
Offloading in CloudVeneto

A. Troja (antonino.troja@pd.infn.it),
F. Fanzago (federica.fanzago@pd.infn.it),
on behalf of CloudVeneto team

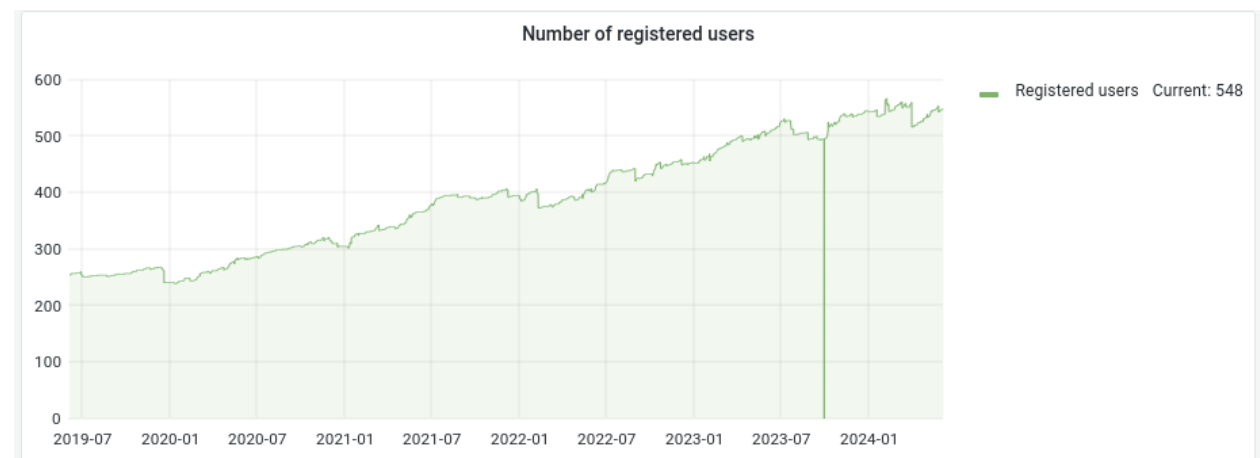
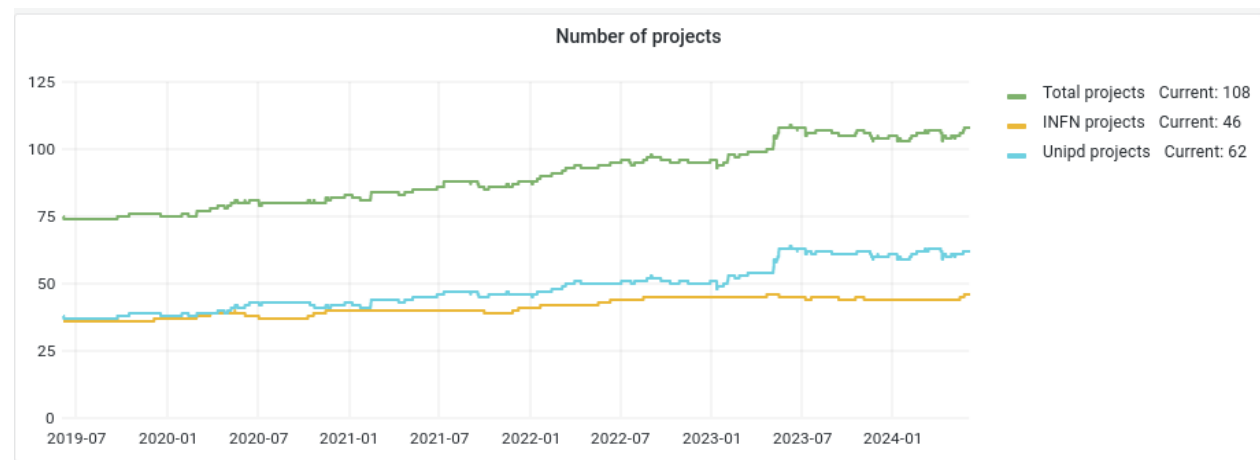
PNRR - ICSC » Spoke 2 » WP2

04/06/2024

CloudVeneto

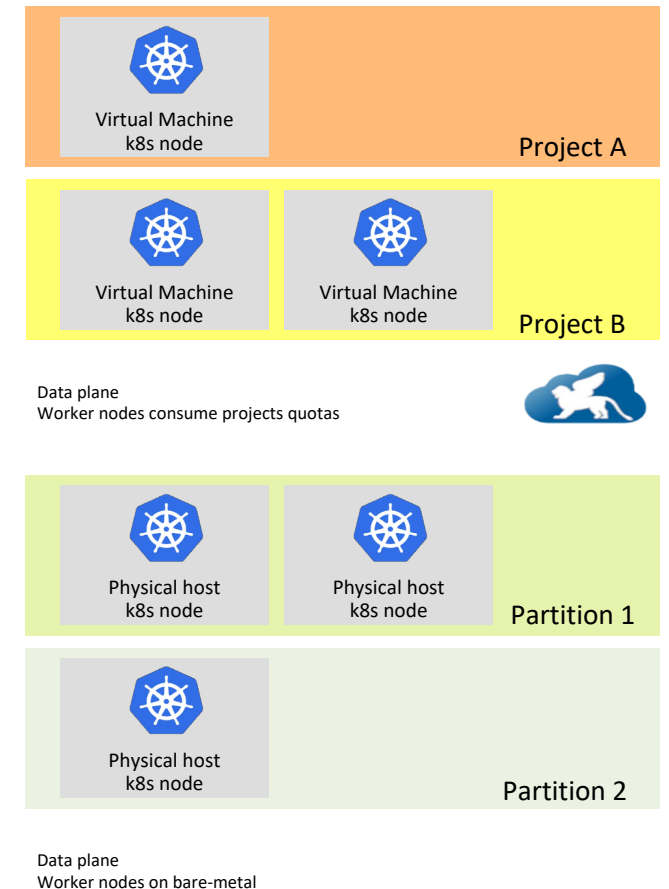
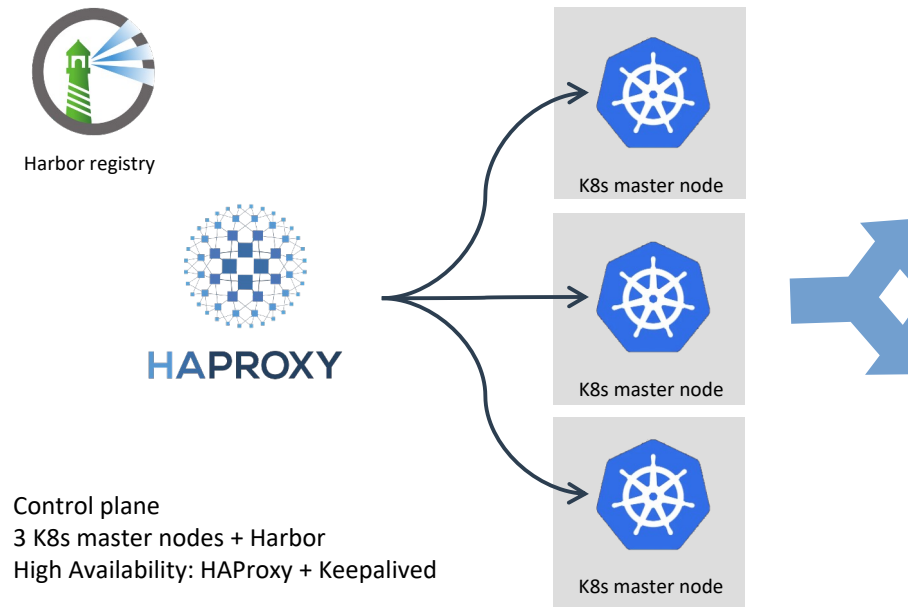
- CloudVeneto.it is an OpenStack based IaaS;
- Funded by INFN and ten departments of the University of Padova, since 2018;
- More than 500 users and 100 projects;
- Its hypervisors and storage nodes are spread across PD and LNL, with a dedicated connection;
- INFNCloud federated:
 - “catch-all” projects on private network;
 - AI_INFN.

Owner	Cores	GPUs	RAM (GB)
INFN-PD	1200	13	5792
INFN-LNL	864	0	5120
UNIPD	1152	0	6144
Other	1688	8	8352
Total	4904	25	25408



Container-as-a-Service (CaaS)

- Fully centrally managed orchestration platform in the cloud;
- Centralized Kubernetes Control Plane without worker;
- Users can create and customize nodes using resources from CloudVeneto within their tenants;
- K8s nodes (private or shared), deployed as VMs;
- Administrators can also deploy nodes on bare-metal.



Offloading on CloudVeneto

- We're testing offloading on CaaS nodes;
- Two tenants involved:
 - CMS
 - AI_INFN
- Custom sidecar plugin implemented in Python:
 - Incorporates a Flask server and translates the calls into K8s commands;
- Currently validating on CMS AF and AI_INFN FlashSim.

