

*Unil*

UNIL | Université de Lausanne

CERN - UNIL

18.06.2026



CENTRE INFORMATIQUE

# AGENDA

1. NETWORK

2. VMWARE / ARIA


3. KUBERNETES

4. HPC (DCSR)



# NETWORK

Grégory Moix



## NETWORK @ unil

- Serving 21k + people (incl 17k students)
- Campus
  - 40 buildings
  - 600 switches
  - 800 wifi antennas
- Datacenters
  - 3 sites on campus
  - from a network perspective = 1 logical DC

## DC technological stack

- Nvidia Cumulus Linux
- Palo Alto
- F5 BigIP
- Prometheus/Grafana + ElasticSearch/Kibana

# datacenter fabric

- Leaf-Spine Topology

- with 40Gbs uplinks (soon to be 100Gbs)
- around 20 switches


- VXLAN + BGP EVPN

- Overlay/Underlay
- Allows us to provide extended vlans between our DCs
- ECMP loadbalancing + redundancy

Spine

Leaf

Servers



# nvidia cumulus linux

- In production for 6+ years
- Managed like a Linux system
  - Really easy to automate
  - Generate new config + `systemctl reload ...`
  - Linux best practice and tools are available (syslog, grep,...)

# automation philosophy

- Goal:
  - Shift from repetitive to high value work
  - Do more with same resources
- How:
  - Scale-able architecture
  - Standardize services as much as possible
  - Ease of automation = requirements for new products



## network as code

- Everything in git (code + issue tracking)
  - Mandatory process of PR + review before any change
- Deploying with ansible
- Describe desired target state as data structures in yaml
- Self-Service: will modify those data structures as well

question



question

Thank you

The logo for the University of Lausanne (Unil), featuring the word "Unil" in a blue, cursive script font.



VMWARE

Arnaud Burkhalter



*Unil*

# VIRTUAL MACHINE INFRASTRUCTURE


- **Mission:** Provide infrastructure to host VMs for the faculties and administrative purpose. Don't include VM for research projects which are hosted in dedicated infras (DCSR)
- Hyperconverged Infrastructure VxRail with vSAN
- Stretch clusters

- **Main cluster - VxRail / vSAN Cluster**

- Hosting all non-research UNIL VMs
- VMware ESXi
- 16 Nodes (2 sites x 8 nodes)
- Cluster Capacity
  - Storage : 670 TB full flash vSAN
  - Memory : 28 TB RAM
  - CPU: 1824 cores
- 1200 Virtual Machines
- Linux & Windows Servers

- **SAP - VxRail / vSAN cluster**

- Dedicated to SAP applications VMs
- VMware ESXi
- 6 Nodes (2 sites x 3 nodes)
- Cluster Capacity
  - Storage : 84 TB full flash vSAN
  - Memory : 12 TB RAM
  - CPU: 384 cores
- 20 VMs
- Linux Servers



# VIRTUAL MACHINE INFRASTRUCTURE

- Some other smaller standard vSphere clusters.
- No VxRail, No vSAN, non-stretched

- **DAS01\_test – Non-stretched**

- Hosting infra test VMs
- VMware ESXi
- 2 Nodes
- Cluster Capacity
  - Storage : 7 TB
  - Memory : 1 TB RAM
  - CPU: 80 cores
- 37 Virtual Machines
- Linux & Windows Servers

- **DAS02\_phone – Non-stretched**

- Hosting Avaya phone application VMs
- VMware ESXi
- 2 Nodes
- Cluster Capacity
  - Storage : 5 TB
  - Memory : 255 GB RAM
  - CPU: 40 cores
- 11 VMs
- Linux appliances

- **MAX**

- Hosting Witness VMs
- VMware ESXi
- 1 Node
- Cluster Capacity
  - Storage: 21 TB
  - Memory : 1.3 TB RAM
  - CPU: 64 cores
- 8 VMs
- Host Linux & Windows Servers

# VM CREATION - PORTAL




UNIL provide a [portal to deploy VM](#) in a fully automated way. The faculty members can create their VM without needed contacting IT.

Based on product VMware ARIA Automation (previously vRealize Automation)

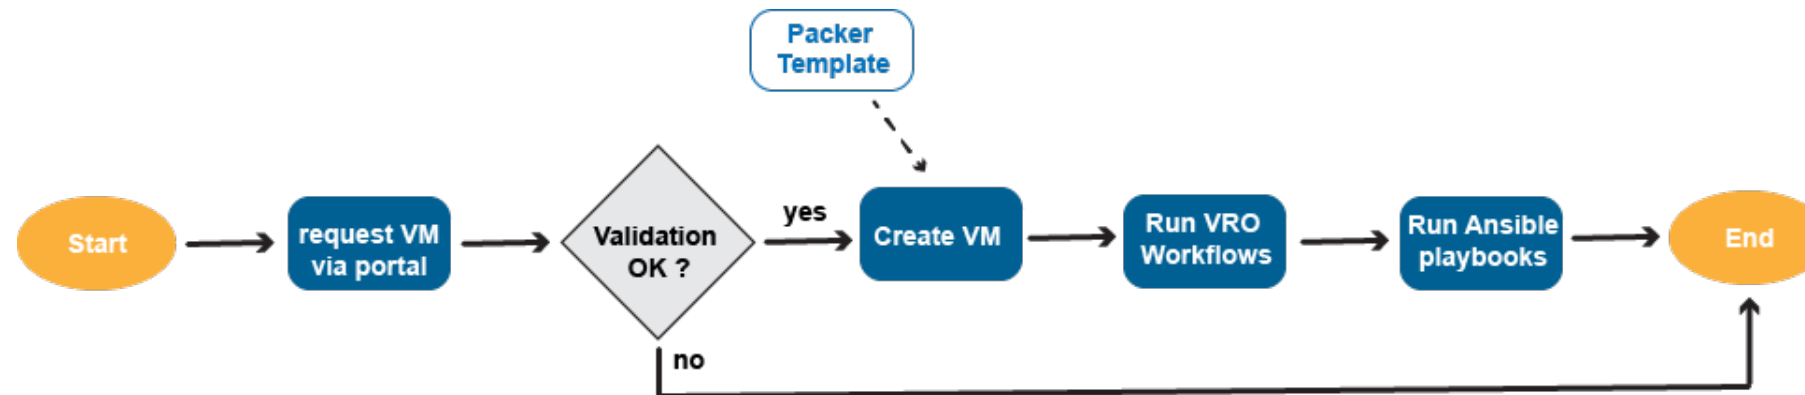
Portal usage:

- VM request
- VM management: replace the vcenter for the basic tasks on VM (for non-IT users)



## VM CREATION - PORTAL

- Template: In ARIA, we use OS Image template, created with HashiCorp [Packer](#).
- VM Creation: ARIA is linked with vCenter for VM creation. Use of YAML to define blueprints.
- VM Creation – Customization: [Workflow](#) scripting (JavaScript, Python, Node.js) for customization tasks during deployment.
- Post-deployment: call [ansible playbook](#) to execute customization tasks.





question



question

Thank you

The logo for the University of Lausanne, featuring the word "Unil" in a blue, cursive script font.



# KUBERNETES

Mikael Doche  
Nicolas Montes



## 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



Cluster Dashboard

Provider: RKE2    Kubernetes Version: v1.26.11+rke2r1    Created: 20 hours ago    [Install Monitoring](#)    [Add Cluster Badge](#)

178	Total Resources	9	Nodes	12	Deployments
-----	-----------------	---	-------	----	-------------

Capacity

Resource	Used	Reserved	Percentage
Pods	37 / 660		5.61%
CPU	0.22 / 24 cores	6.2 / 24 cores	0.90% / 25.83%
Memory	16 / 92 GiB	4.99 / 92 GiB	17.39% / 5.42%

Cluster Tools

Etcd    Scheduler    Controller Manager

Events

Full events list

v2.7.9

UNIL Kubernetes Service Portal




# 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


The screenshot shows the UNIL ARIA Service Broker interface. The top navigation bar includes the UNIL logo, 'ARIA', and 'UNIL | Université de Lausanne'. Below this, there are tabs for 'Service Broker CHANGE', 'Consume', 'Content & Policies', 'Infrastructure', and 'Inbox'. The main content area is titled 'New Request' and shows a request for 'rhel-kube' with version '59'. The 'Cluster' tab is selected, displaying the following configuration options:

Option	Value
Number of Workers *	2
Worker Size *	Medium - 2VCPU 8Gb
Cluster Storage *	XSmall - 50Go
Cluster Admin Groups *	acces-otobo-bas-g acces-soft-g all-users-portail-g bookstack-ci-g
OS Update Day	Thursday
OS Update Time	02:00
K8S Upgrade Run On	1st




# 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](http://portal.uks.unil.ch))



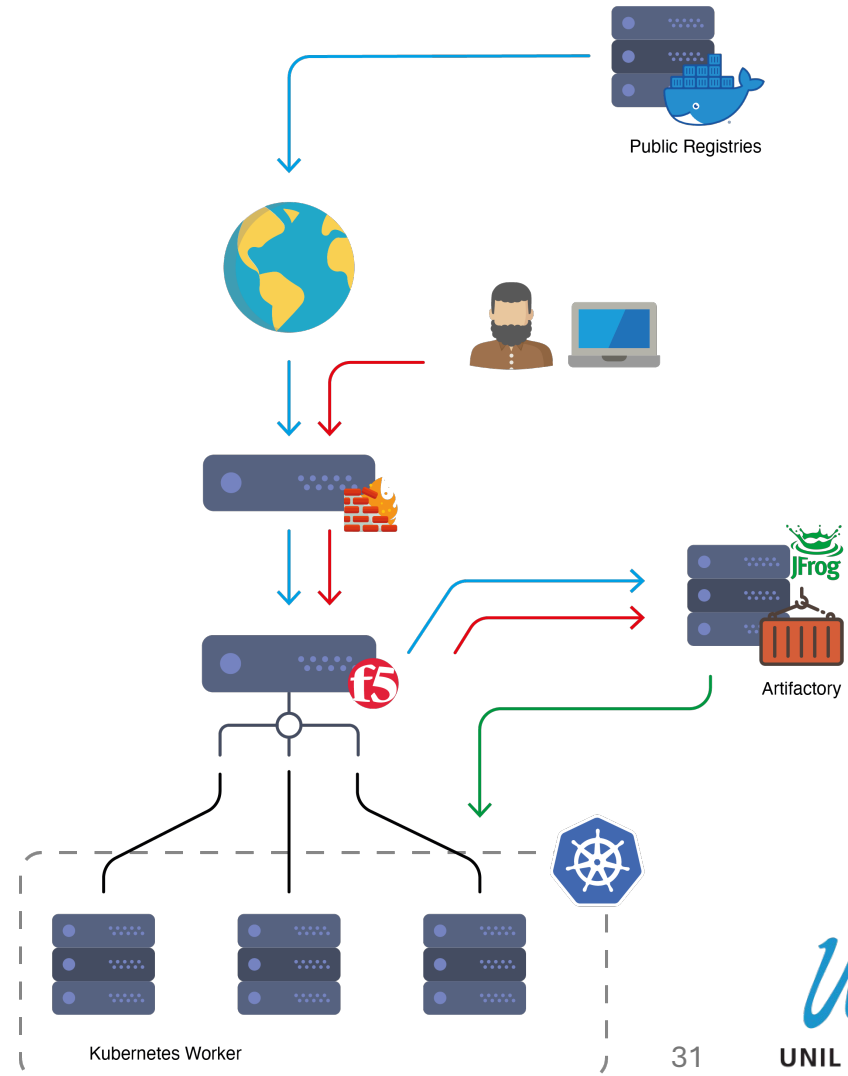
# KUBERNETES - Architecture - storage

- 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
- $104\text{TB available storage} / 3 (\text{x3 replicas}) = 34\text{TB usable for Kubernetes clusters}$



# KUBERNETES - Architecture - image registry

- Image registry relies on Artifactory from JFrog
- Each cluster has a dedicated cluster image registry
- Cluster registry has two default 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






question

Thank you

The logo for the University of Lausanne (Unil), featuring the word "Unil" in a blue, cursive script font.



# HPC (DCSR)

Division de **C**alcul et de **S**outien à la **R**echerche  
(Scientific Computing and Research Support Unit)

Volker Flegel



# mission & goals

- Provide support to **research projects** for all faculties
- Computational resources
  - Storage (NAS, Object Store, Cluster FS, Encrypted Storage, VM datastore)
  - High Performance Computing
  - Secure Computing Environment for sensitive data
  - Virtual Machines (Servers, Desktops)
- Logistical support for research projects
  - Infrastructure support
  - Software optimization & development
  - Research Analysis multi-level support
  - Machine Learning
  - Courses & Training

# storage resources

- Generic storage for **research data**

- **NAS**

- ~5 PB (Isilon)
- Accessible from Campus
- Non-sensitive or personal data

- **S3 ObjectStore**

- ~1.5 PB (Scality)
- Accessible worldwide
- Non-sensitive or personal data
- Data exchange with external partners

- **Long Term Storage**

- $\infty$  (StorNext – Tape HSM)
- No direct access
- Non-sensitive or personal data
- Archival space for finished projects

- **Tresorit (cloud)**

- External storage partner
- Sensitive data
- Encrypted

# computational resources

- Virtual Machine Infrastructure

- Virtual Desktop Infrastructure

- VMware VDI (Horizon)
  - 4 Nodes
    - 128 cores
    - 1 – 2 TB RAM
    - GPU: Tesla T4, Tesla V100
  - Linux & Windows Desktops
  - Accessible from Campus
  - Non-sensitive or personal data
- Security:
    - NSX-T Network Isolation

- Virtual Server Infrastructure (VSI)

- VMware ESXi
  - 3 Nodes
    - 128 cores
    - 2 TB RAM
  - Host Linux & Windows Servers
  - User Managed services
  - Expose services Internal or External
  - Non-sensitive data
- Security:
    - NSX-T Network Isolation

- VSI - Sensitive

- VMware ESXi
  - 2 Nodes
    - 64 cores
    - 512 GB RAM
  - Host Linux & Windows Servers
  - User Managed services
  - Expose services Internal or External
  - Sensitive data
- Security:
    - NSX-T Network Isolation

PowerVault

- ~400 TB



# computational resources

- HPC Infrastructure

- HPC cluster - Curnagl

- Standard HPC
- 88 Nodes + 8 GPU Nodes
  - 48 cores AMD Epyc
  - 84 nodes: 512 GB RAM, 12 nodes: 1TB RAM
  - GPU: 2x A100 / node
- Non-sensitive or personal data
- Slurm scheduler

- Security:

- Network separation & Firewalled

- HPC cluster - Urblauna

- HPC for sensitive data
- 16 Nodes + 2 GPU Nodes
  - 48 cores AMD Epyc
  - 1 TB RAM
  - GPU: 2x A100 / node

- Sensitive data
- Slurm scheduler

- Security:

- “Air-gapped” Network separation & Firewalled
- 2FA
- Guacamole WebRDP
- JumpHost for SSH & data upload (SFTP)
- POSIX Access rights

- OpenStack - SENSEA

- Cloud Computing for Medical data
- 18 Nodes + 4 GPU Nodes
  - 20 - 52 cores
  - 384 – 512 GB RAM
  - GPU: 6x A100, 2x RTX2080
- Research project isolation (Tenants)
- Tenants managed by DCSR
- Encrypted Filesystems per Tenant
- Sensitive / medical data

- Security:

- “Air-gapped” Network separation & Firewalled
- 2FA
- Guacamole WebRDP
- Encrypted data upload
- Per Tenant encrypted Filesystems

GPFS (~2 PB)

NAS

- On login-node

GPFS (~1 PB)

StorNext HSM

- At-rest tape encryption

CEPH

- At-rest encryption

WEKA.io

- At-rest & in-transit per Tenant encryption
- Tiering (S3)

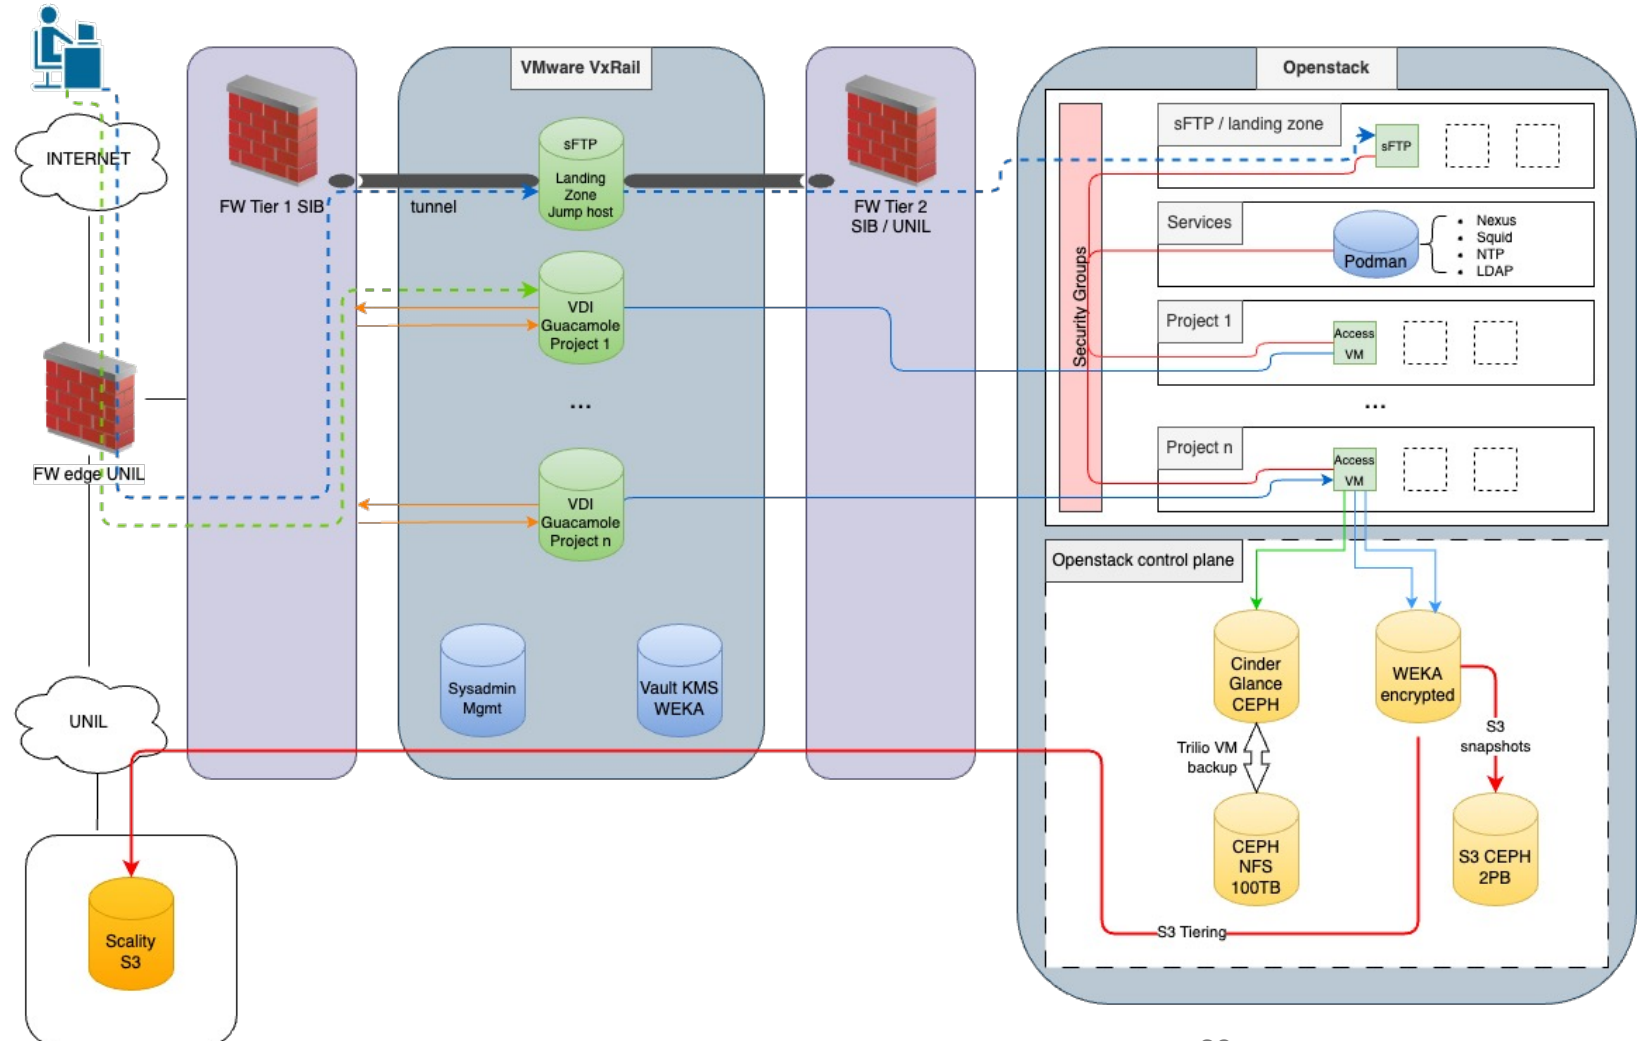
# SENSA - openstack for sensitive / medical data

- Access restrictions

- Allowlists (IP, VPN, ...)
- 2FA / eduID
- Incoming data via transfer requests
- Outgoing traffic only to proxied resources
- No Admin rights on Tenant or VM

- Project level isolation

- Web-RDP service
- Security Group network isolation
- Encrypted storage (in-transit, at rest)



question





question

Thank you

The logo for the University of Lausanne (Unil), featuring the word "Unil" in a blue, cursive script font.