


The CloudVeneto initiative: 10 years of operations to support interdisciplinary open science



P. Andreetto, F. Costa, A. Crescente, **F. Fanzago**, R. Lenzo, M. Sgaravatto, S. Traldi, M. Verlato, L. Zangrando (INFN - Padova)
S. Fantinel (INFN - LNL)

P. E. Mazzon (University of Padova, PNC), L. Lazzaro, M. Menguzzato, (University of Padova, DFA), G. Sella (University of Padova, DISC)

CloudVeneto is a private cloud targeted to scientific communities, based on **Openstack software**, designed in 2013 to support INFN projects, mainly HEP ones.

During the last 10 years it evolved integrating resources and use cases of ten Departments of the University of Padova. **It currently supports several scientific disciplines of different domains**, but it is also used in other scenarios (e.g. to support teaching activities)


Resources spread across two different sites, approximately 10 Km apart, connected with a **high speed optical link**.



Open Source tools and technologies allow to implement basic compute and storage services, higher level services, deployment automation, monitoring, accounting, security incident handling, etc.

Total capacity

- 66 Compute Nodes
- About 3500 Cores (in HT)
- About 16 TB of RAM
- 24 GPUs
- About 1.3 PB of Storage



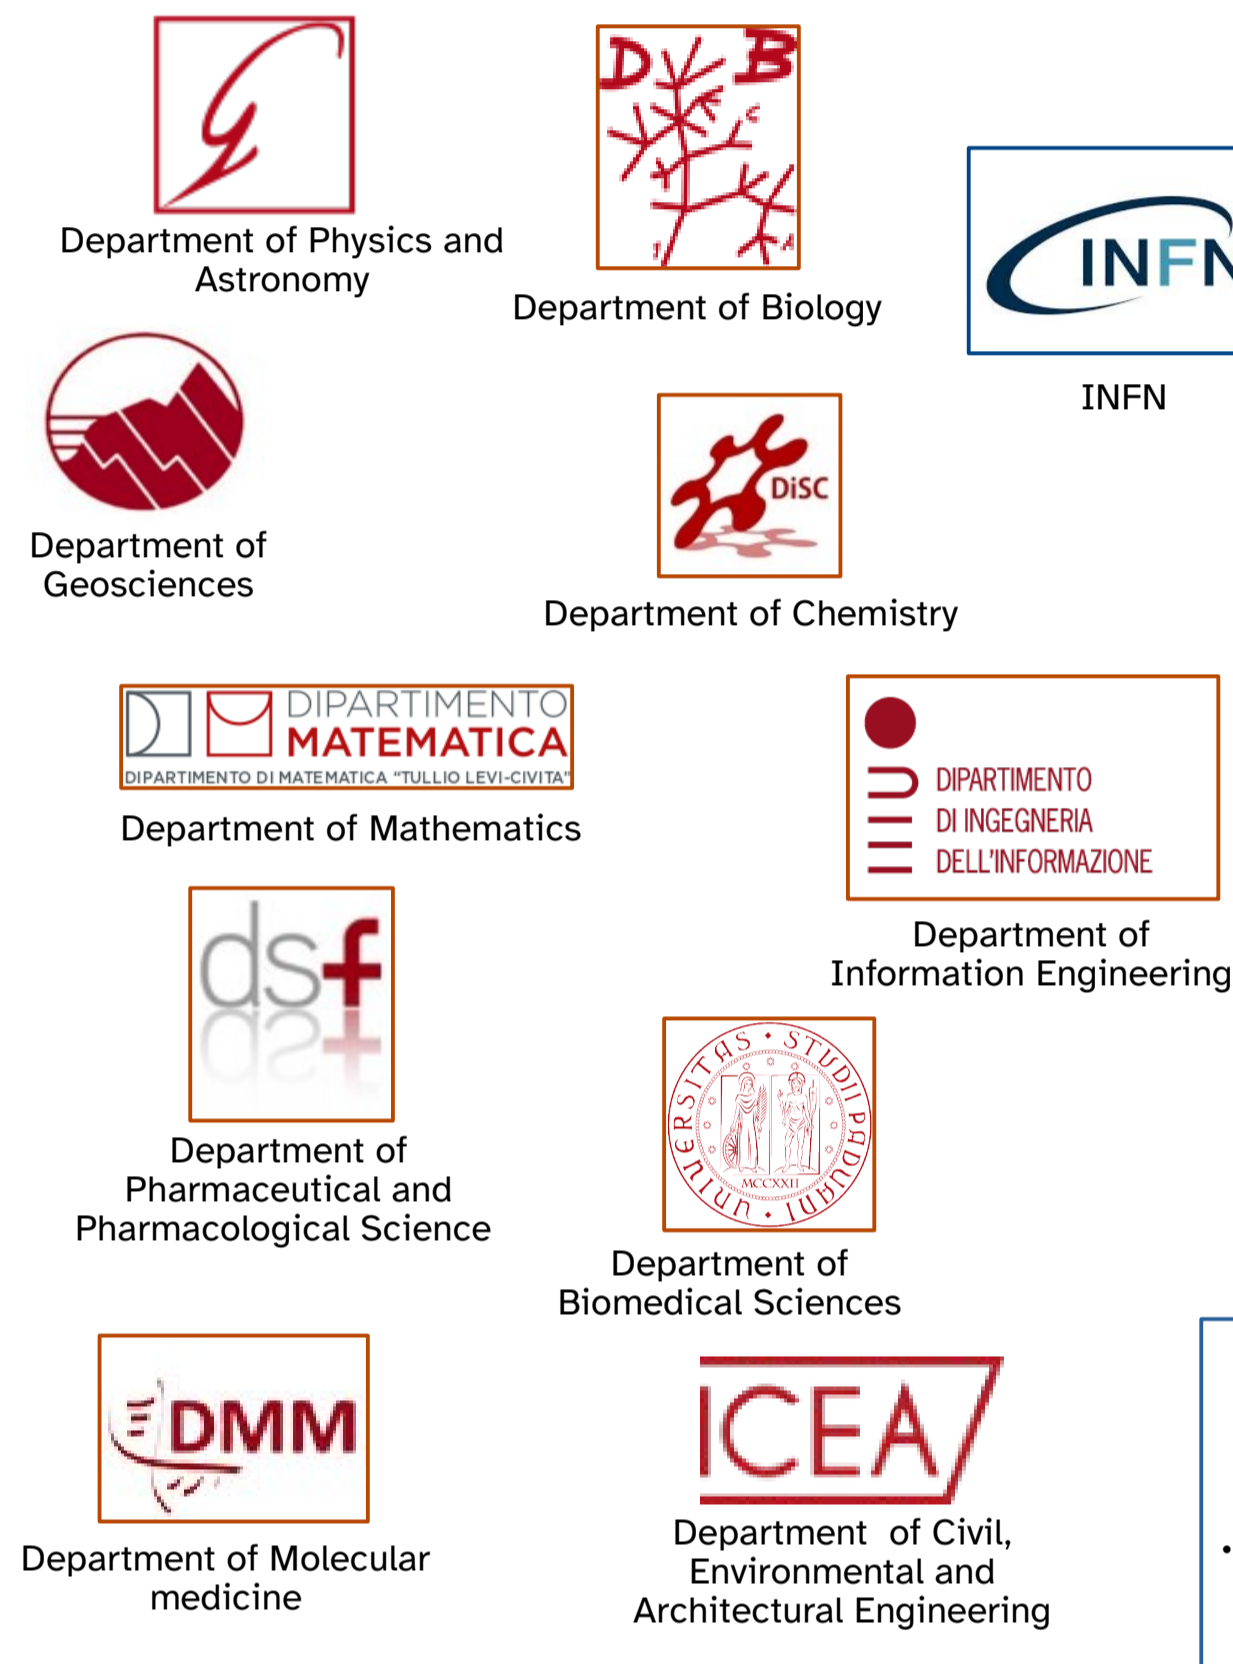
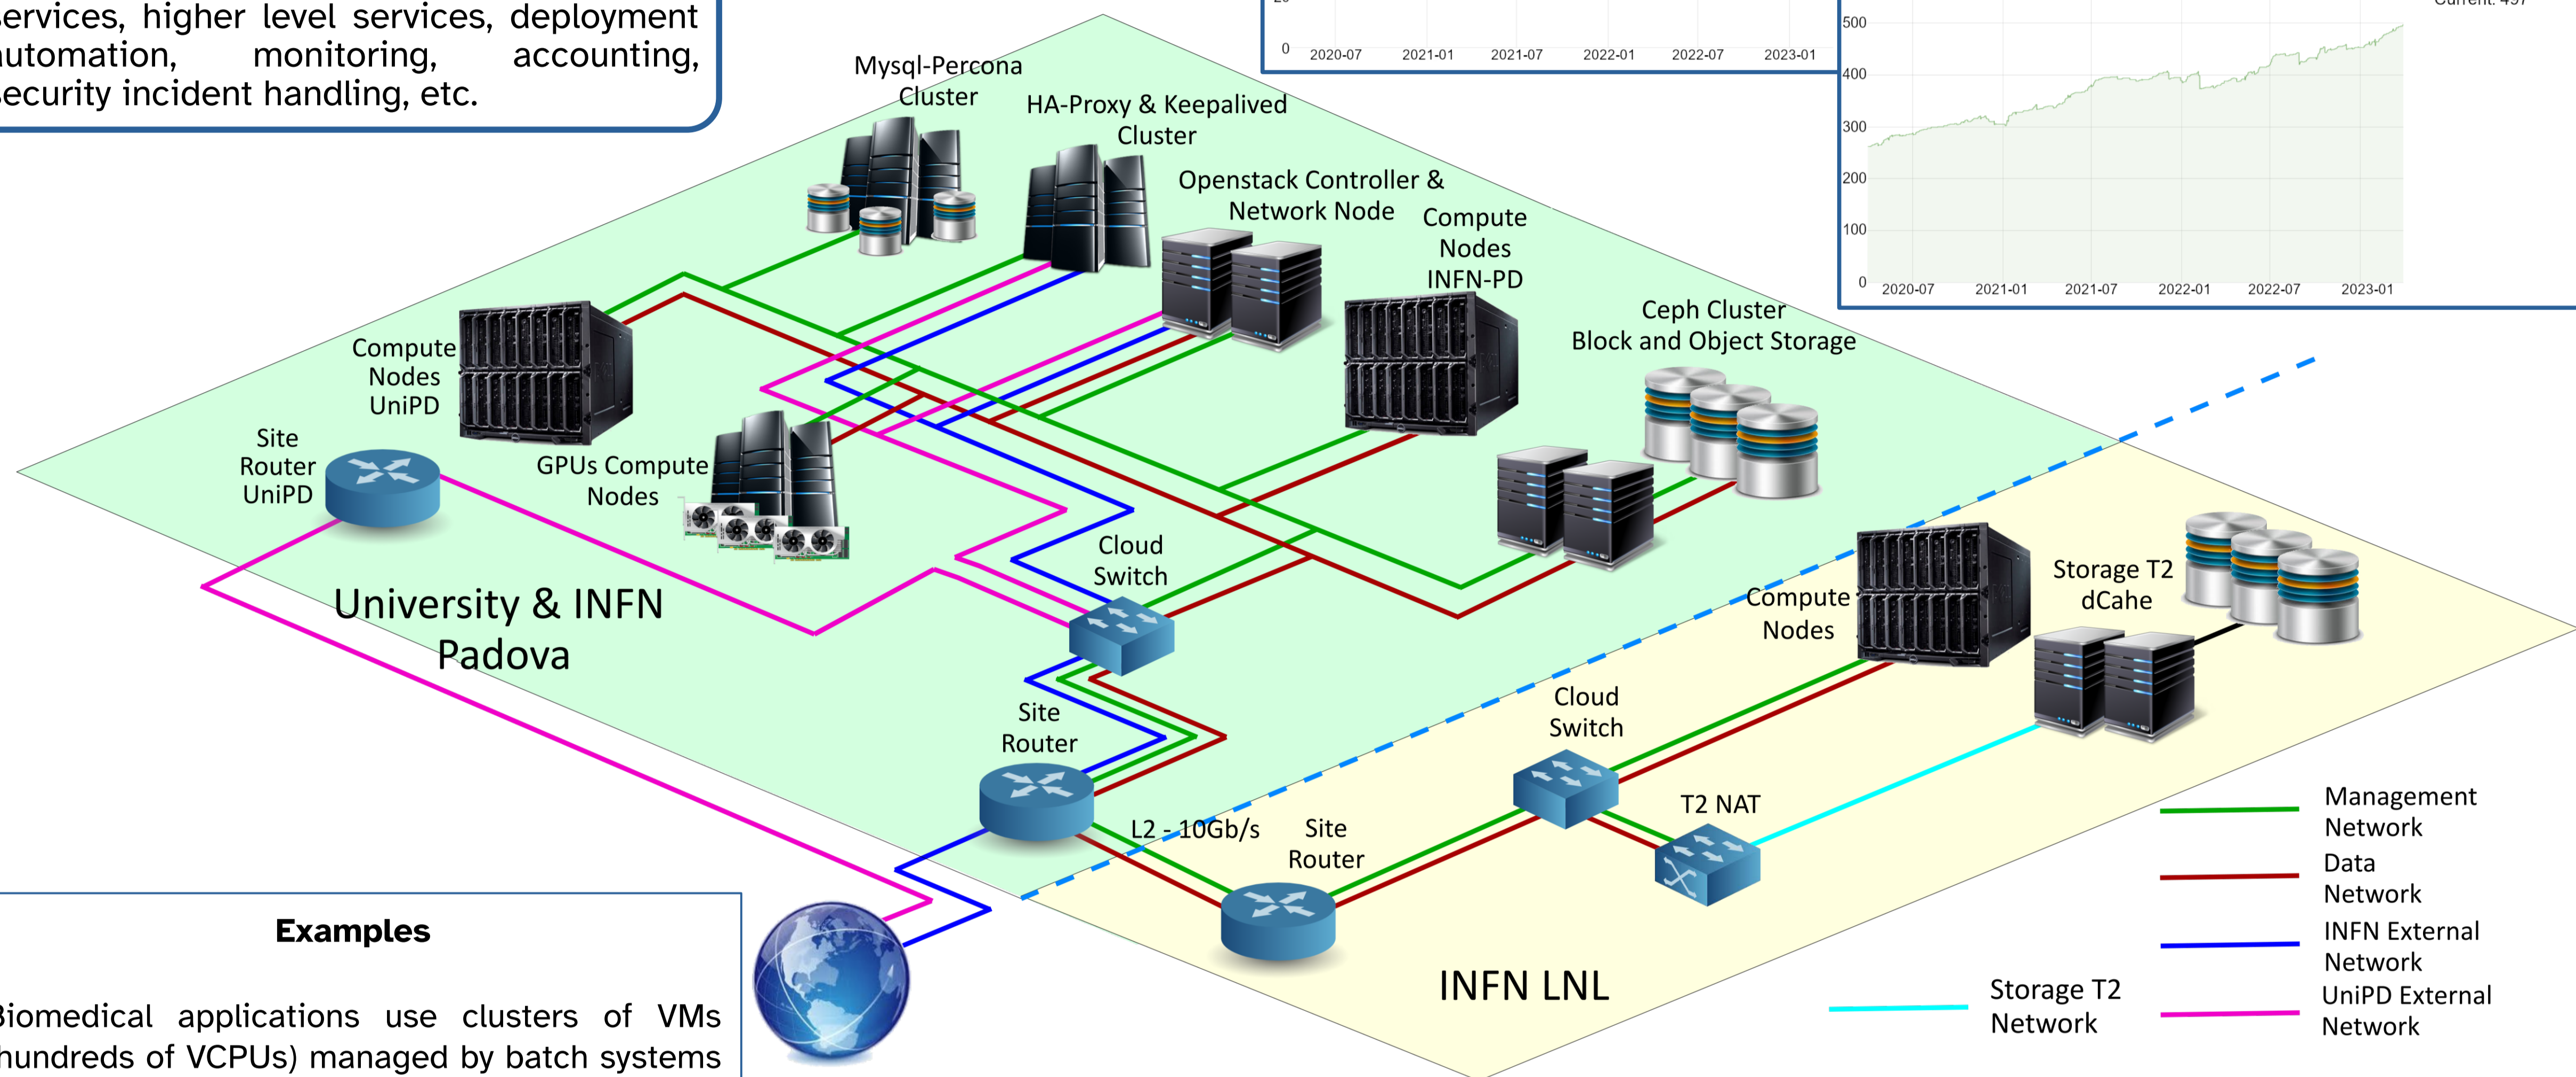


Services deployed in high availability

- Cloud services deployed on 2 controller nodes
- Database implemented by a percona XtraDB cluster
- HA-Proxy/Keepalived 3 hosts cluster to manage availability and load balancing of services
- Message system implemented by a RabbitMQ cluster composed of 3 instances

The Cloud software

- CloudVeneto is based on **Openstack software**
 - Current installed version is Yoga
- Other software and technologies are used to address some specific needs and to provide higher level services
- Integrated also some in-house developments

About 100 research projects
About 500 users



Examples

- Biomedical applications use clusters of VMs (hundreds of VCPUs) managed by batch systems (SGE, Slurm) to process data with I/O from/to NFS shared cloud block storage.
- Chemistry and Statistical Physics applications execute simulations exploiting parallel processing inside large multi-VCPUs single VMs.
- Radio Pharmaceutical and Nuclear Physics applications execute containerized simulations on the CaaS service.
- Quantum computing applications use the CaaS service for their quantum circuit simulations.

Use cases common patterns

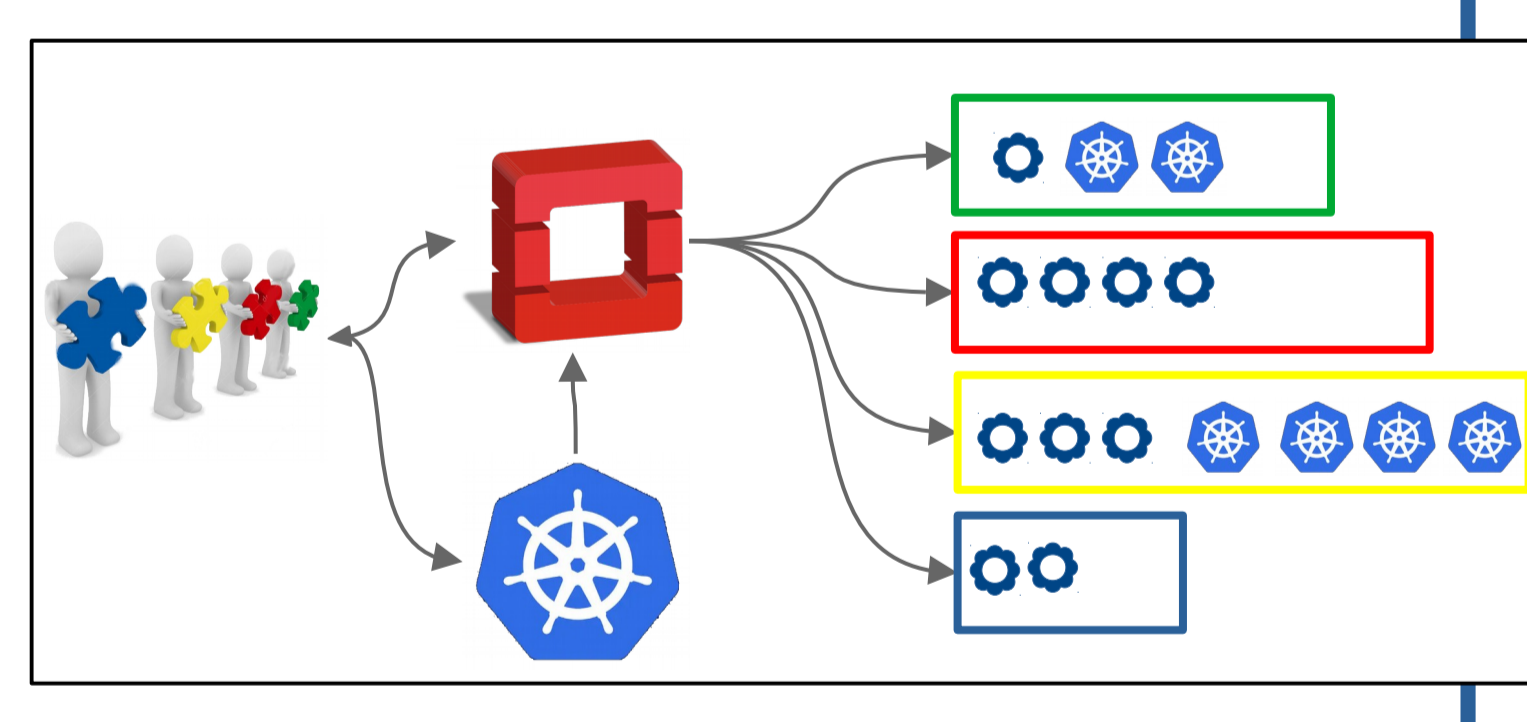
- On demand provisioning of resources.
- Orchestration of containers.
- Data-intensive applications.
- ML workflows on GPUs.

Container as a Service (CaaS)

CaaS provides a **fully managed orchestration** platform of **containers** based on Kubernetes as a **“central” cloud service**.

Single Multitenant cluster: our solution is a **centralized service, administered by CloudVeneto staff but shared among its users**

- users no longer have to create and administer their own cluster
- resource usage optimized by avoiding the proliferation of K8s clusters dedicated to individual users
- **users are grouped in namespaces** which map to the Openstack projects and access the service with their Keystone credentials.
- the **fully user isolation** is guaranteed enhancing the Kubernetes' multi-tenancy with advanced webhooks for both authN and authZ.
- **users can only operate on their own containers (pods)** and consume the quota assigned to their CloudVeneto project.
- **users create their K8s nodes via the standard kubectl command-line interface.** These nodes are CloudVeneto VMs automatically instantiated and configured by our CaaS (operator) and joined to a centralized Control Plane. They live until the owner removes them.



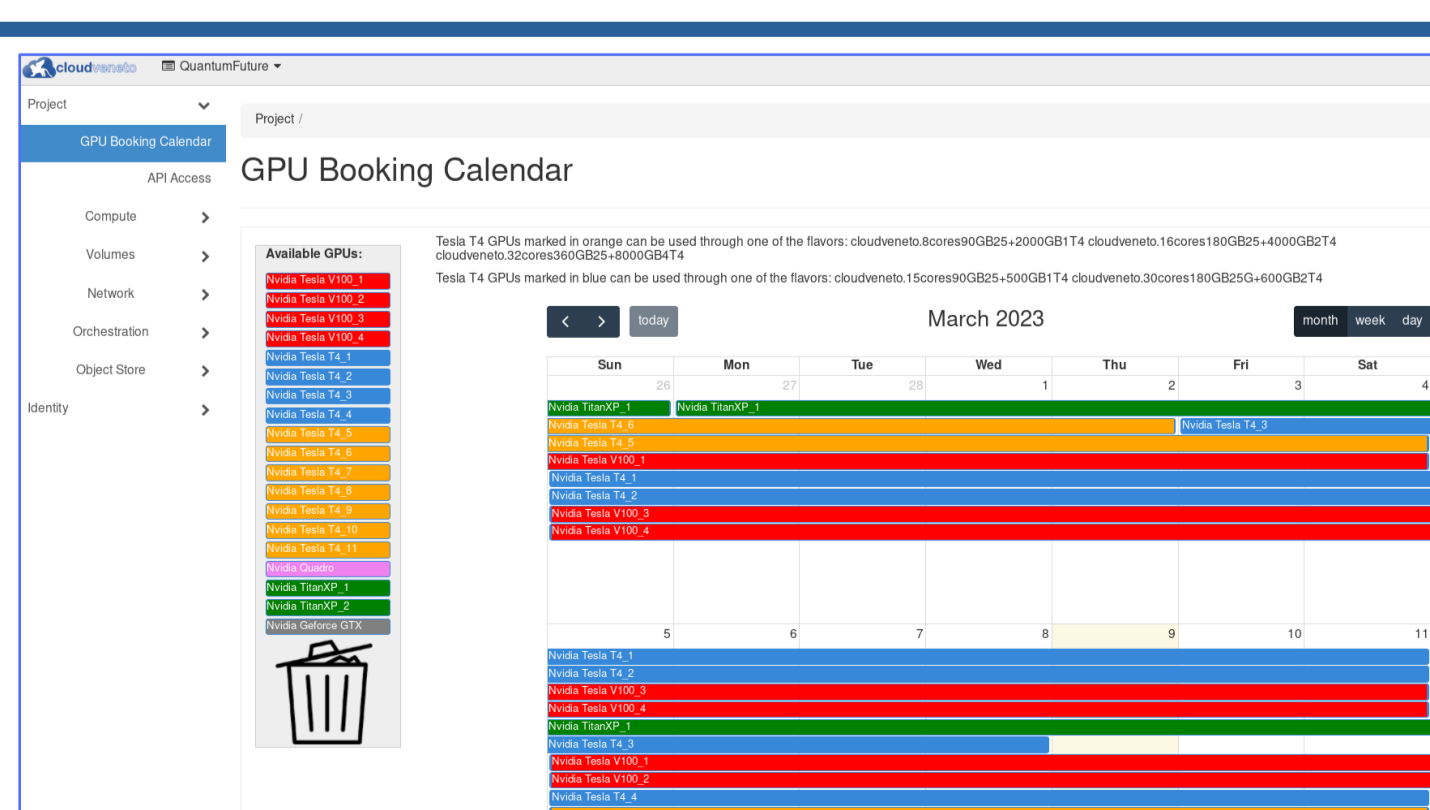
GPUs

CloudVeneto integrates **several GPU cards of different models**. They are used to address the special needs of communities whose computations involve machine learning training.

GPUs integrated in CloudVeneto

- 2 NVIDIA TITAN Xp
- 1 NVIDIA Quadro RTX 6000
- 1 NVIDIA GeForce GTX TITAN
- 12 NVIDIA Tesla T4
- 4 NVIDIA Tesla V100
- 1 NVIDIA A30
- 2 NVIDIA A2
- 1 NVIDIA RTX A4000

To manage the **allocation** of such resources among the different user communities, a **booking calendar system** has been implemented and integrated in the CloudVeneto dashboard.




CloudVeneto integrated in the INFN Cloud federation

The INFN Cloud is a federation of distributed infrastructures with the target to exploit all the computational resources available in the institute and to implement a “national Cloud infrastructure” for the research. Its services are based on modular components, covering IaaS, PaaS and SaaS and can be deployed in a simple way in the federated resources.

CloudVeneto is fully compliant with the INFN Cloud “rules of participation”, and part of CloudVeneto resources are federated in INFN Cloud.

Authentication of INFN Cloud users is implemented through the INDIGO Identity and Access Management (IAM) service. Federated users are then mapped locally to dedicated OpenStack projects.



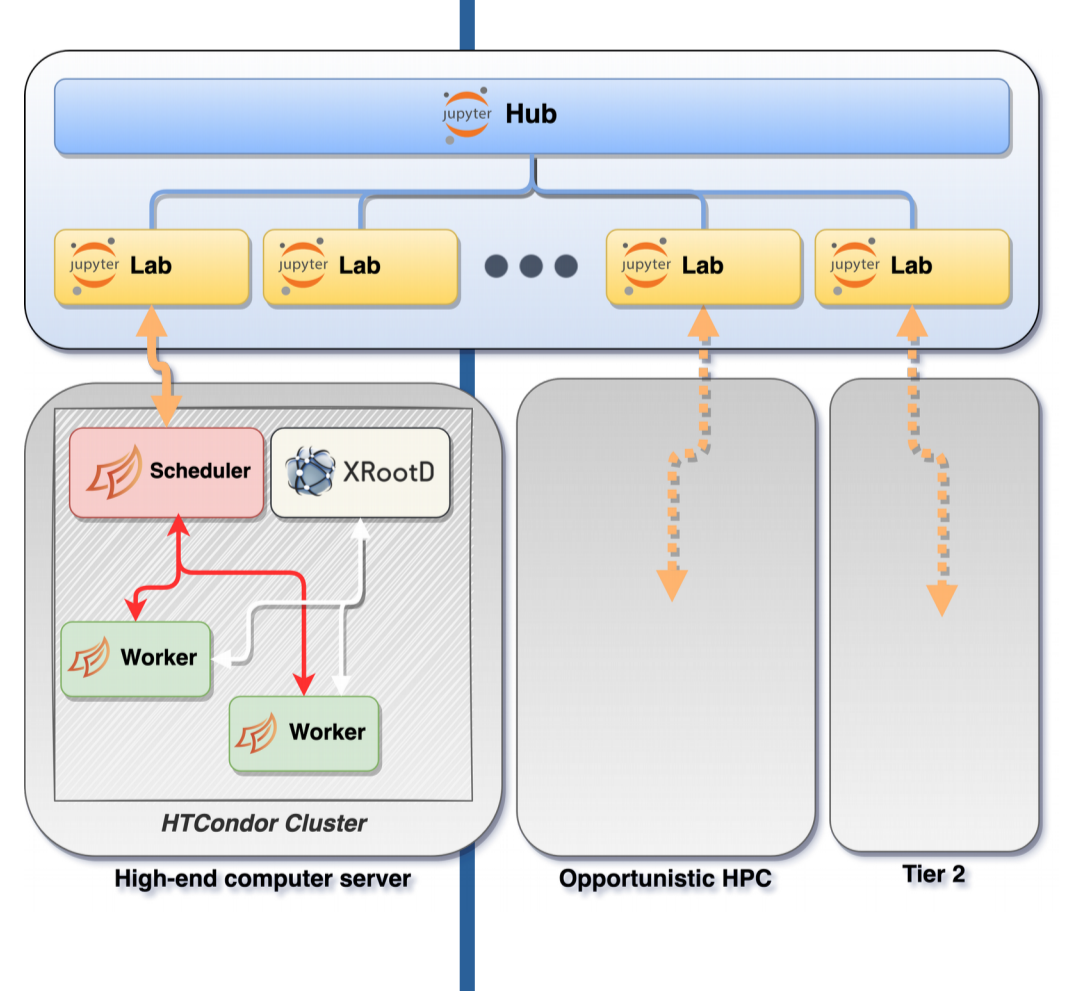
CloudVeneto and the INFN Analysis facility R&D activity

CloudVeneto is involved in the INFN Analysis Facility R&D project.

The INFN Analysis facility is conceived as a testbed for analysis that can be used in batch mode or for interactive analyses also using new frameworks.

Dask, Jupyter and HTCondor are among the enabling technologies used in this Analysis Facility.

Some resources instantiated on CloudVeneto are made available to the Analysis Facility pool.



Storage

A **unique system**, implemented through **Ceph**, is used to provide storage for different clients and for different use cases.

It provides **block storage** to:

- Image (glance) and block storage (cinder) OpenStack services
- Proxmox cluster hosting several Cloud and ancillary services

It also provides, through a highly available configuration, **object storage**, exposed through S3 and swift.

Different hardware classes and customized pool configurations (replicated pools, erasure code pools) can address multiple use cases.

A dedicated network setup was configured to allow Cloud instances to efficiently access the data of LHC experiments stored in the dCache instance of the LNL-Padova Tier2.

