KUBERNETES

Mikael Doche
Nicolas Montes

UNIL | Université de Lausanne

KUBERNETES - History

- 2018, The two technologies was neck and neck, and Docker was also more popular at the time
- Beginning of 2019, Docker Swarm was introduced at UNIL for a limited group of users/departments
- In 2020, K8S gained in popularity and was more mature, so, UNIL started digging more into it
- In 2021, due to license limitation and the freeze of features in Docker Swarm, UNIL decided to make the move to K8S
- Mid 2022, Docker instances was moved to Rancher K8S clusters
- Meanwhile, the K8S Infra team beginning redesigning K8S automatic creation workflow
- November 2023, UKS (UNIL Kubernetes Service) was officially launched for UNIL community

KUBERNETES - UKS solution

UNIL Kubernetes Service (UKS) consist of the following

- Full stack Kubernetes cluster using Rancher API calls
- Cluster nodes up and down scaling support
- CSI driver with CEPH Storage support
- Image Registry using Jfrog Artifactory
- Backup solution with Velero and Cohesity S3 buckets
- Load balancing support using F5 solution
- Firewalling solution using custom automation processes

KUBERNETES - UKS solution

UNIL provides Kubernetes cluster solution for its community members

Any experienced user/department can create a cluster

Clusters are isolated from the others and are spread into four environments

- Sandbox (AKA BAS)
- Development
- Testing (UAT like)
- Production

Users can manage their clusters as they want and are responsible for the management of various applications running in a cluster

IT department are responsible for maintaining the clusters up to date and for the underling part that compose a cluster, such as VMWare, CEPH Storage, F5 Load balancer, etc.

KUBERNETES - security

- Kubernetes security relies on firewall rules, K8S network policies, Pod Security Admission, rBAC and custom image scanning process using XRay for Artifactory.
- Clusters could be exposed to Internet via Ingresses that are behind firewall and the F5 Load balancer, which limit access, which prevent DOS, DDOS and other network attacks.
- K8S components are automatically updated when a minor releases is available.
- OS Update occur on daily basis.

KUBERNETES - used technologies

Our Kubernetes clusters relies on those technologies

- Vmware/ARIA (vRealize)
- RedHat
- AWX
- Artifactory
- Docker/RKE2
- Rancher
- F5
- Cohesity
- Ceph

KUBERNETES - Rancher

RANCHER BY SUSE

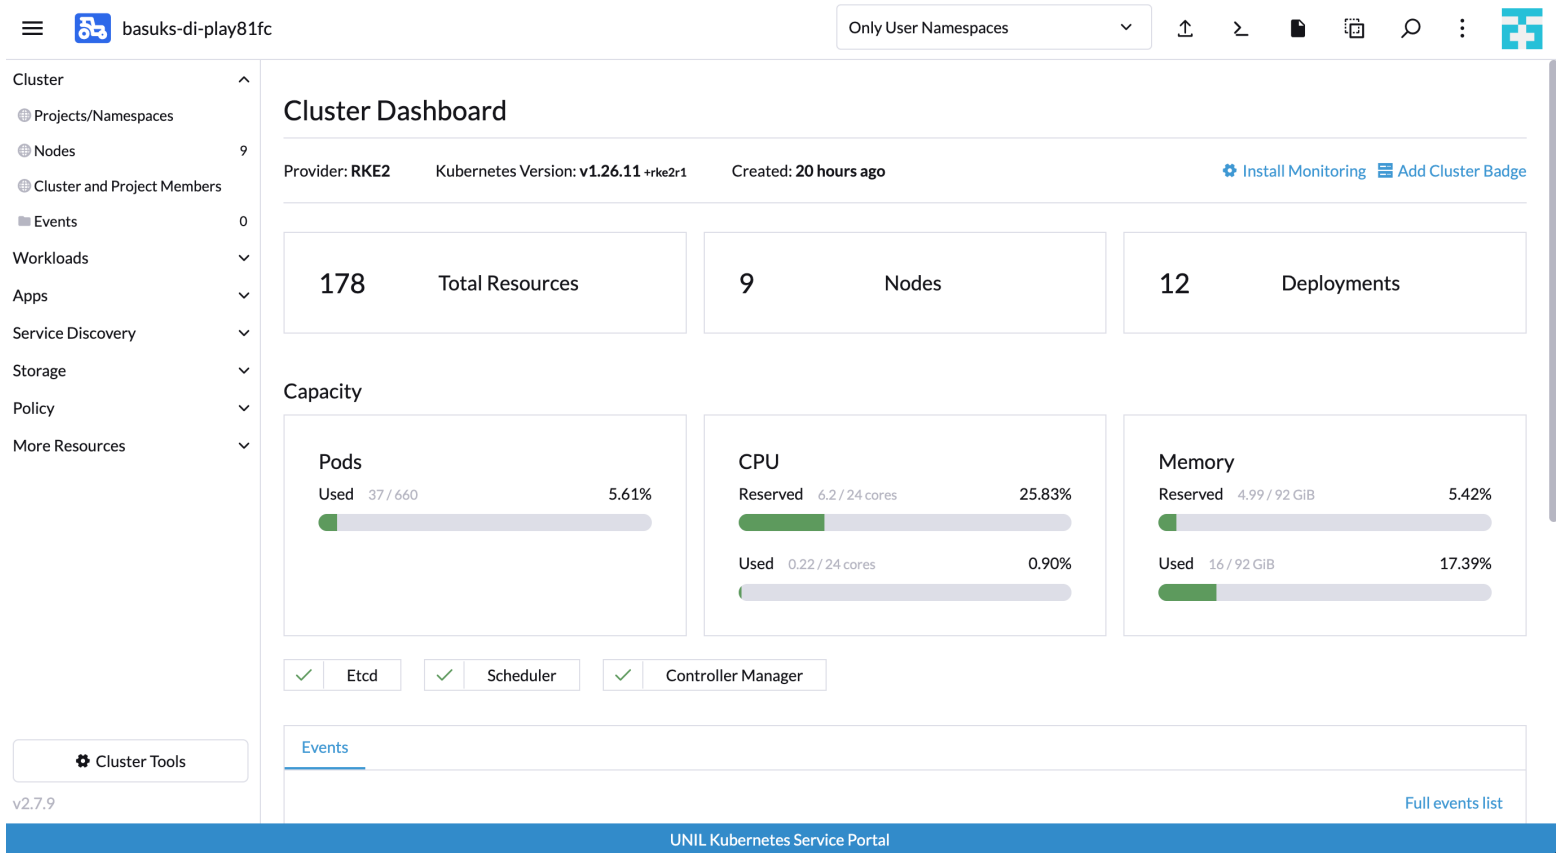

UNIL | Université de Lausanne

Kubernetes - Rancher

- Kubernetes web GUI
- Central Multi-cluster management
- Allow easy ACL management including LDAP support
- Paid Suse Rancher support available
- Easy/One click cluster upgrade
- Public clouds integration
  - Openstack
  - Vmware
  - AWS
  - Etc...
- Third-party tools integration
  - Prometheus/Grafana
  - Istio
  - OPA Gatekeeper
  - NeuVector
  - Etc…
- Rancher continuous delivery that automatically deploy helm chart on managed clusters
  - ceph-csi-cephfs
  - Ingress NGINX or Traefik
  - Velero
  - Etc.

KUBERNETES - CREATION workflow

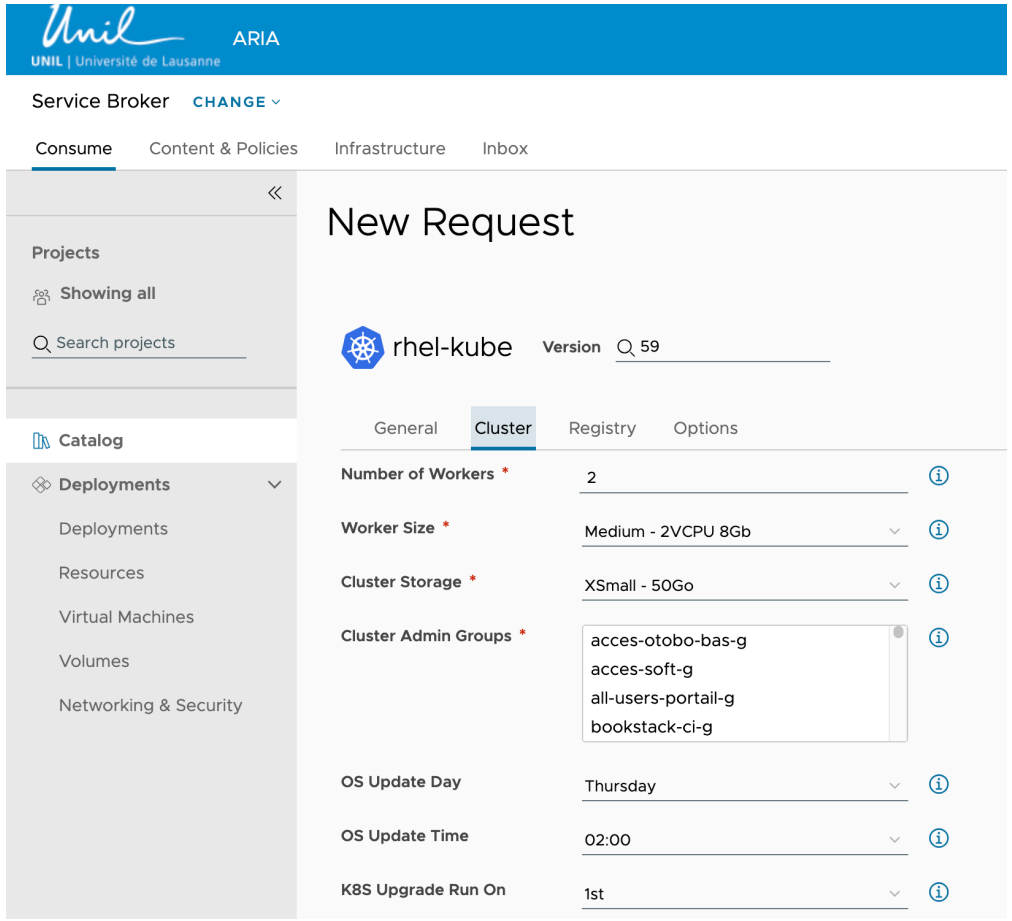

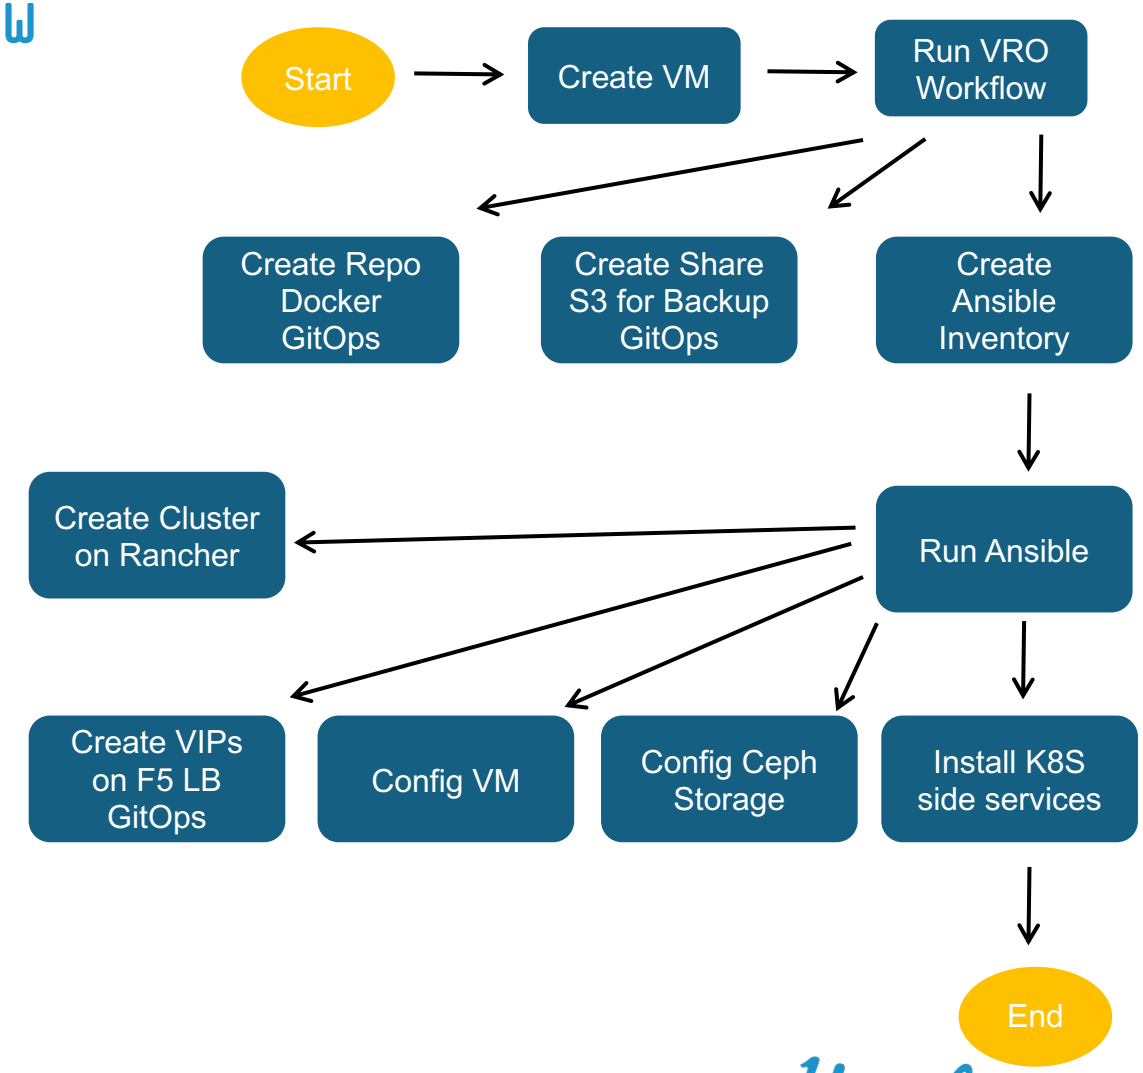

Kubernetes - architecture - network

- Ingress traffic are routed through out the F5 Load balancer
- F5 Load balancer send traffic to Kubernetes nodes which could run Nginx or Traefik local ingress system
- Client have two VIPs
  - External VIP which could be accessible from the outside
  - Internal VIP only accessible inside UNIL network
- Currently only HTTP and HTTPS traffic is allowed
- SSL Certificates are automatically generated on the F5 Load balancer

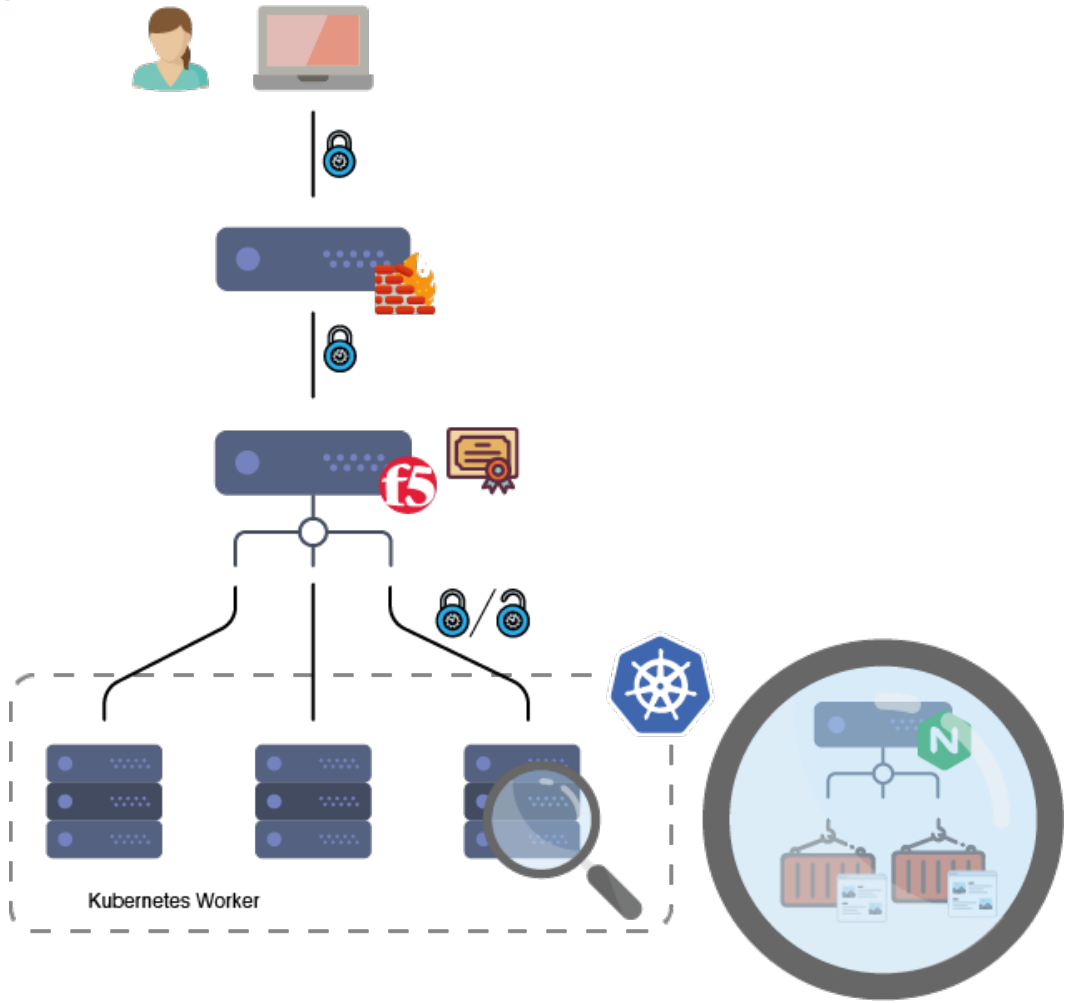

Kubernetes - architecture - rancher

- Dedicated K8S cluster for Rancher Server running on top of RKE2 (Rancher Kubernetes Engine v2)
- About 20 clients K8S clusters manager by Rancher Server
- Each client cluster running on top of RKE2 (Rancher Kubernetes Engine v2)
- Client's cluster contacts Rancher Server via the F5 Load Balancer (portal.uks.unil.ch)

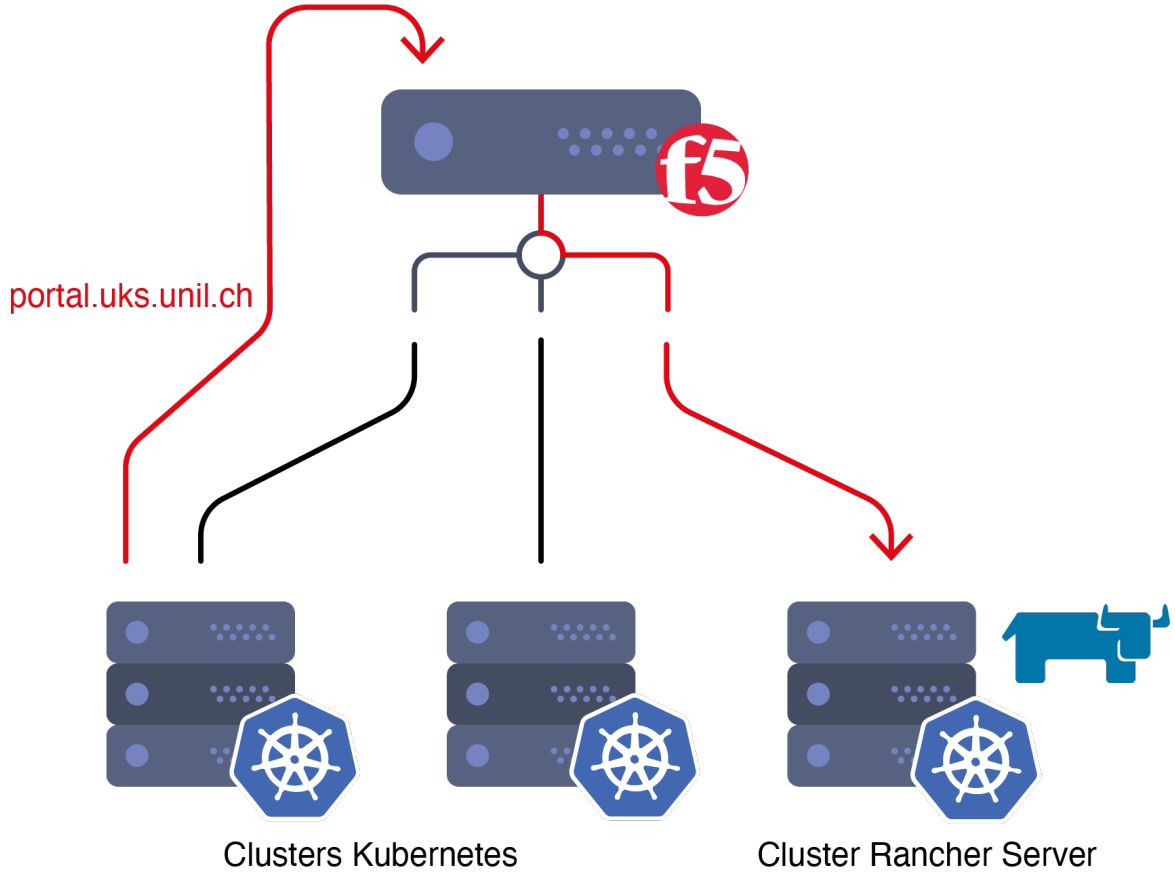

Kubernetes - architecture - storage

- Dedicated CEPH cluster is used as for the Kubernetes backend storage
- CEPH cluster is split across 3 datacenters
- Each DC, has 2 storage nodes (OSD) and 1 monitor
- The manager is a virtual machine that relies on the VXRail VMWare infrastructure
- Each K8S cluster has a dedicated cephfs subvolume
- 104TB available storage / 3 (x3 replicas) = 34TB usable for Kubernetes clusters

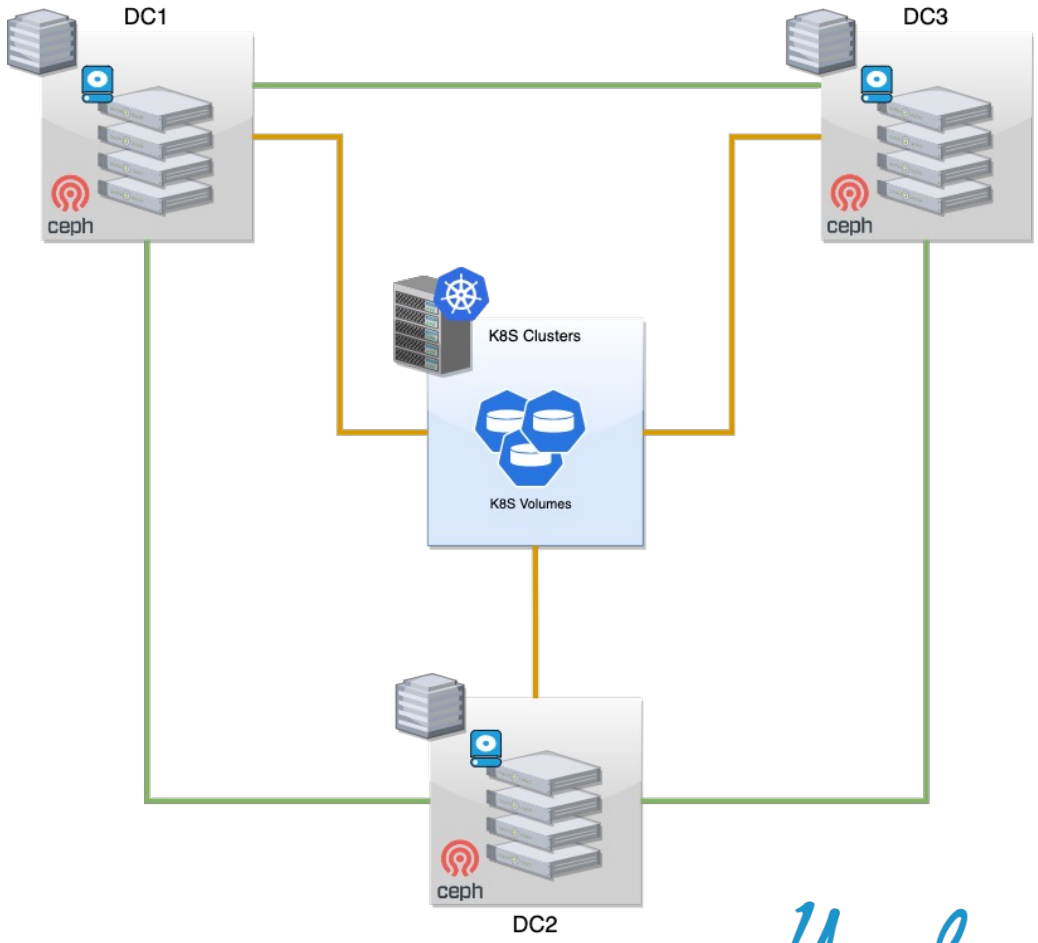

Kubernetes - Architecture - image registry

- Image registry relies on Artifactory from JFrog
- Each cluster has a dedicated cluster image registry
- Cluster registry has two defaults credentials
  - Read Write (DevOps)
  - Read Only (Cluster Deployment)
- User can access the cluster image registry with their user/group LDAP account
- The cluster image registry is automatically linked with the on demand K8S cluster
- Artifactory and cluster image registries are behind F5 Load balancer

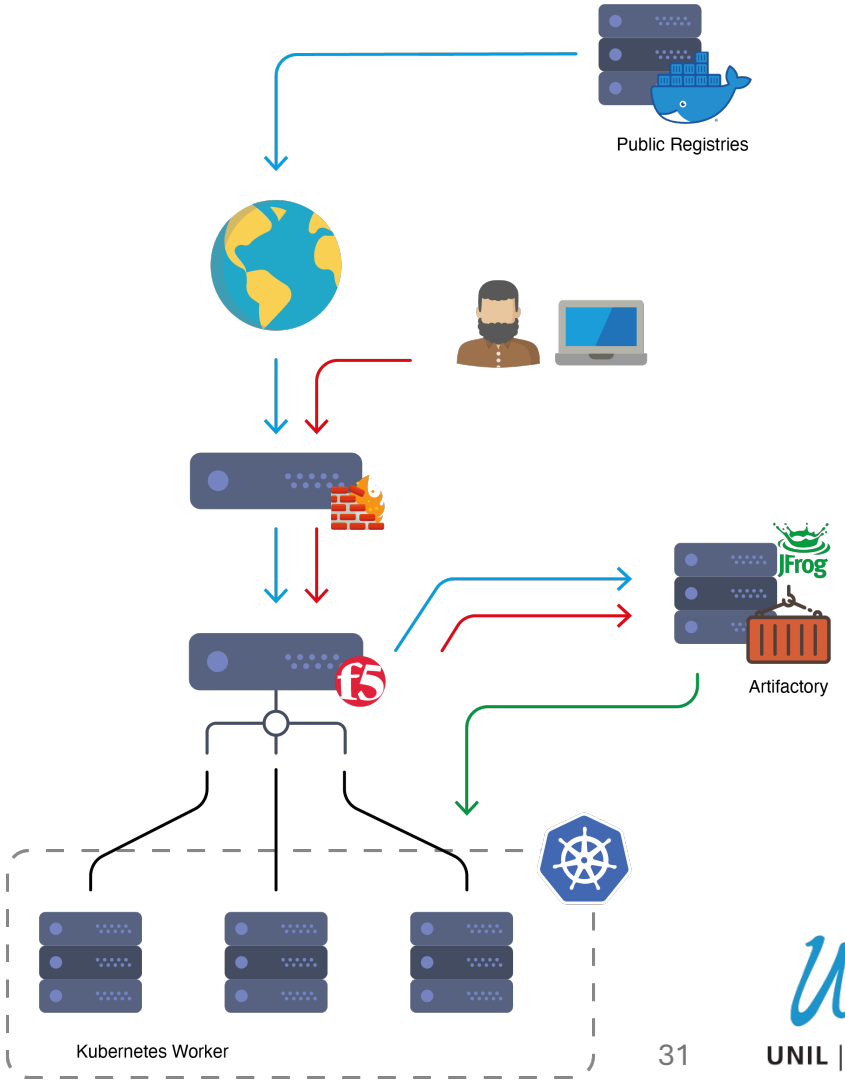

question

Thank you