Virgo Computing Centres Workshop - November 28-29, 2019

INFN TORINO
Istituto Nazionale di Fisica Nucleare
Sezione di Torino

Virtual Clusters

Sara Vallero
For the INFN Torino team.

Small and/or not dedicated Computing Centers that would like to contribute to the IGWN computing assets

Small or not dedicated centers

- I already have an HTCondor pool
- I don’t, but I do have a CE…

  … then GlideinWMS
- I don’t have any of the above…

  … then follow this talk

Goal setting

Enable Virgo offline computing on diverse infrastructures (scalability)

Uniformity of execution environment

- Portable virtualized computing infrastructure/platform

Smooth deployment and management of the computing platform

- Infrastructure as Code (IaC) paradigm

What is a Virtual Cluster

A way to hide to the end-user the complicated underlying infrastructure and provide a well known **interface**.

What is a Virtual Cluster

A way to hide to the end-user the complicated underlying infrastructure and provide a well known **interface**.

Application

Pipelines

Physicist Domain

What is a Virtual Cluster

A way to hide to the end-user the complicated underlying infrastructure and provide a well known **interface**.

Application — **Pipelines**

Physicist Domain

Platform — **HTCondor, Spark…**

Abstraction layer — **Virtualization and Orchestration**

Computing resources — **CPUs, GPUs, storage, network**

Computing Expert Domain

Offline Analysis

Requirement for the target Computing Center:

a IaaS Cloud
(OpenStack or OpenNebula)

Dynamic **O**n **D**emand **A**nalysis **S**ervice

- one of the *Thematic Services* of the EOSC-hub H2020 project

- allows to easily configure and orchestrate a complex pool of container based microservices → build a **customized platform for software execution**

- currently supports:

  - **HTCondor** Batch-System-as-a-Service

  - Big Data platform for ML-as-a-Service (**Spark+HDFS**)

  - extension of the two integrating community specific needs

Dynamic On Demand Analysis Service

INFN TORINO
Istituto Nazionale di Fisica Nucleare
Sezione di Torino

ADMIN

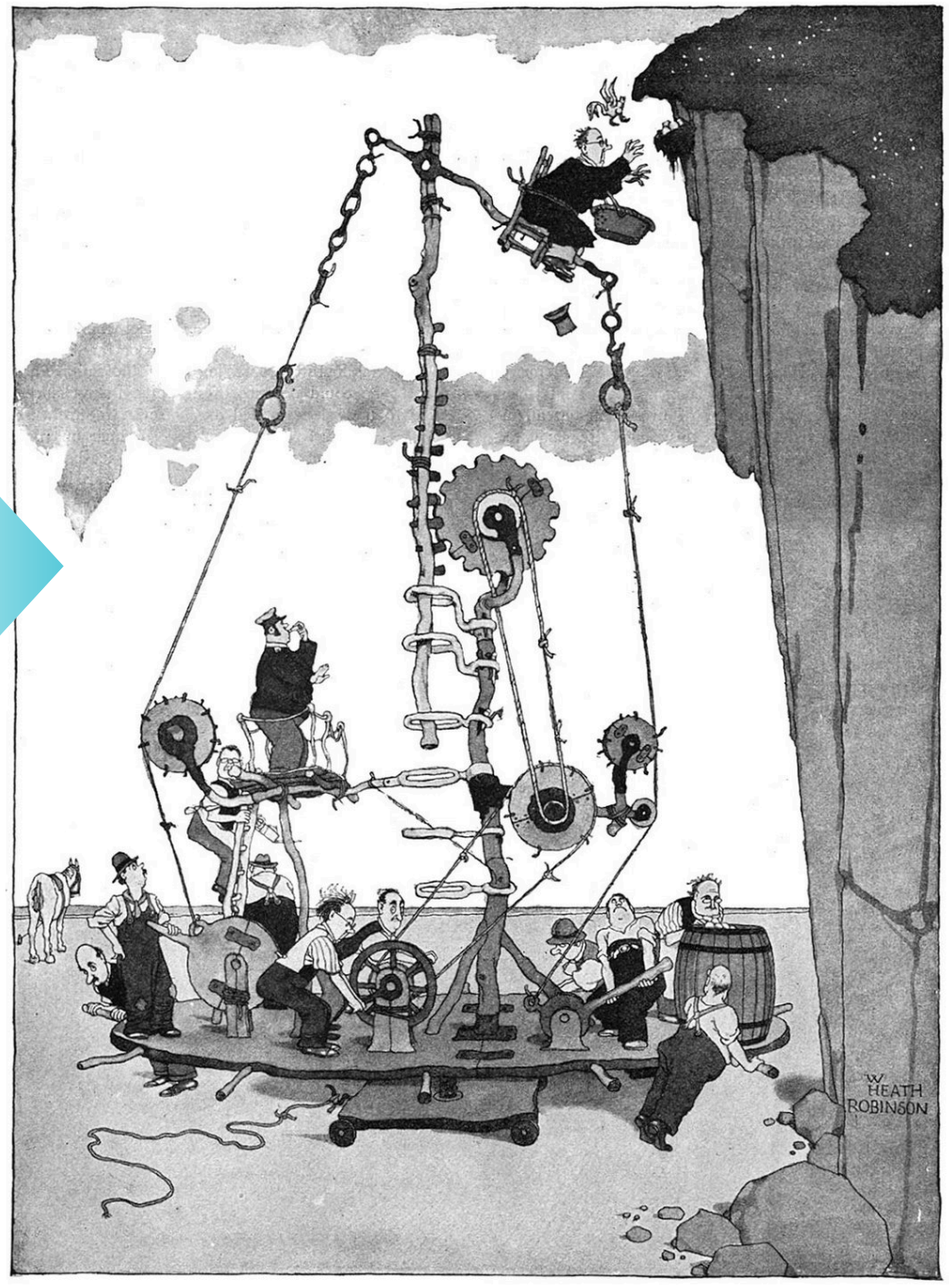

USER

CVMFS

DATA (CACHE)

HTCondor
High Throughput Computing

APACHE Spark™

TOSCA template

DODAS machinery

Exposed services

Dynamic On Demand Analysis Service

INFN TORINO
Istituto Nazionale di Fisica Nucleare
Sezione di Torino

TOSCA
template

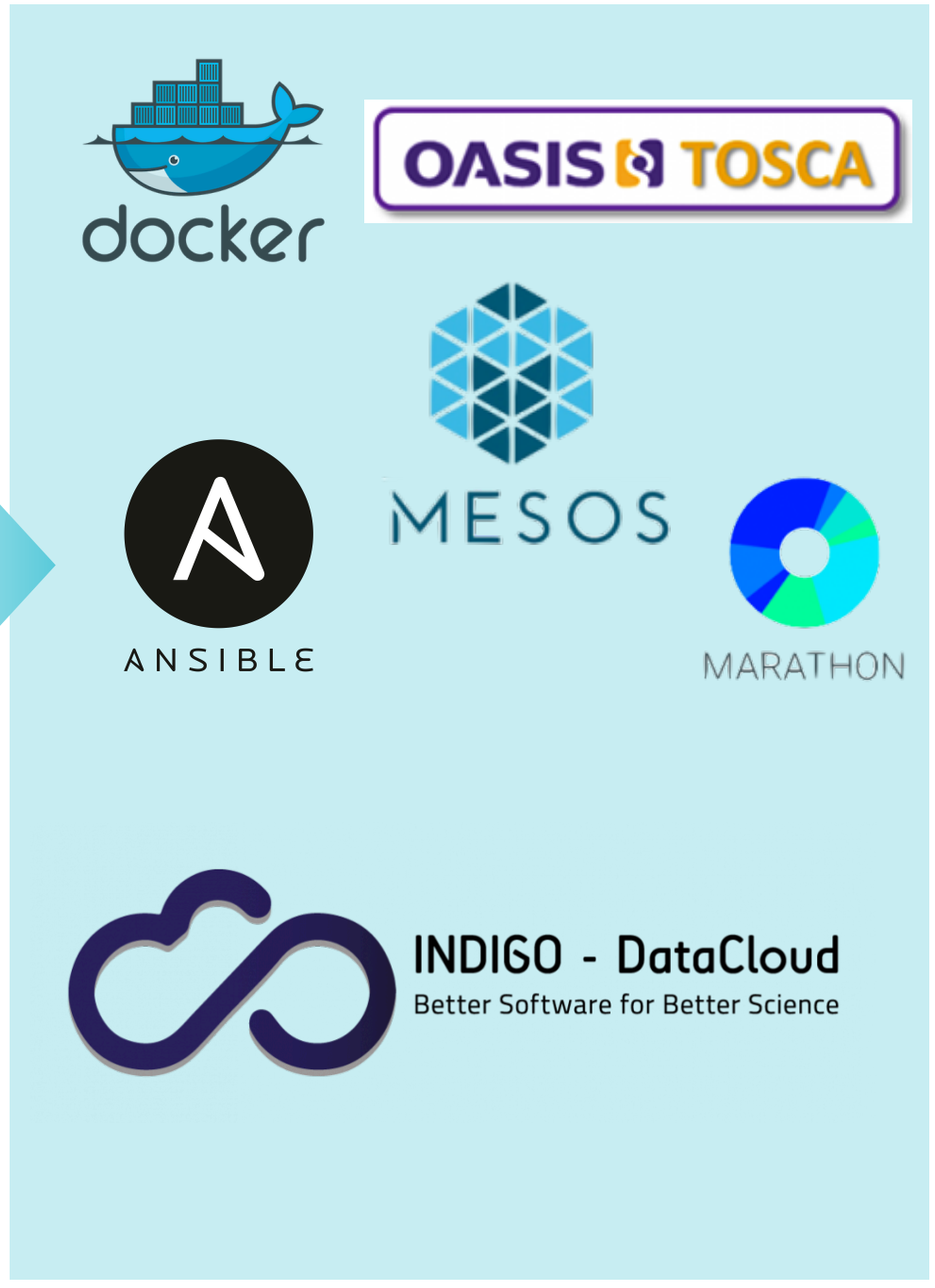

DODAS
machinery

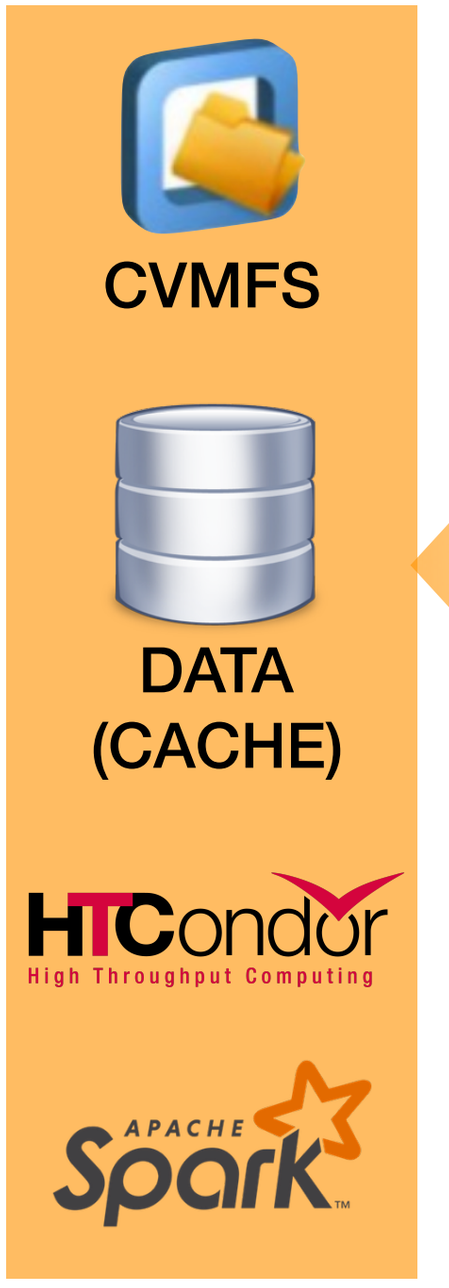

Exposed
services

Dynamic On Demand Analysis Service

INFN TORINO
Istituto Nazionale di Fisica Nucleare
Sezione di Torino

TOSCA template

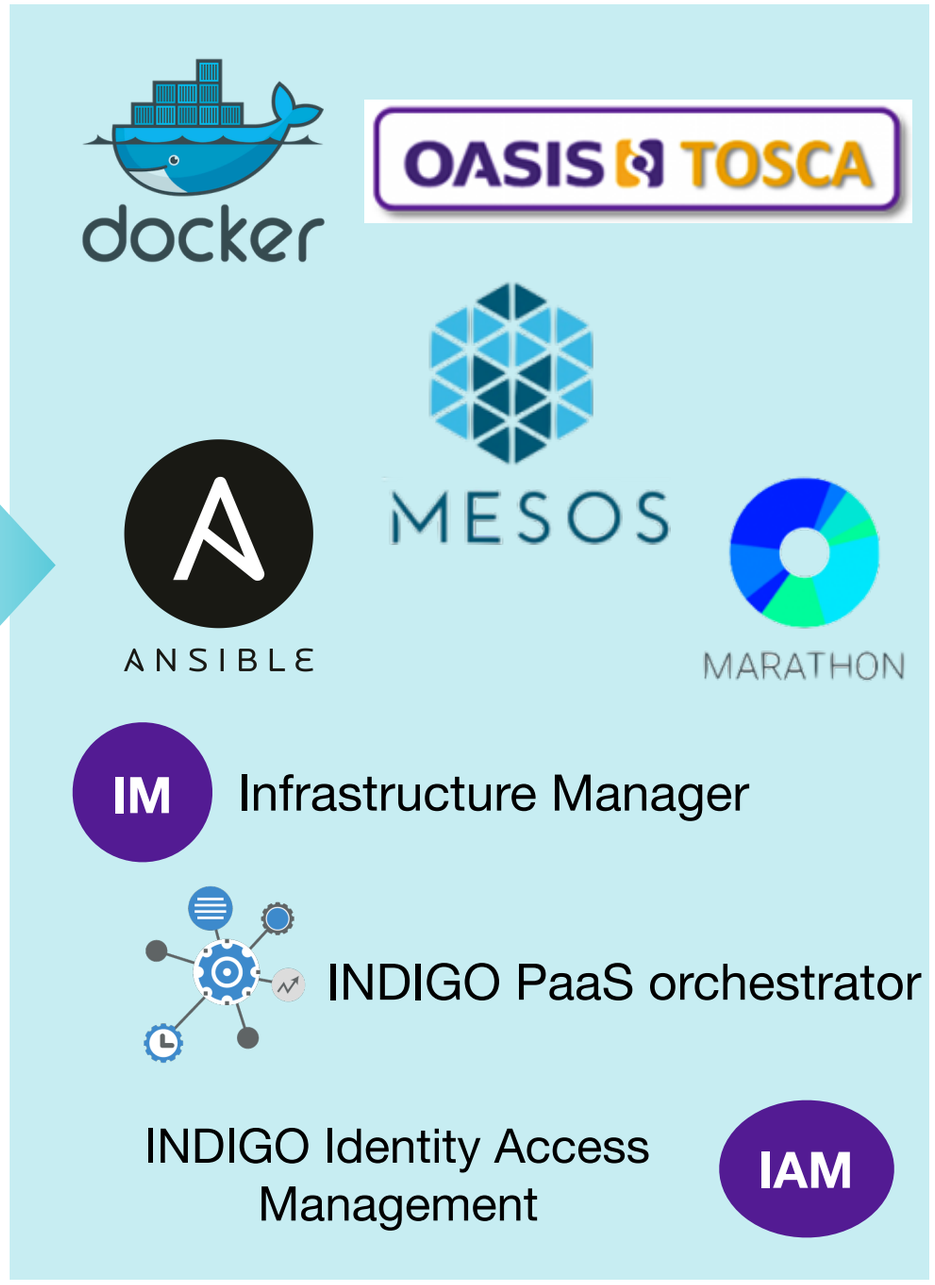

DODAS machinery

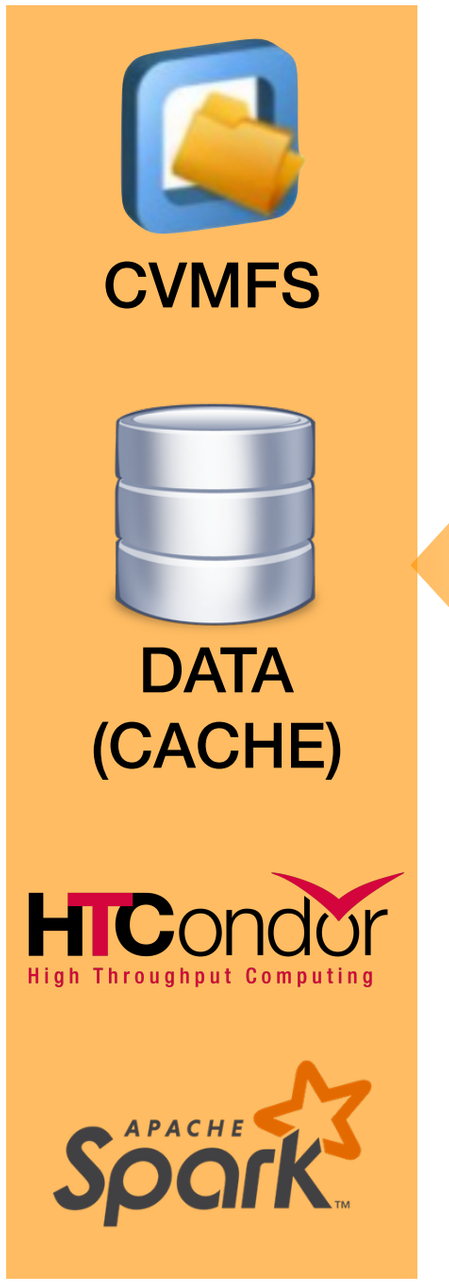

Exposed services

Discussion

- Any site interested in DODAS?

Virgo specific template available for HTCondor

- Does anybody need Spark?

Low Latency Searches

(workflow offload proposal)

Requirement for the target Computing Center:

a production grade orchestrator
(i.e. Kubernetes)

Low-latency workflow (simplified)

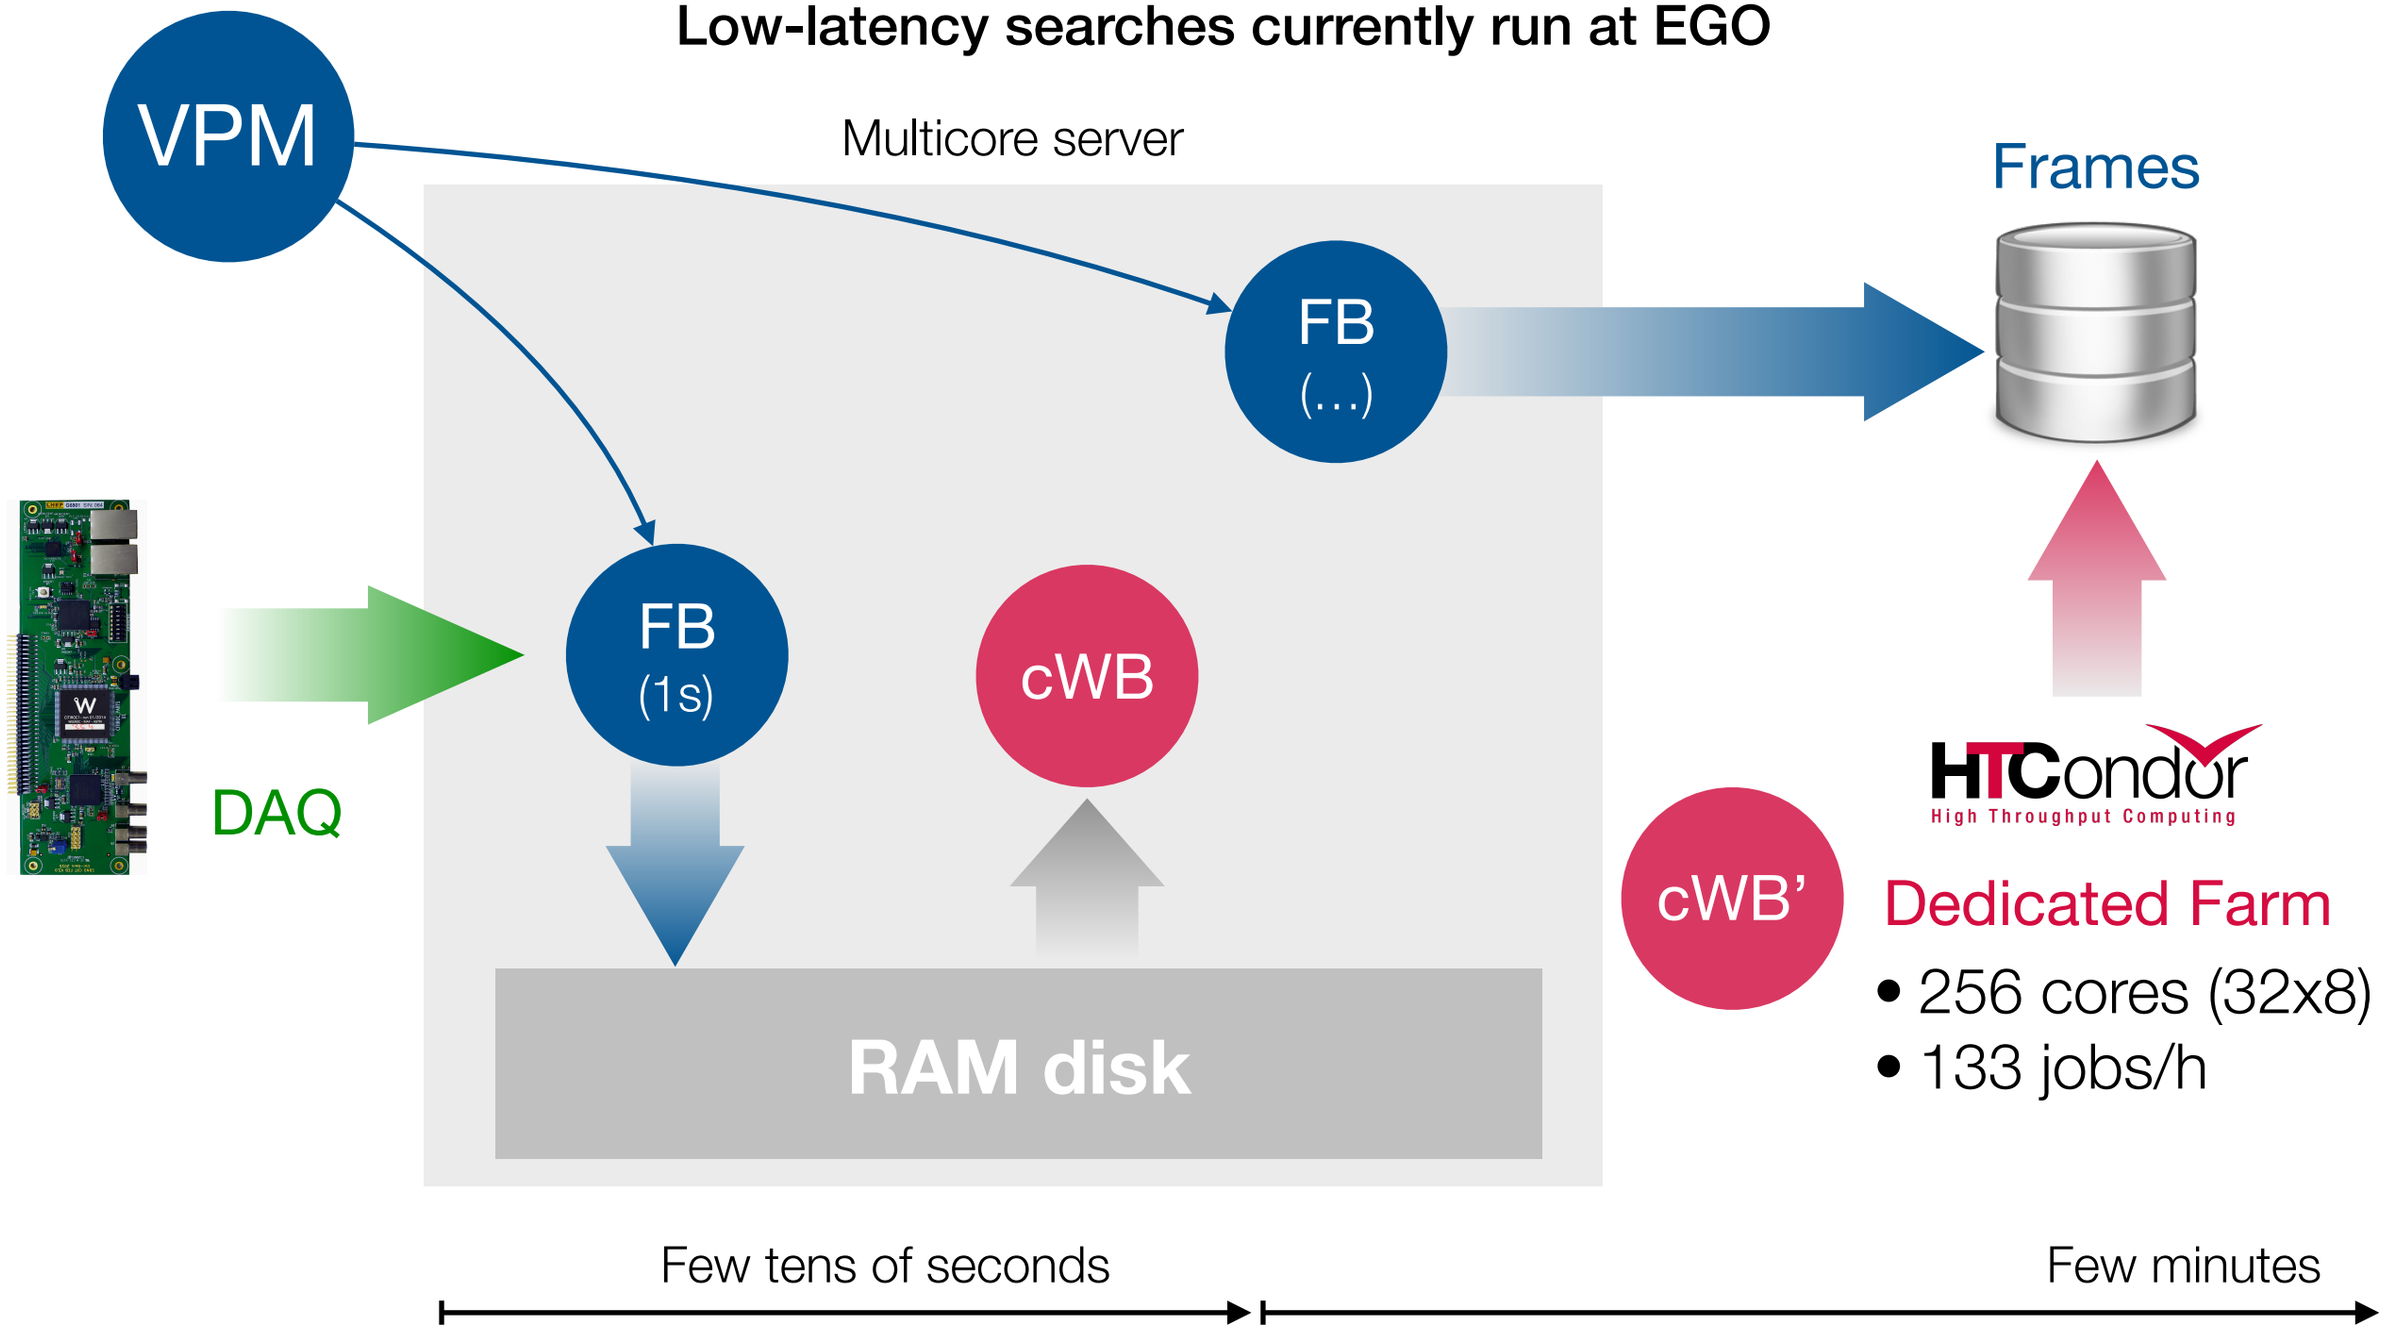

VPM = Virgo Process Monitor
FB = Frame Builder

Requirements

- Dedicated resources for low latency and continuous batch processing
- Immediately available (few seconds) resources to accomodate higher loads in batch processing:
    - Shared pool is not ok (unknown latency due to queue)
    - Virtual Machines (VMs) on-demand might be too slow, therefore Linux Containers on top of: bare metal or running virtual machines
- Fault-tolerance
- Scalability
- Portable solution (deploy at any site, high availability)
- Agile updates (make changes, roll-back)

Idea

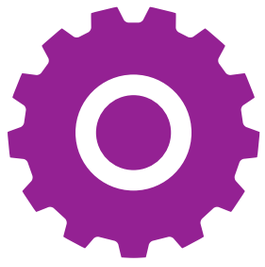

PROCESS

LINUX CONTAINER
=
PROCESS + ENVIRONMENT

ORCHESTRATION
=
Coordinate the execution of several containers to reach the desired goal

Example implementation

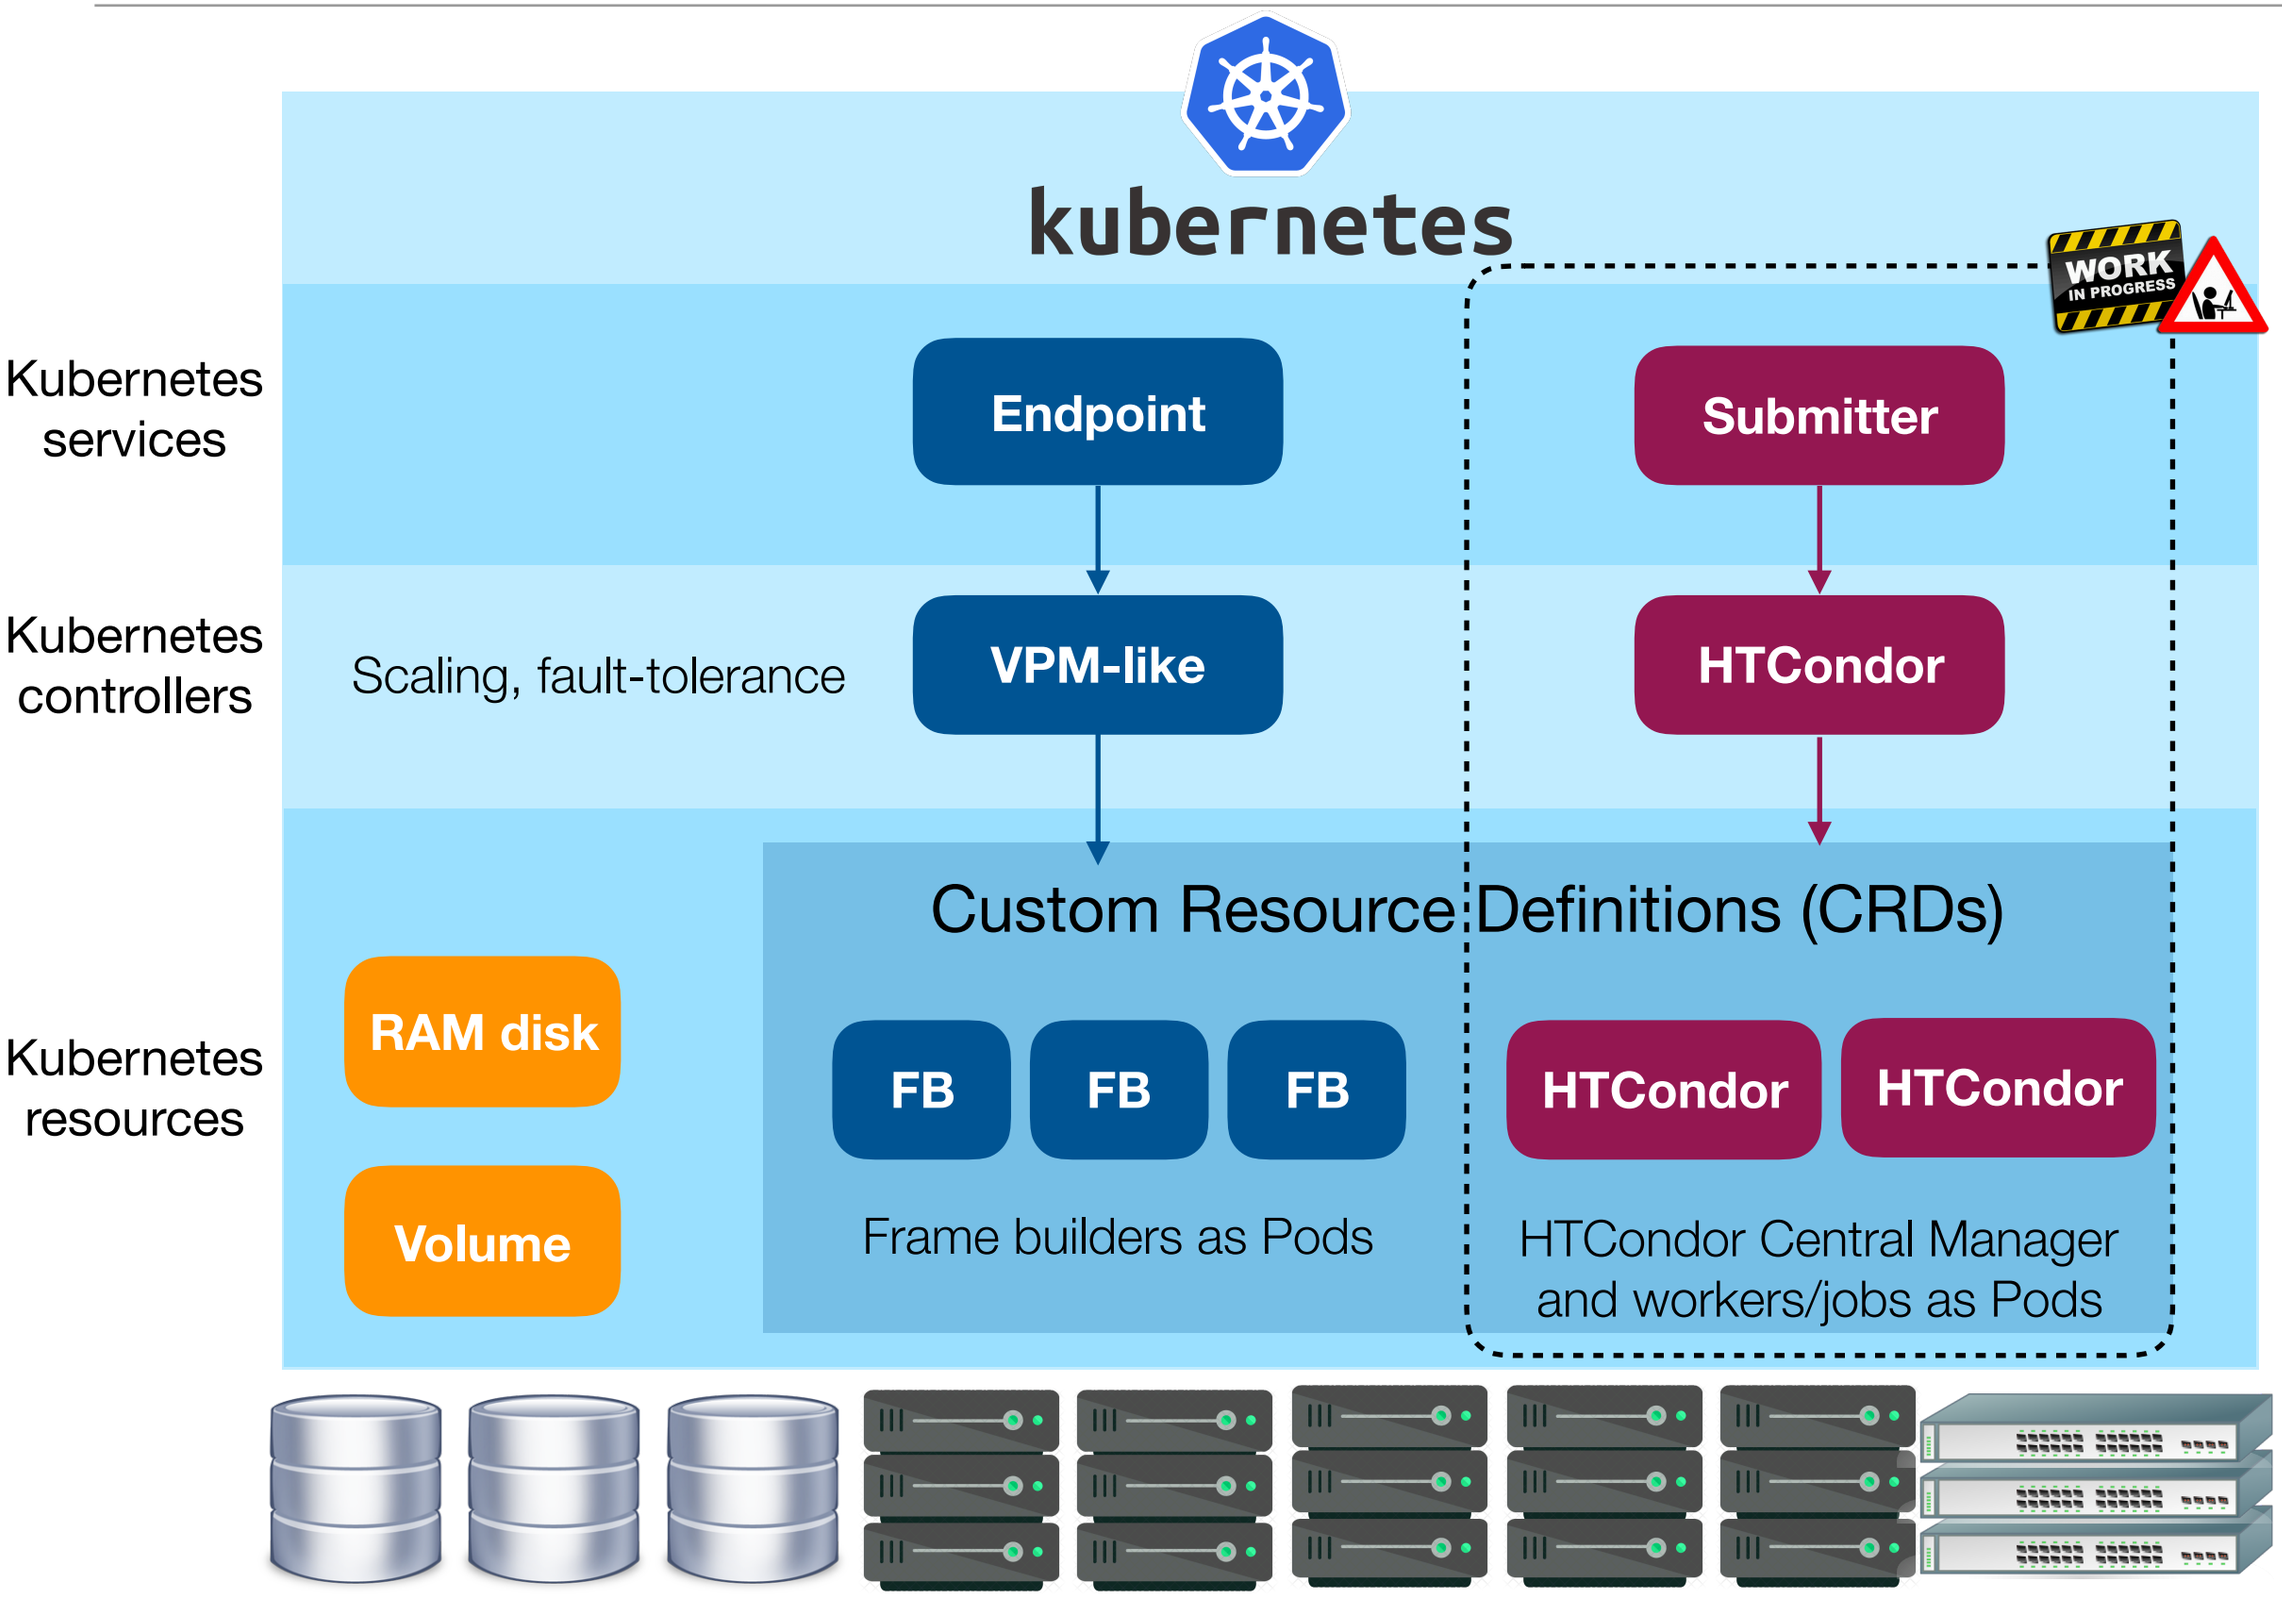

Discussion

- Do we want to offload? YES
- Why taking the burden of a virtualized solution?
  (for the reasons outlined in the Requirements)
- Candidate site? CNAF
- Anyone willing to do the job?