

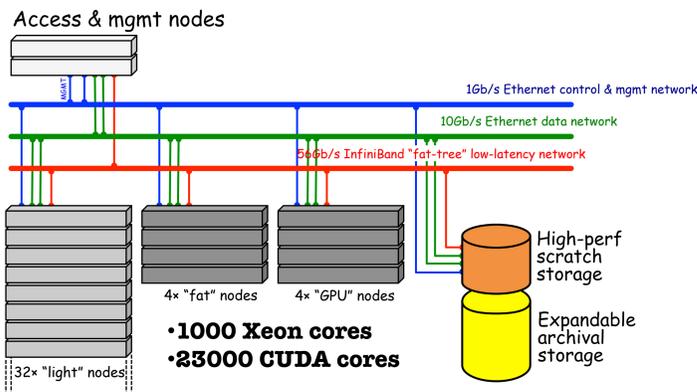
- The **Open Computing Cluster for Advanced data Manipulation (OCCAM)** is a HPC facility designed and operated by a collaboration between Università degli Studi di Torino and INFN-Torino. It is aimed at providing a **flexible** and **multi-purpose** infrastructure to cater to a broad range of scientific computing needs, as well as a platform for R&D activities on computational technologies themselves.
- The broad and variegated range of use cases challenges for great flexibility in providing resources to the final users. We choose to borrow some ideas from **Cloud Computing** technologies introducing the concept of **Computing Applications**, each defined by its runtime environment, resource requirements and an execution model.

PILOT USE CASES AND ARCHITECTURE

HPC: batch-like, multi-node workloads using MPI and inter-node communication

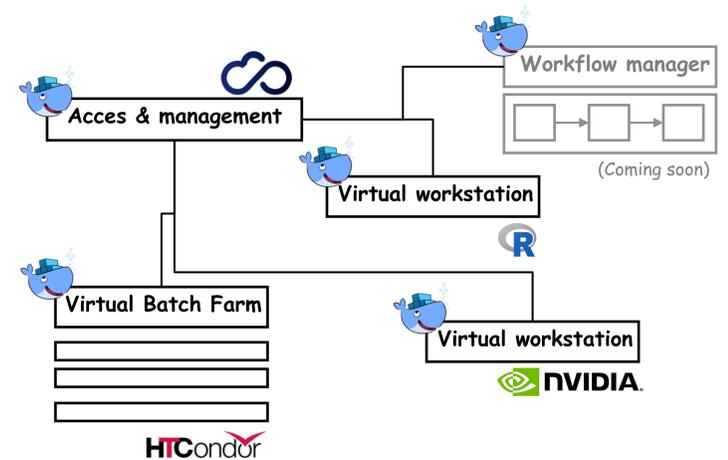
Genomic pipelines: multi-step data analysis requiring high-memory large single-image nodes

Virtual workstation: code execution (e.g. R or ROOT) in a single multicore node, possibly with GPU acceleration



The system comprises **different types of computing nodes** (standard, high-memory 4-socket and GPU-accelerated), a high-performance scratch storage for intensive random access, powered by a parallel filesystem, and a larger partition for data staging. All components are connected by 10Gb/s Ethernet and InfiniBand FDR networks.

By dynamically partitioning the system, we deploy **consistent and sandboxed sub-clusters** tailored to a Computing Application's requirements. The pivotal technologies for the middleware architecture are **Linux containers**, currently managed with Docker, and several building blocks developed by the **INDIGO-DataCloud project**.



OCCAM AND INDIGO

- **HTMesos** is an activity within the INDIGO-DataCloud project aimed at providing a Batch-system-as-a-Service platform for Scientific Computing, delivering a well consolidated computational framework, while complying to modern computing paradigms.
- **udocker** is another INDIGO product, a tool to run Docker containers in userspace. We use it for running user-defined images in the simpler use cases.



- **Docker:** Industry-standard containerization platform
- Used to partition the system into **isolated virtual clusters** to run Computing Applications
- Also, self-packaging decouples infrastructure from application software management

- **Apache Mesos:** resource abstraction and management
- **Mesosphere Marathon:** long-running services scheduling and monitoring
- Used to schedule, deploy and manage Computing Applications

- **Calico:** enable secure IP communication between containers
- Used to **manage isolated networks** for Virtual Clusters

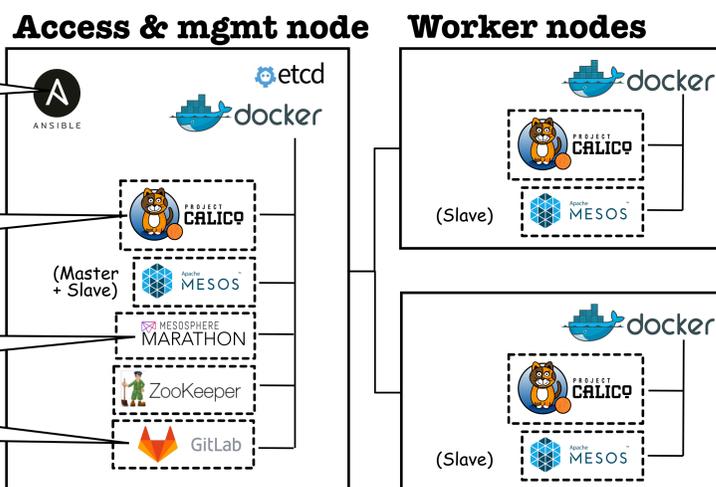
- **HTCondor:** a batch scheduler widely used in the scientific community
- Used to provide a familiar user experience for batch-like use cases

Ansible is used throughout the system for configuration management, using INDIGO-developed roles and playbooks

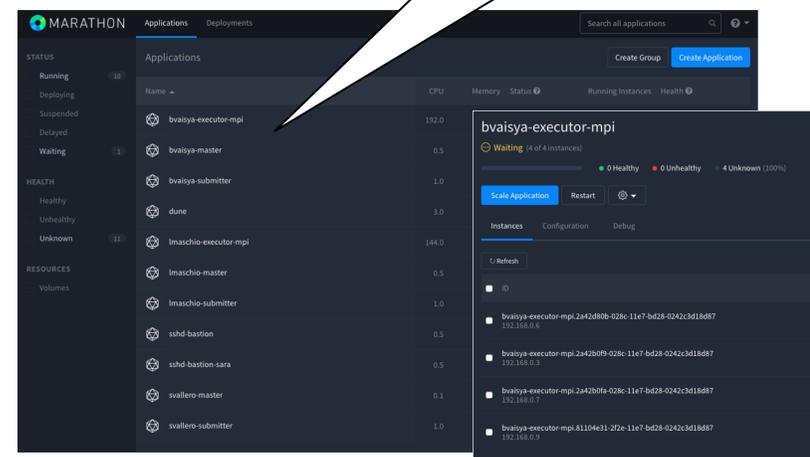
Calico manages isolated networks for each virtual farm. Access is granted through a container in Access node.

Marathon schedules and monitors user- or system- defined containers on worker nodes

GitLab is the access portal, providing user management, private image registry, continuous integration,...



HTCondor component containers (Master, Submitter and Executor) deployed as Marathon applications



Marco Aldinucci^{2,3}, Stefano Bagnasco^{1,3}, Matteo Concas^{2,3}, Stefano Lusso^{1,3}, Paolo Pasteris^{2,3}, Sergio Rabellino^{2,3}, Sara Vallero^{1,3}

¹ Istituto Nazionale di Fisica Nucleare, sezione di Torino; ² Department of Computer Science, University of Torino; ³ Scientific Computing Competence Centre, University of Torino