KIFÜ

A MAGYARORSZÁGI **DIGITALIZÁCIÓ** SZOLGÁLATÁBAN

Cloud for Innovation

SIG-CISS 2022. 09. 16.

MOLNÁR Péter
molnar.peter@kifu.gov.hu

KIFÜ

History

- 2015: project started, one site, educational focus
- 2019: officially prod: regions, zones, multiple volume types
- 2020: customer region at Hungarian National Library
- 2021: phasing out of blades, moving to rack servers
- Now IaaS for: eduID.hu, Edumeet, Kubernetes, MinIO, virtual classroom

Regions and zones

- Region: datacenter site
- Zone: group of rack cabinets
- Big regions contain two zones
- Small regions contain one zone

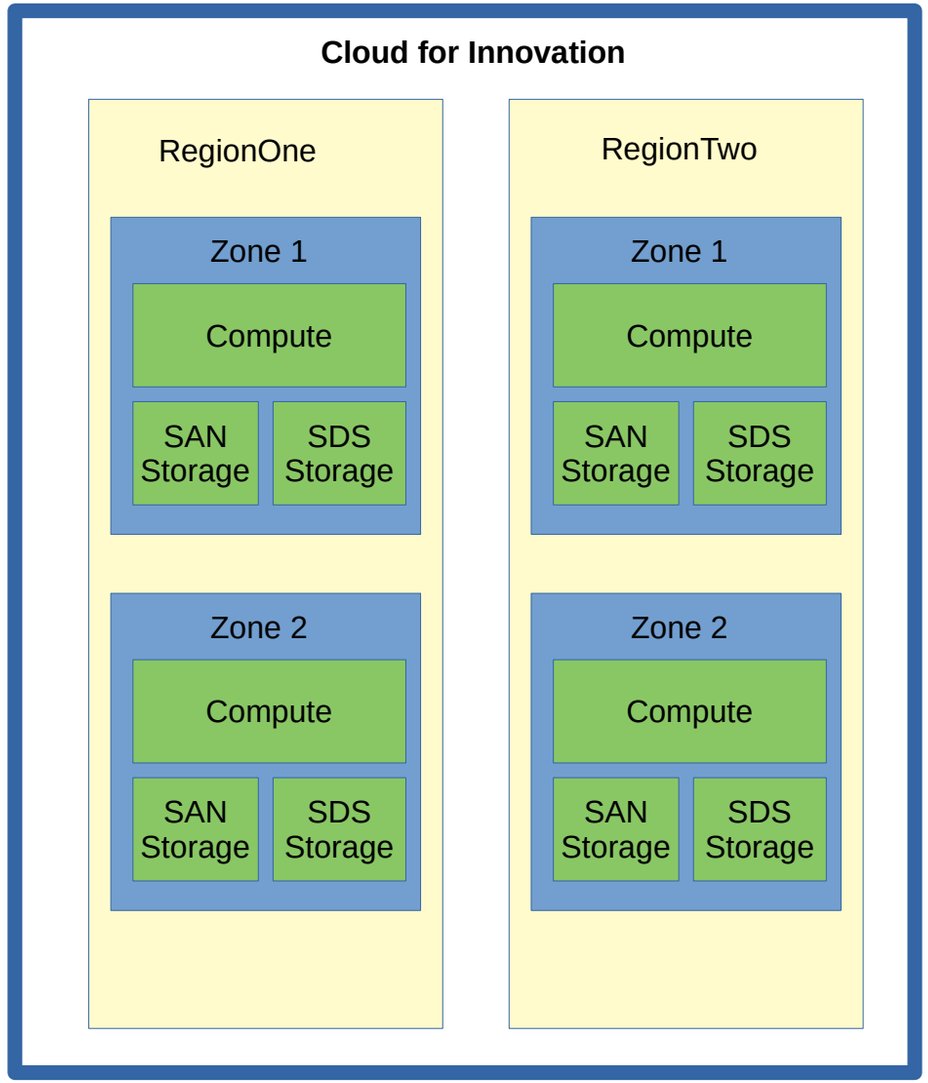

Deployment toolset

- Ubuntu server install using TFTP boot and D-I.
- OpenStack services are installed with custom-made Ansible playbooks.
  - Basic services (MariaDB, RabbitMQ, Keepalived etc.) and each OpenStack service have its own Ansible role.
  - Separate playbook variables for staging and production.
- Ceph clusters are deployed with ceph-ansible

Upgrade process

- Initial version upgrade by hand on development VMs.
- Upgrade Ansible playbook development and test.
  - Version upgrade roles are based on official OpenStack upgrade notes.
- Upgrade of the production environment with playbooks.

Monitoring

- Monitoring is done by Icinga and a lot of nrpe check.
- Logs and metrics are collected by ELK.
- Icinga and Elasticsearch are not part of the OpenStack cloud.

Accounting

- We are not using CPU hour, disk hour or network traffic based accounting.
- Users can buy static amount of vCPU/RAM/disk values. Enforcing is done by OpenStack project quotas.
- Usage statistics are generated weekly.

Project on-boarding

- Authentication and authorization is completely federation based.
- A virtual organization is need to be created by an operator in Hexaa (our federated group manager) for each OpenStack project.
- Managers of virtual organizations can invite users to the VO.
- A PHP application consume the eduPersonEntitlement SAML2 attribute released by Hexaa AA during Horizon login.
- This app creates API password and OpenStack project, assign user to the project.

Security

- BMC (iDRAC, iRMC etc.) devices and management network in a dedicated VLAN or VRF.
- Management traffic has its own MPLS VPN in the backbone network.
- Iptables firewalling.
- Access to internal networks is possible via SSH jumphosts.
- Only a very few people have admin access.

Debugging

- Metric collection in ELK: switches, SAN storages.
- Design goal: no vendor plugin in neutron.
- tcpdump can be run on compute, storage, network node.
- Design goal: dedicated hardware for dev – can be useful during debugging.

Current work

- Upgrading to Ussuri
- Launching of one big and two small regions

KIFÜ

Köszönöm a figyelmet!

www.kifu.gov.hu

molnar.peter@kifu.gov.hu

KIFÜ

A MAGYARORSZÁGI **DIGITALIZÁCIÓ** SZOLGÁLATÁBAN