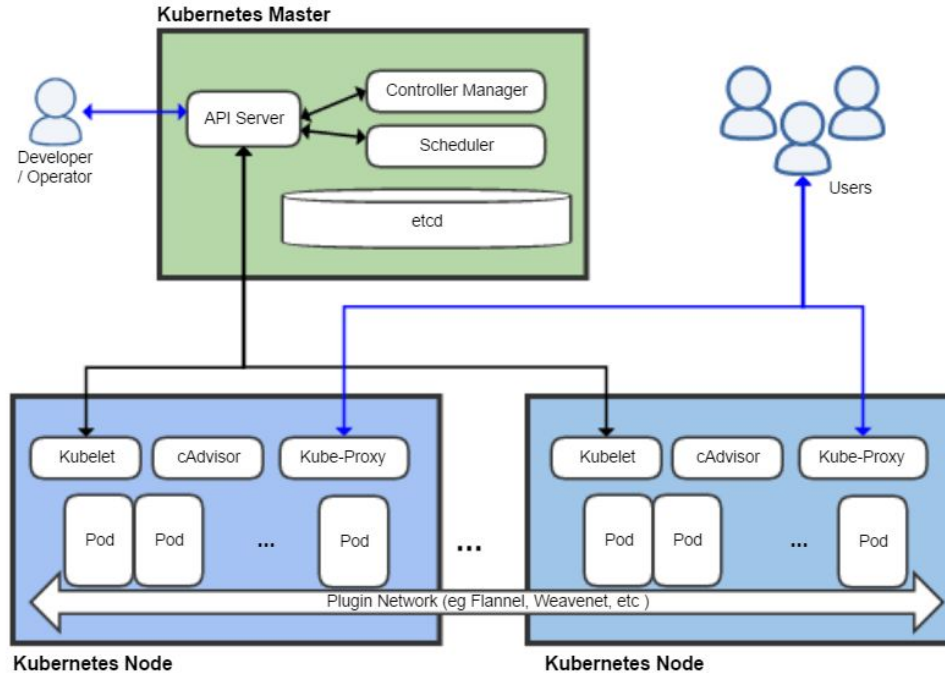
- We need to tell what type of resources are required for a container to run

Ways to express replication, failure domains, disruption budgets, etc

Ways to express policies for auto scaling of workloads

<https://clouddocs.web.cern.ch/containers/README.html>

# Container Orchestration with Kubernetes



# Kubernetes

Lingua franca of the cloud

Managed services offered by all major public clouds

Multiple options for on-premise or self-managed deployments

Common declarative API for basic infrastructure : compute, storage, networking

Healthy ecosystem of tools offering extended functionality



Google Cloud Platform



OPENSIFT



Rancher  
Kubernetes Engine



MAGNUM  
an OpenStack Community Project



kubeadm





**70 TB Dataset**



**OpenStack Magnum**

**25000 Kubernetes Jobs**



**Job Results**



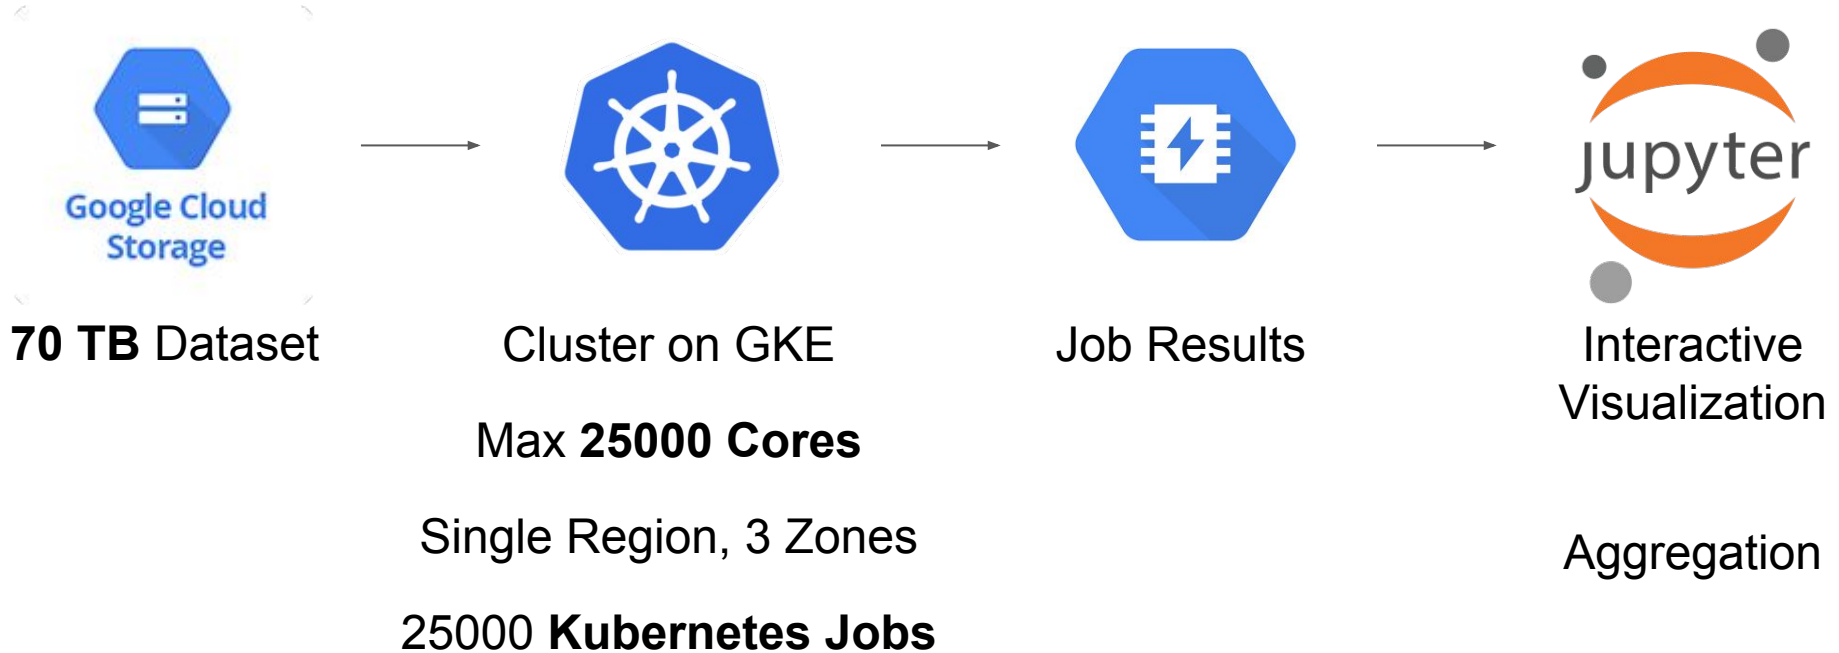
**Interactive  
Visualization**

**Aggregation**





Google Cloud



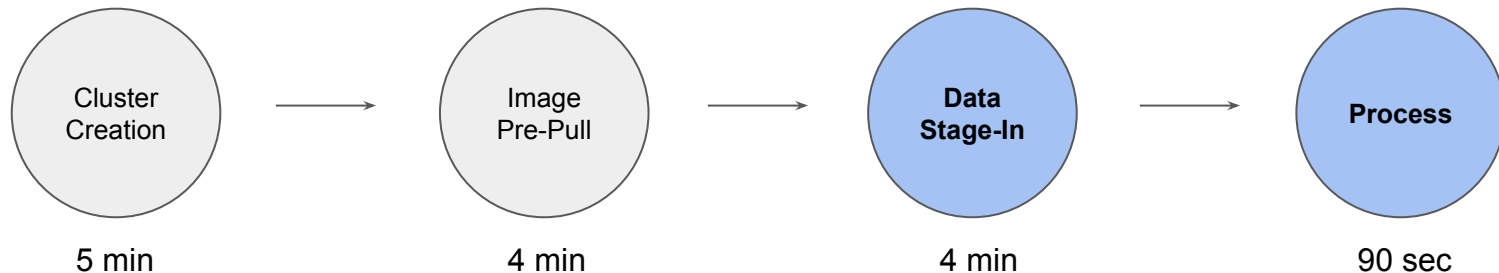
# GCP Analysis Run

Kubernetes clusters on GKE ( Managed Kubernetes service on GCP )

Run included

~1200 nodes: n1-highmem-16, 104 GB RAM

~20000 cores, ~120 TB RAM



# Quick Demo

# Jupyter Notebooks

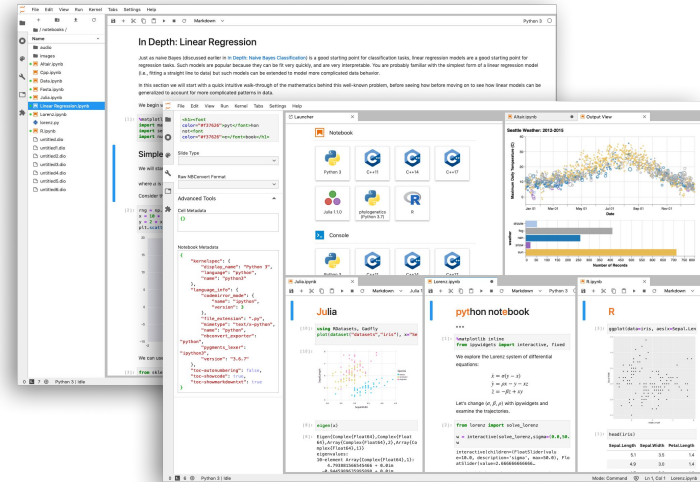
Easy way to create and share documents with code, equations, ...

Can be run locally, but popular to run in a central service: JupyterHub

Started with python, but kernel support for R, Go, C++, etc

Extremely popular for data science

And machine learning...



# JupyterHub

Central web application managing multi-user environments

Popular to rely on Docker containers for isolation

Common Issue: how to manage dependencies not in an image

*pip install --user mypkg* has issues when sharing, reproducibility

Someone needs to rebuild the image, push

Or you just rely on an alternative to get extra dependencies (shared fs?)

Some automation is required, we'll see another option later

# Dask

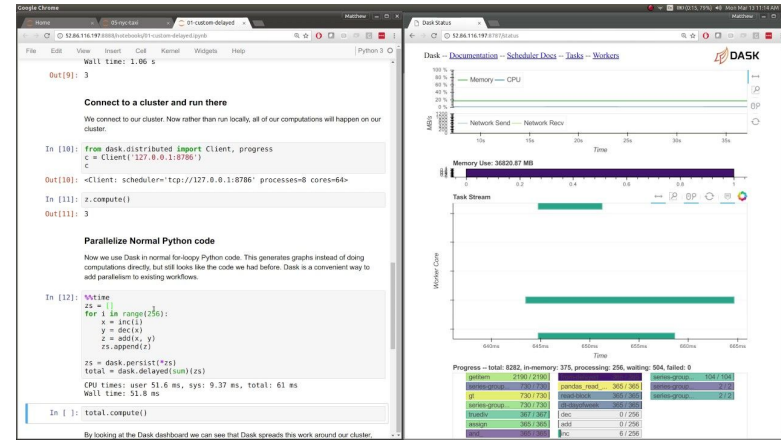
Advanced parallelism for analytics

Integrates well with Pandas, numpy, scikit-learn

Familiar if you're used to Python

Extreme scalability: thousands of nodes, tens of thousands of workers

Backends for Kubernetes, SLURM, HTCondor, among others



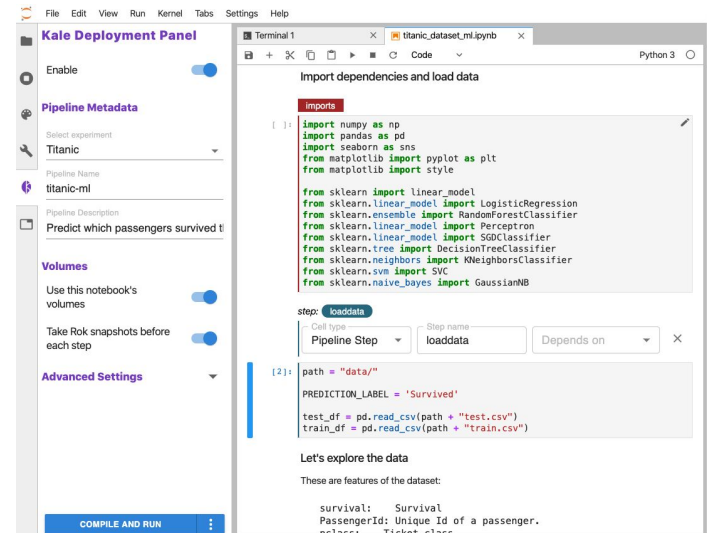
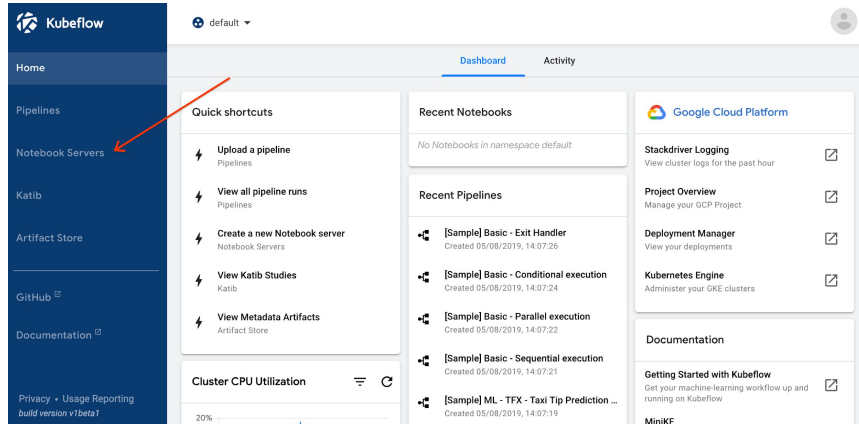
# Quick Demo

# Kubeflow + Kale

Kubeflow is a platform to do machine learning / data science on Kubernetes

Kale is a Jupyter extension to ease moving from notebooks to pipelines

Early access almost ready for a CERN instance





# Quick Demo

# Summary

Docker helps with sharing computing environments and reproducibility

Kubernetes is taking a big place in orchestrating computing workflows

And much more in the service area

Lots of options for data analysis, with notebooks being a very popular option

Platforms like Dask and Kubeflow try to help scaling out these workloads

# Questions?

<https://clouddocs.web.cern.ch/containers/README.html>

<https://binder.cern.ch>

<https://ml.docs.cern.ch>

<https://swan.cern.ch>

<https://reana.cern.ch/>

<https://batchdocs.web.cern.ch/>