



UH-iaaS

Background

- Academic community cloud i Norway based on Openstack
- Run by a team of 6 (3 USIT, 3 UiB)
- Started in 2013, in production since november 2016
- Since 2019, (single node) HPC workloads for CERNs Alice and Atlas projects
- Funded and run by UiB and UiO with additional funding from various research entities
- Open for several universities in Norway (UiT, NTNU, UiA, NMBU)

Some stats (gross numbers)

- 2 regions (OSL and BGO)
- 10 hypervisors in each site for ordinary usage (Intel-based, 48 or 56 CPUs each)
- 24 hypervisors for Alice project, 10 for Atlas and an additional 8 for general HPC (AMD-based, 128 CPUs each)
- A total 6928 cores, 34.56TB RAM, 3.3PB of Ceph block storage + 200TB object storage
- Example usage from the general use hypervisors
 - 985 VMs - 2390 vcpus used (568 available), ~ 7.5 TB RAM used (6.5 TB available)
 - 1110 projects and 839 users
 - Overcommitting and CPU scheduling is very useful for general usage, not efficient for HPC

Who use UH-IaaS

- Researchers
- Students
- IT departments
- CERNs Alice and Atlas projects

Wide array of usage spanning from HPC to teaching and web services

General terms of usage

- The UH-iaas team is responsible for the operation of the cloud infrastructure service (Openstack)
- Within the frame of this service the user is given freedom to choose the operating system and application stack
- The user is responsible for the maintenance of the operating system and application stack
- Users can apply for a project and resources, which are given by current capacity and limited by a quota
- Currently free to use within the limits of our capacity

Security

- The UH-iaaS team is responsible for the security of the service
- The user (project owner) is responsible for the security of firewalls (called security group), volumes, access and permission within the project as well as the operating system and application stacks on the instances.
- All general resources are shared between all users and projects
- Currently not supporting encrypted volumes
- Approved for “yellow” and “green” data

UH-iaaS and automation

There are several ways Openstack can be automated, for example:

- Terraform for defining and provisioning a cloud-based infrastructure
- Ansible for configuration management and orchestration on instances
- Elasticcluster for cloud-based HPC e.g. a slurm cluster

Note that Terraform needs to be configured for the UH-iaaS networking model, which is pure layer 3.

Kubernetes on Openstack

Kubernetes runs great on Openstack. There are several ways to set it up, but we recommend one of the following:

- Terraform with Ansible or kubeadm/Kubespray
- KubeNow, which is a cloud agnostic platform for microservices, based on Docker and Kubernetes. It is for the most part developed by researchers and IT personnel at Uppsala University in Sweden.

The future of UH-IaaS

- GPU, ASIC and other ML dedicated hardware with Openstack Cyborg
- More resources for HPC
- Further increased capacity for non-HPC general usage
- More Openstack services available for end-users
- Object storage ready for production