INDIGO - DataCloud
Better Software for Better Science

Geographically distributed Batch System as a Service: the INDIGO-DataCloud approach exploiting HTCondor

Speaker: Vallero Sara (INFN Torino)

CHEP 2016 - October 10/14 - San Francisco (CA)

Authors

- intro ad INDIGO
- link al talk generale
- lista autori

The scientist's way

A typical end-user experience…

ssh to the remote head-node and submit jobs

The end-user would like to:

- have full control over system and software (know the root password)
- use all resources… now!

or alternatively…

- a system tailored to her needs
- easy to use and modify
- forget about computing and do research!

The administrator point of view…

- have full control over system and software (do not reveal the root password)
- prevent users from breaking the system and disrupt other people's work (isolation)
- do not tailor the system to everybody's needs (keep it simple)
- a system easy to maintain (IaaS)

Front-end

Batch system, MPI

Operating System

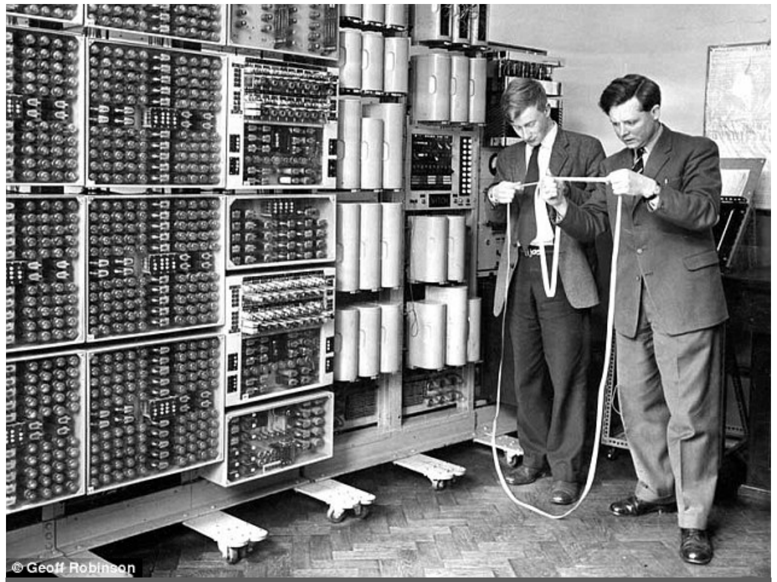

The WITCH computer, first used in the 1950s

Infrastructure

Virtualisation is the way to go, but a mindset shift is needed…

INDIGO DataCloud

A new computing paradigm

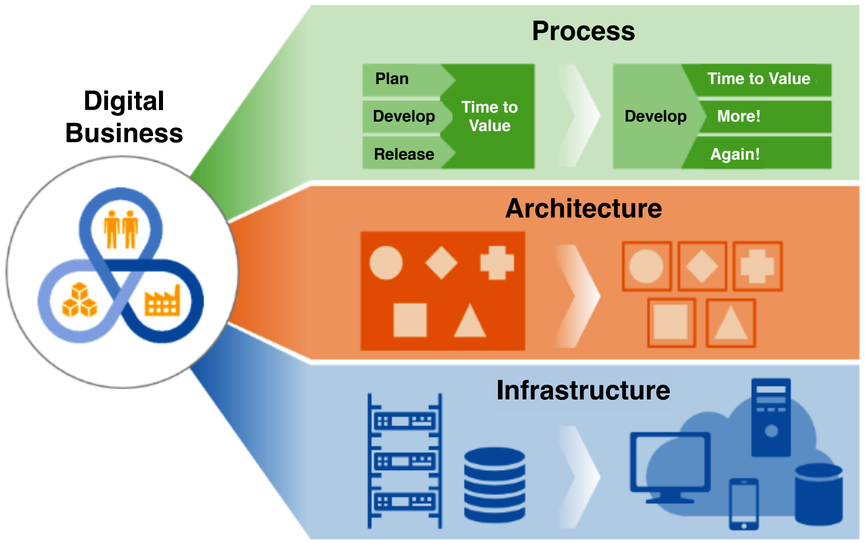

Image from Gartner (March 2016)

Encapsulating tasks into higher and higher abstractions is being shown by various large companies to provide a competitive advantage.

According to GARTNER:
Cloud Computing is a style of computing in which **scalable and elastic** IT-enabled capabilities are delivered **as a service** using Internet technologies.

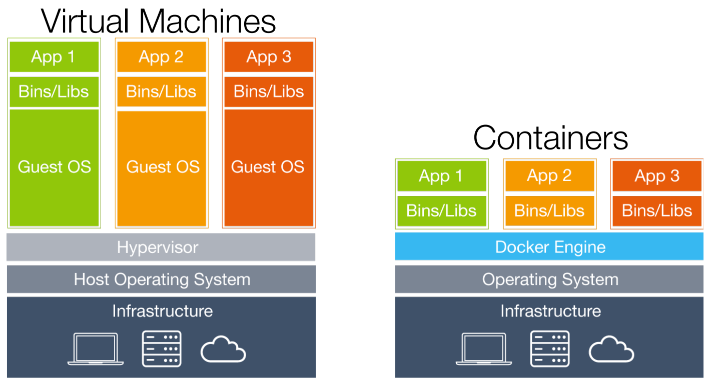

From Virtual Machines to Containers:

- package, ship and run distributed application components with guaranteed platform parity across different environments
- *democratising* virtualisation by providing it to developers in a usable, application-focused form
- access to virtual machine virtualisation tends to be provided through, and governed by, gatekeepers in infrastructure and operations
- Docker is being adopted from the ground up by developers using a DevOps approach

INDIGO DataCloud

Toolbox

- **Docker:** https://www.docker.com/

the leading software containerisation platform. Docker containers wrap a piece of software in a complete filesystem that contains everything needed to run: code, runtime, system tools, system libraries - anything that can be installed on a server. This guarantees that the software will always run the same, regardless of its environment.

- **Apache Mesos:** http://mesos.apache.org/

abstracts CPU, memory, storage, and other compute resources away from machines (physical or virtual), enabling fault-tolerant and elastic distributed systems to easily be built and run effectively.
Mesos is built using the same principles as the Linux kernel, only at a different level of abstraction. The Mesos kernel runs on every machine and provides applications (e.g., Hadoop, Spark, Kafka, Elasticsearch) with API's for resource management and scheduling across entire datacenter and cloud environments.

- **Calico:** https://www.projectcalico.org/

a pure L3 approach to data center networking… enable secure IP communication between virtual machines, containers, or bare metal workloads.
Based on the same scalable IP network principles as the Internet, Calico implements a highly efficient vRouter in each compute node that leverages the existing Linux kernel forwarding engine without the need for vSwitches. Each vRouter propagates workload reachability information (routes) to the rest of the data center using BGP - either directly in small scale deployments or via BGP route reflectors to reach Internet level scales in large deployments.
Calico peers directly with the data center's physical fabric (whether L2 or L3) without the need for on/off ramps, NAT, tunnels, or overlays.
Calico supports rich and flexible network policy which it enforces using bookended ACLs on each compute node to provide tenant isolation, security groups, and external reachability constraints.

- **Gartner:** http://www.gartner.com/technology/home.jsp

INDIGO
DataCloud

Our mission

Key components of the INDIGO PaaS framework.

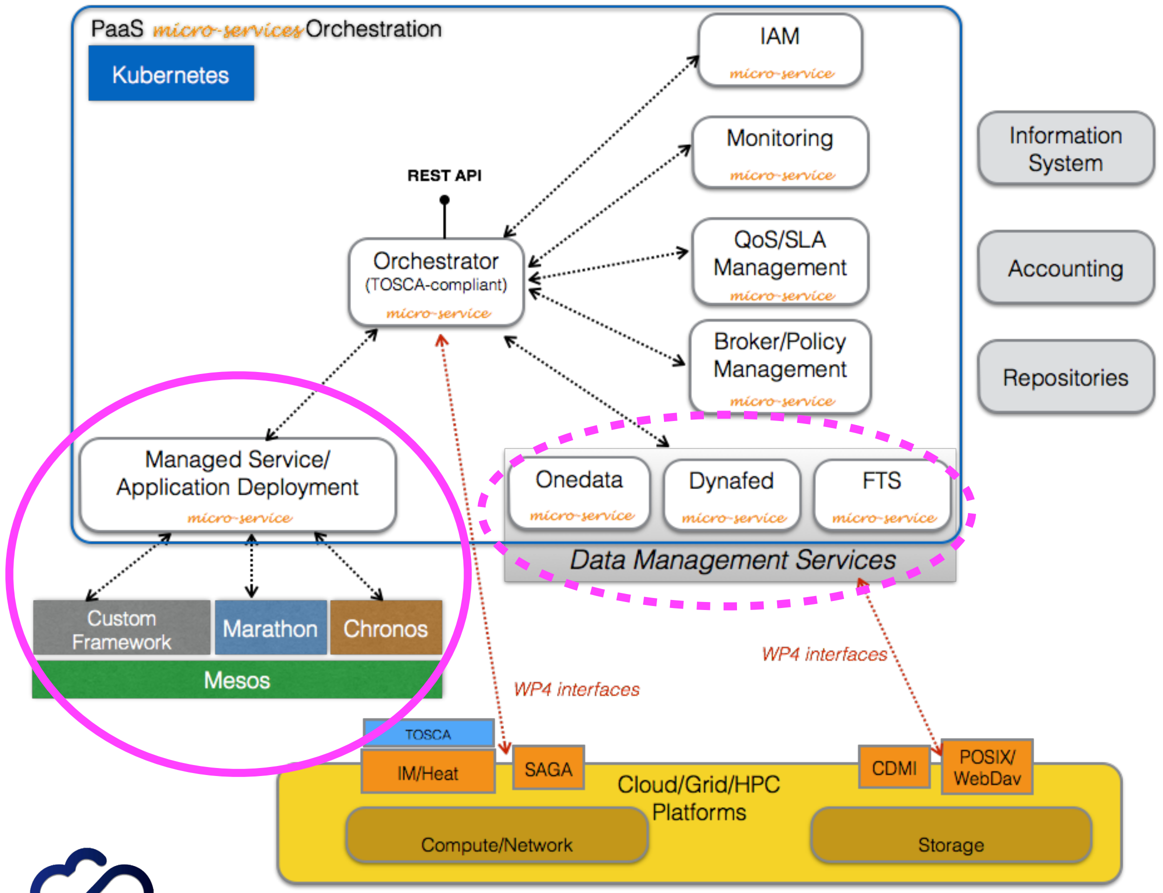

INDIGO DataCloud

The goal:

- keep delivering a well consolidated computational framework, while complying to modern computing paradigms
- ease system administration to all levels (hardware/applications)
- provide a smooth end-user experience

Batch System as a Service:

- automatically and dynamically deploy a complete batch system cluster (with appropriate user interfaces) in highly-available and scalable configurations
- use HTCondor:
  - widely used within the scientific community
  - cloud aware
- networking:
  - HTCondor shared port mode
  - overlay networks
  - span the cluster over multiple sites (Connection Broker)
- storage:
  - scalability, performance and reliability
  - CVMFS (using Parrot)
  - integrate with the INDIGO Data Services (Onedata, Dynafed, FTS)

Two approaches

Mesos

Framework

- fine grained control on the application's tasks
- ad hoc authorisation and scaling rules
- isolation and packaging achieved with the Docker Containerizer module of Mesos
- components:
    - master
    - scheduler (implements policies on the resource offers)
    - executors (to launch tasks on the slave nodes)
- can be used in standalone or as a component in the INDIGO PaaS system

Marathon

- using Marathon framework (health checks, failover capabilities)
- master, scheduler and executors packaged in separate pre-configured Docker containers
- containers are deployed ad Long Running services

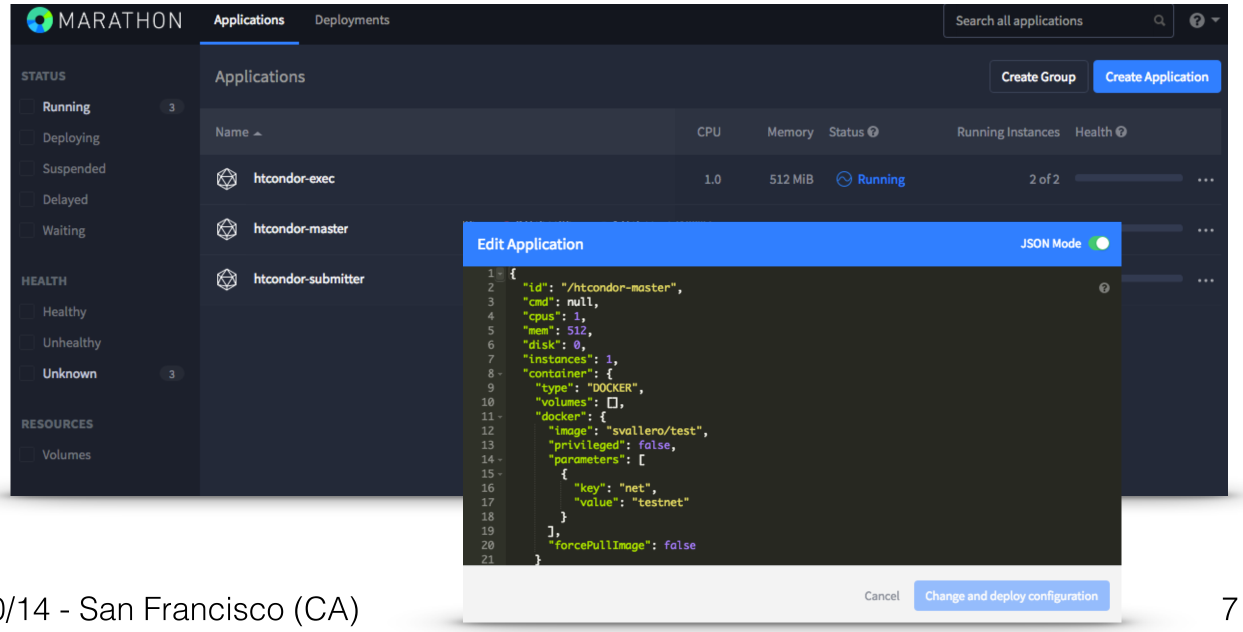

Name	CPU	Memory	Status	Running Instances	Health
htcondor-exec	1.0	512 MiB	Running	2 of 2	
htcondor-master					
htcondor-submitter					

```
{
  "id": "/htcondor-master",
  "cmd": null,
  "cpus": 1,
  "mem": 512,
  "disk": 0,
  "instances": 1,
  "container": {
    "type": "DOCKER",
    "volumes": [],
    "docker": {
      "image": "svallero/test",
      "privileged": false,
      "parameters": [
        {
          "key": "net",
          "value": "testnet"
        }
      ],
      "forcePullImage": false
    }
```

INDIGO DataCloud

GEOGRAPHICAL DEPLOYMENT

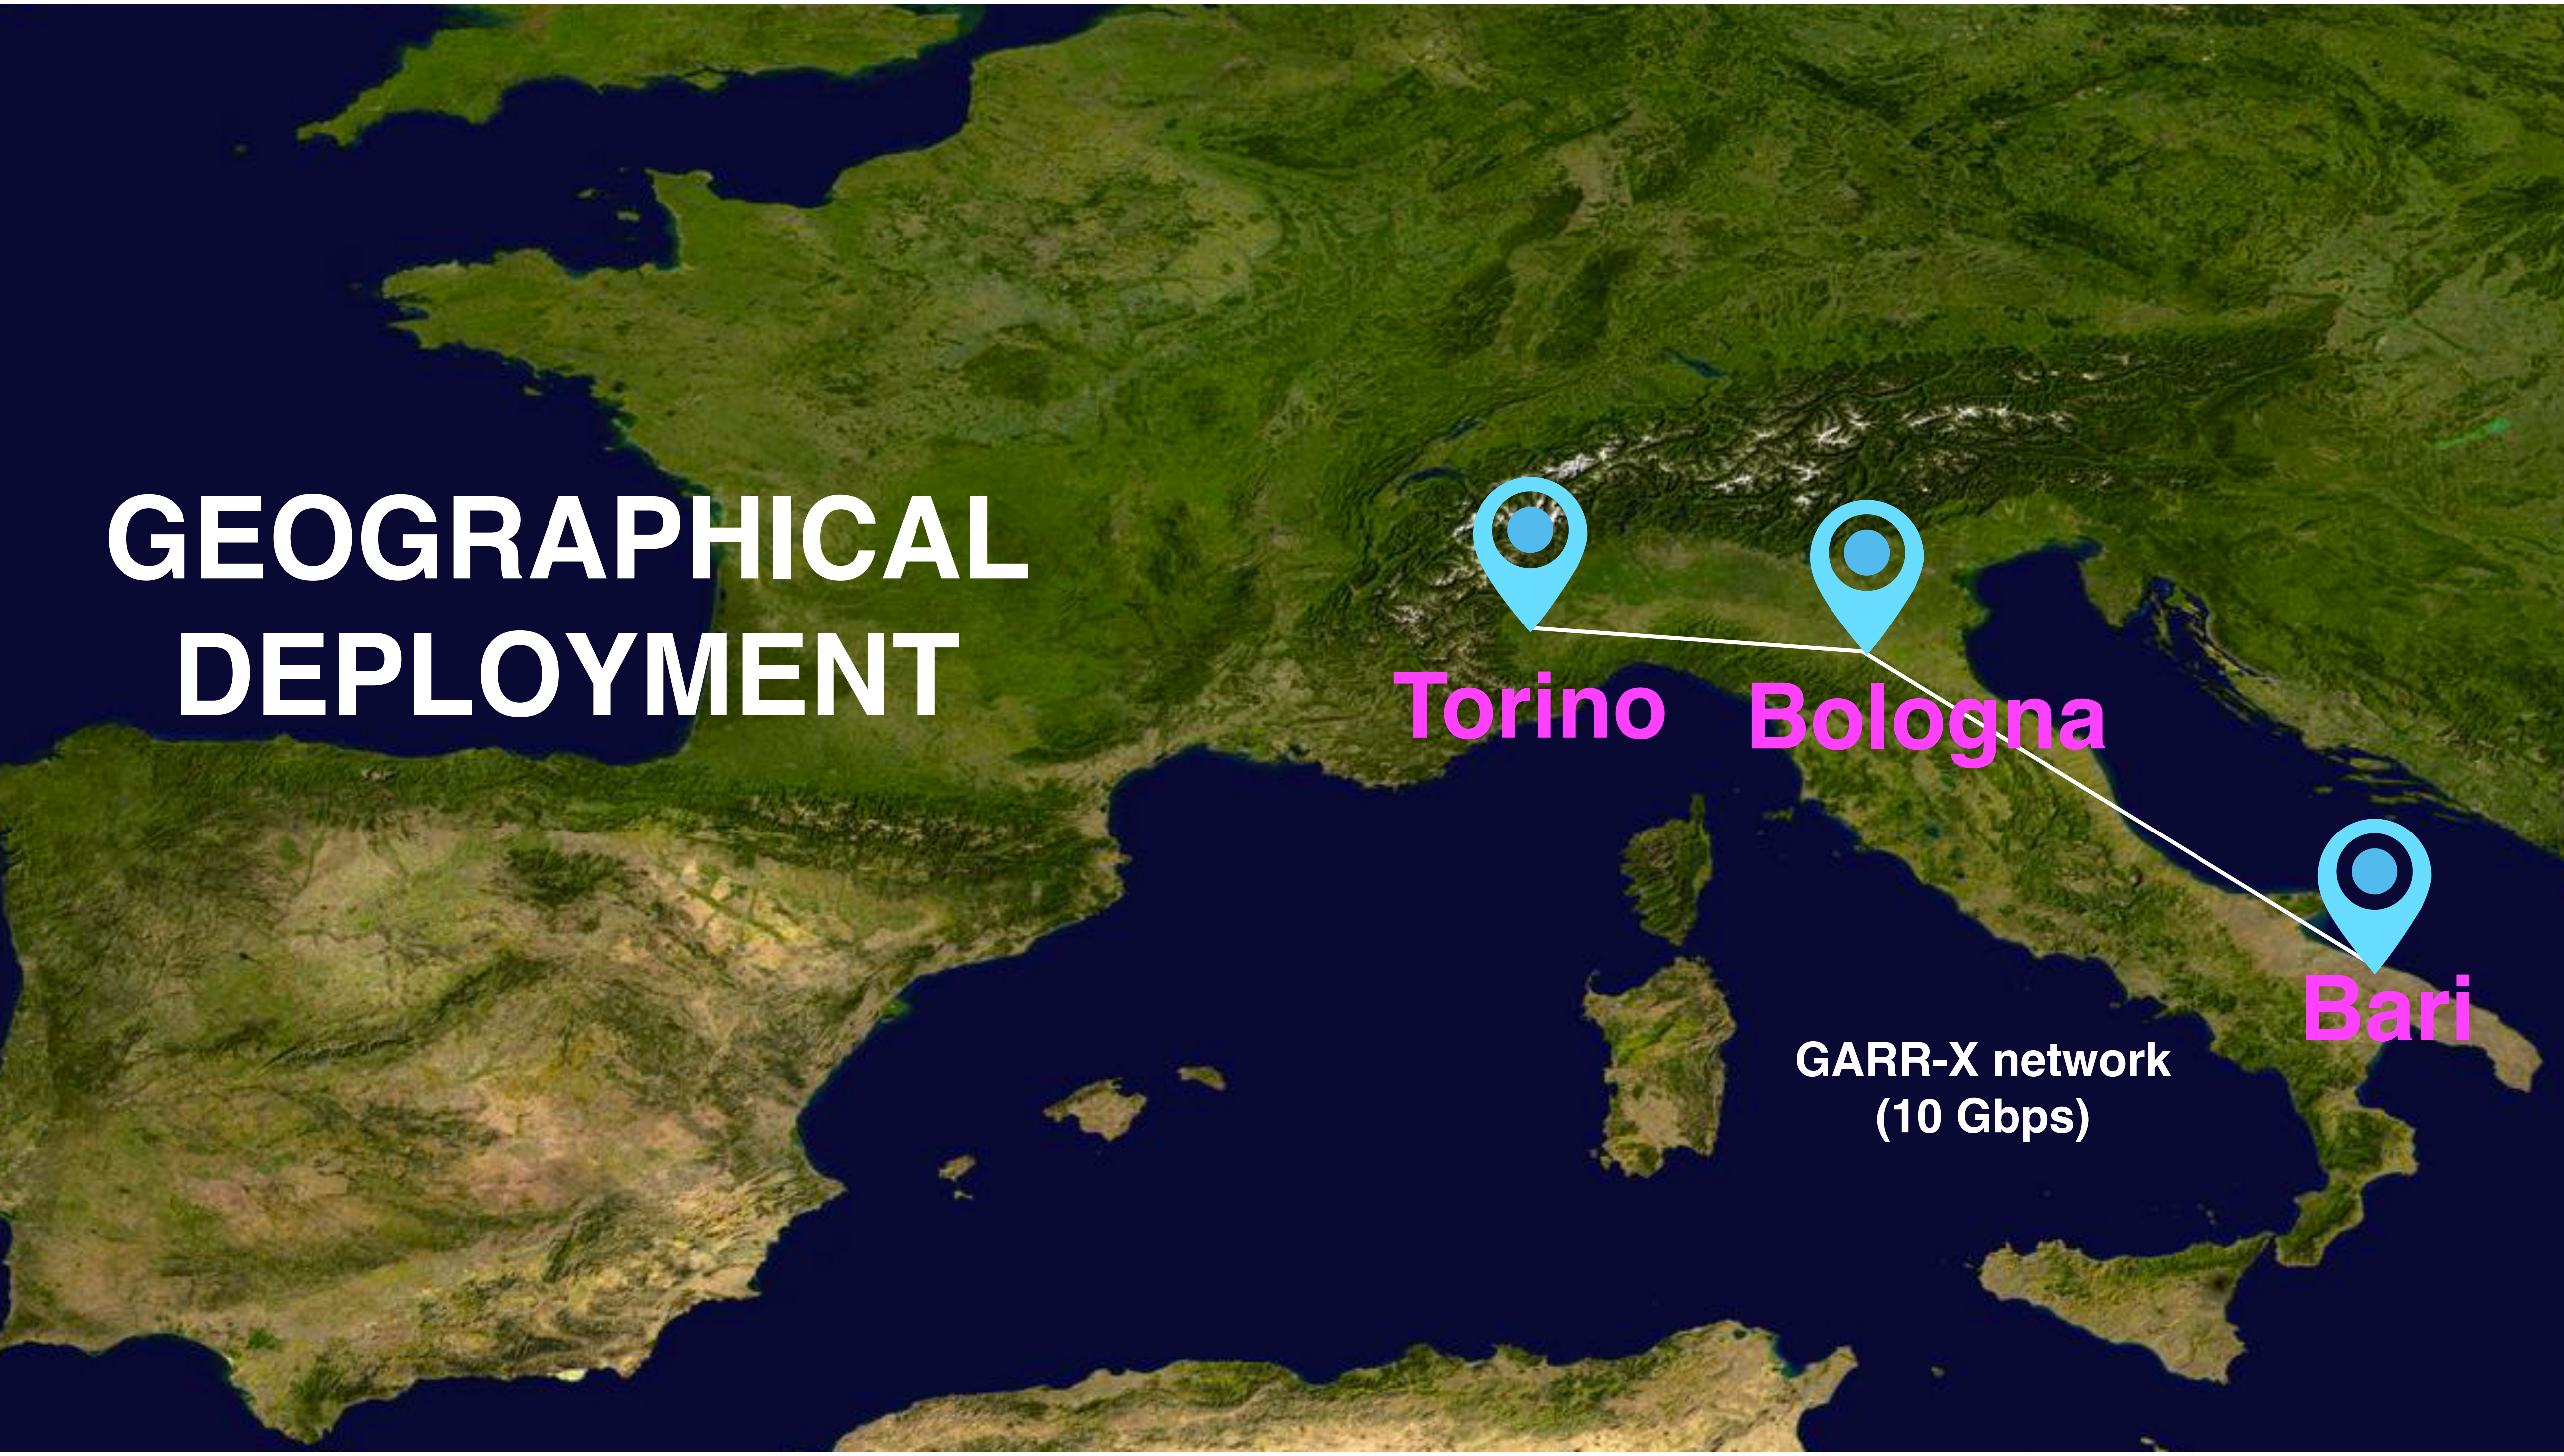

Networking

- describe network topology
- ideas on how to proceed for final version

INDIGO
DataCloud

Services

- service deployment (Ansible)
- topology

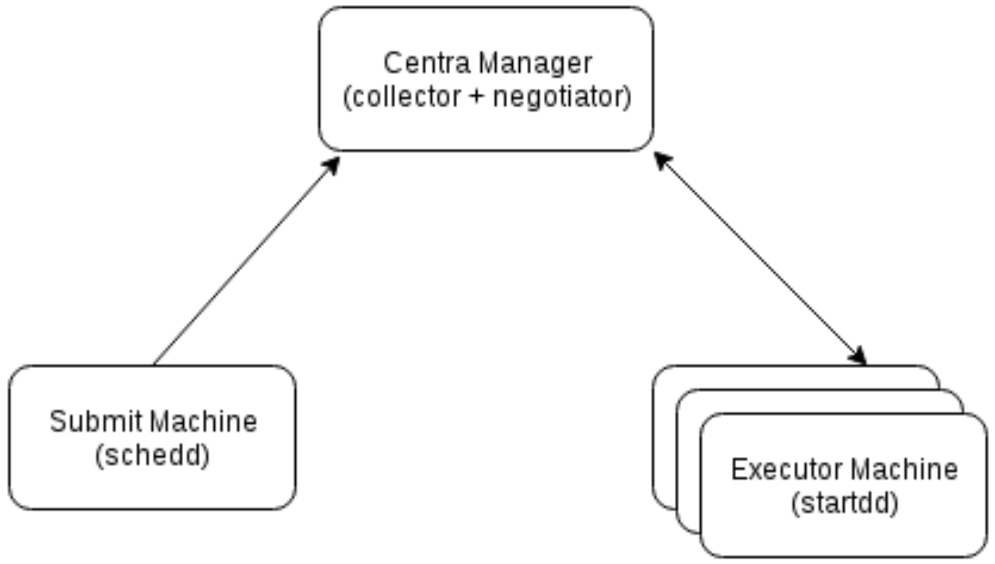

Snapshots

- snapshot Mesos, Docker containers e reti, Calico indirizzi

First results

- smoke tests
- condor benchmarking

Next

- summary cose fatte
- cose da fare
- link a repo vari
- mention service catalogue

INDIGO
DataCloud