

A stage for Samba in the era of the container platform!?!...

sambaXP 2020

Michael Adam

Red Hat / Samba Team

What we'll discuss today

Kubernetes

- ▶ Intro
- ▶ Storage
- ▶ CSI
- ▶ Operators

Storage Location

- ▶ Platform
- ▶ Infra vs Platform
- ▶ Kubevirt


Enter Samba

- ▶ Use cases
- ▶ Container
- ▶ Operator
- ▶ Provisioner?



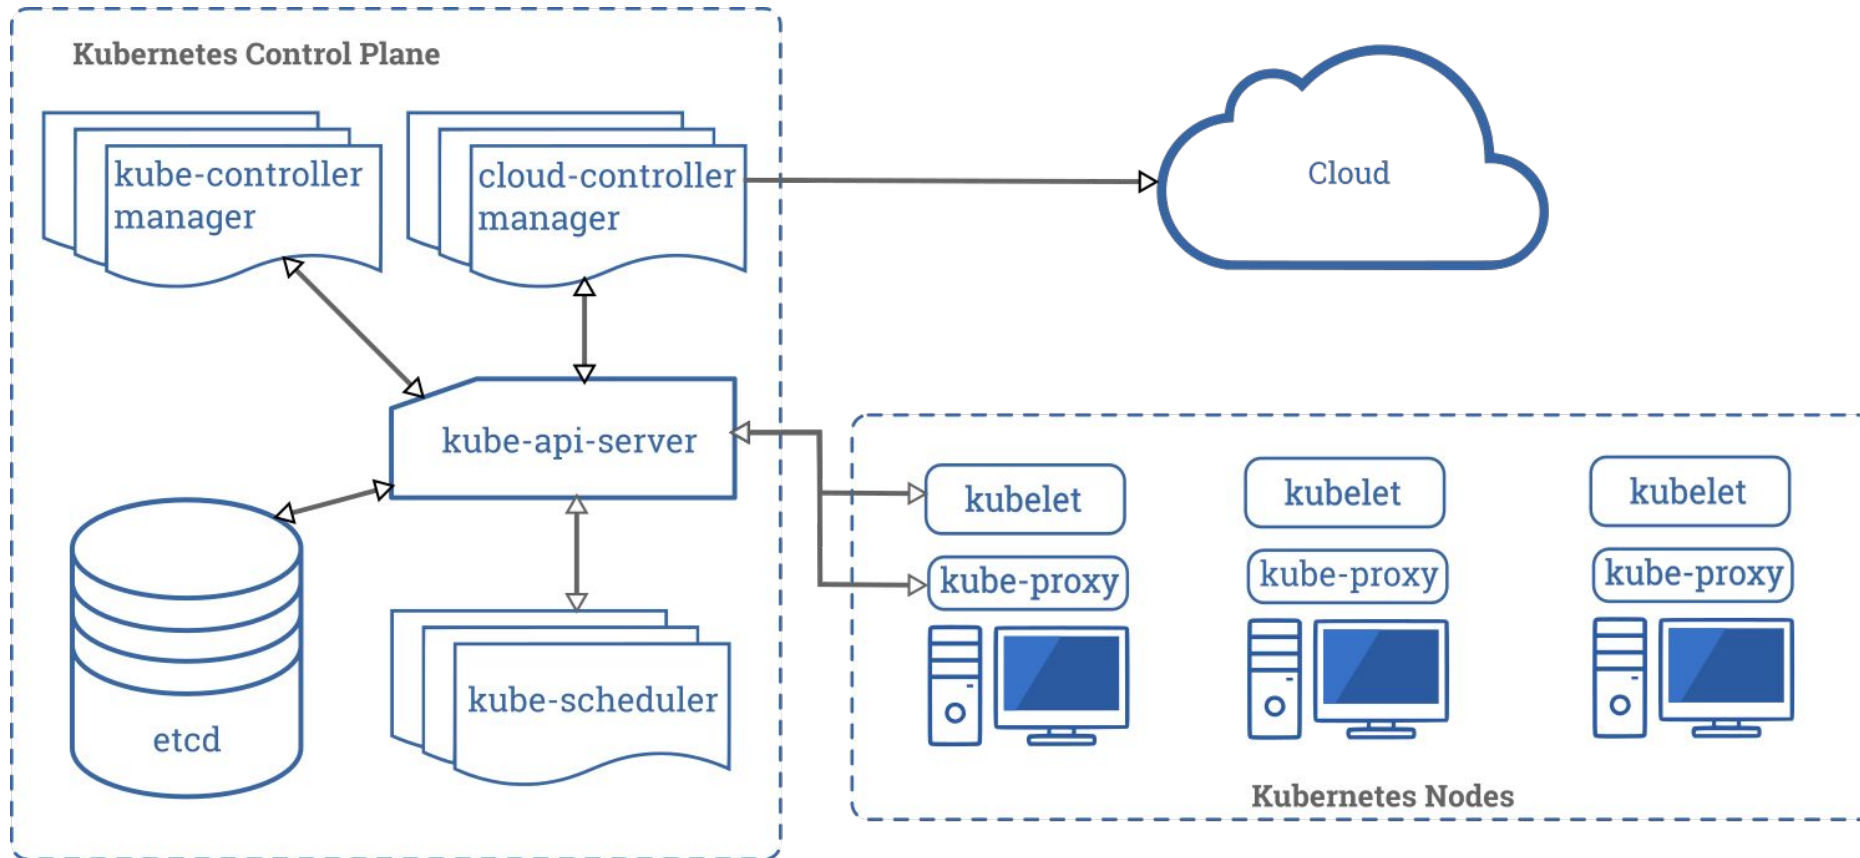
kubernetes

- ▶ Intro
- ▶ Storage
- ▶ CSI
- ▶ Operators



Kubernetes (K8s) is an open-source system for automating deployment, scaling, and management of containerized applications.

<https://kubernetes.io/>



Storage abstraction in k8s

Some terminology



PersistentVolume (PV)


K8s object wrapping a backend storage volume. This could live inside of k8s or somewhere outside.

Types: *FileSystem* and *Block*



PersistentVolumeClaim (PVC)

Handle for a user to consume a PV in an application / namespace.



Mount plugin

Volumes are mounted on the host node and bind-mounted into the container. (Block PVs: device in container)



Dynamic Provisioning: Volume self-service for users

More terminology...



(Dynamic) Provisioner

Server component responding to PVC create requests



StorageClass (SC)

K8s resource pointing to a dynamic provisioner, configured with a certain backend storage pool.



Flow

- ▶ User issues ``kubectl create -f pvc.yml``
- ▶ PVC ⇒ StorageClass ⇒ provisioner
- ▶ Provisioner reacts and creates PVC



Small Demo



Operators

Automating Application Management

- ▶ Server component: controller
- ▶ Custom resources (CR): new type k8s APIs
- ▶ CR: Desired state
- ▶ Operator reconciles to achieve desired state

- ▶ Initial deployment
- ▶ Keeps application running (fixes outages)
- ▶ Changes in scaling (replica)
- ▶ etc...

Infrastructure:

- ▶ <https://github.com/operator-framework>
- ▶ <https://operatorhub.io/>
- ▶ [operator-lifecycle-manager](#)

Examples:

- ▶ Rook-ceph (<https://github.com/rook/rook>)
- ▶ [prometheus-operator](#)
- ▶ ...

CSI

Container Storage Interface



CONTAINER
STORAGE
INTERFACE

- ▶ Further standardizing storage abstractions beyond k8s.
 - (kubernetes, cloud foundry, mesos, nomad, ...)
- ▶ The new standard in kubernetes.

- ▶ <https://github.com/container-storage-interface/spec>

- ▶ Example: <https://github.com/ceph/ceph-csi>

File, Block, and... Object?

Object access (S3) not currently part of core k8s / CSI

Extensions possible with CRDs:

- ▶ ObjectBucket
- ▶ ObjectBucketClaim
- ▶ bucket provisioner

Discussed as native k8s enhancement:

<https://github.com/kubernetes/enhancements/pull/1383>

Storage Location

- ▶ Platform
- ▶ Infra vs Platform
- ▶ Kubevirt



Storage in the platform (k8s)

Converged with the platform

- ▶ Block storage coming from somewhere
 - local, SAN, cloud, ...
- ▶ SDS layer as containers in k8s
- ▶ Scaling with k8s
- ▶ Managed by operator(s)

Red Hat OpenShift Container Storage (OCS)

Evolution

1st Generation

- ▶ Based on Gluster
- ▶ Openshift 3.X
- ▶ Pioneer in dynamic provisioning
- ▶ Started 2016



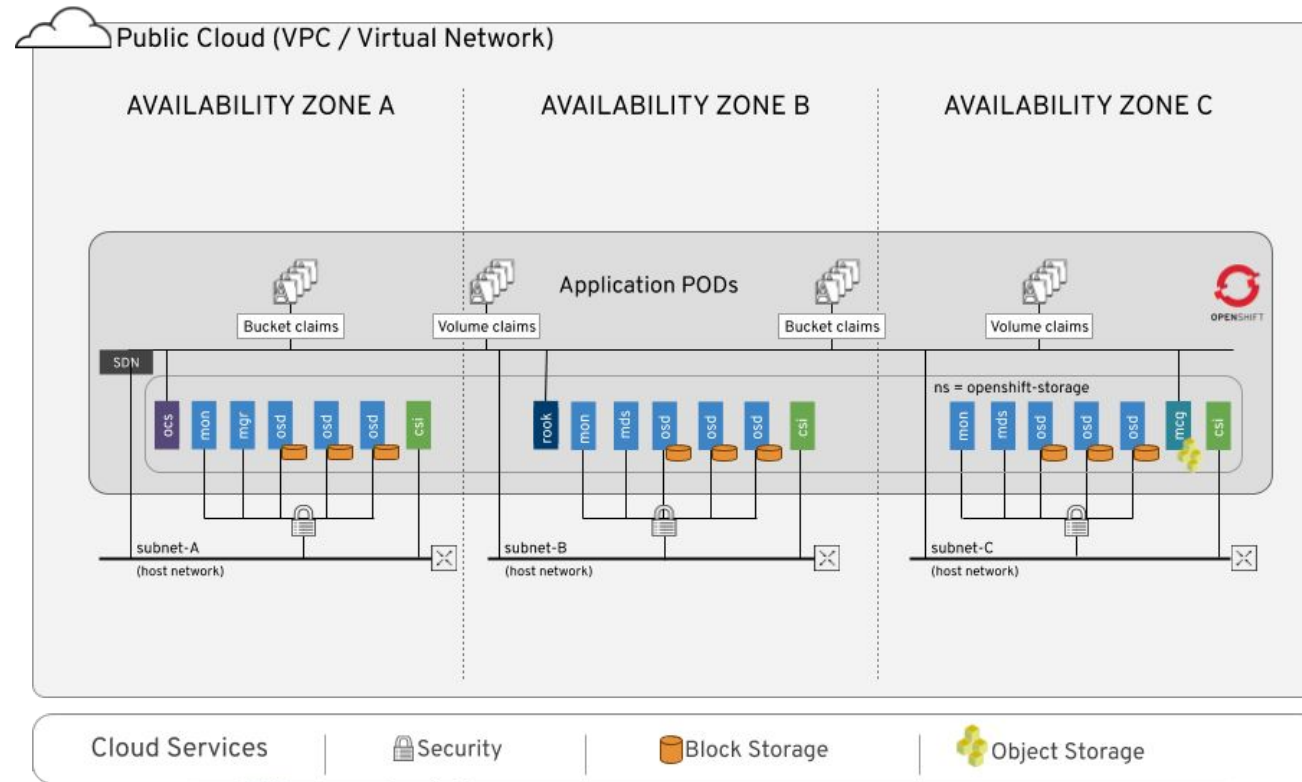
2nd Generation

- ▶ Based on Ceph
- ▶ OpenShift 4.X
- ▶ Operator driven
- ▶ CSI
- ▶ GUI ("Console")
- ▶ 2019/2020

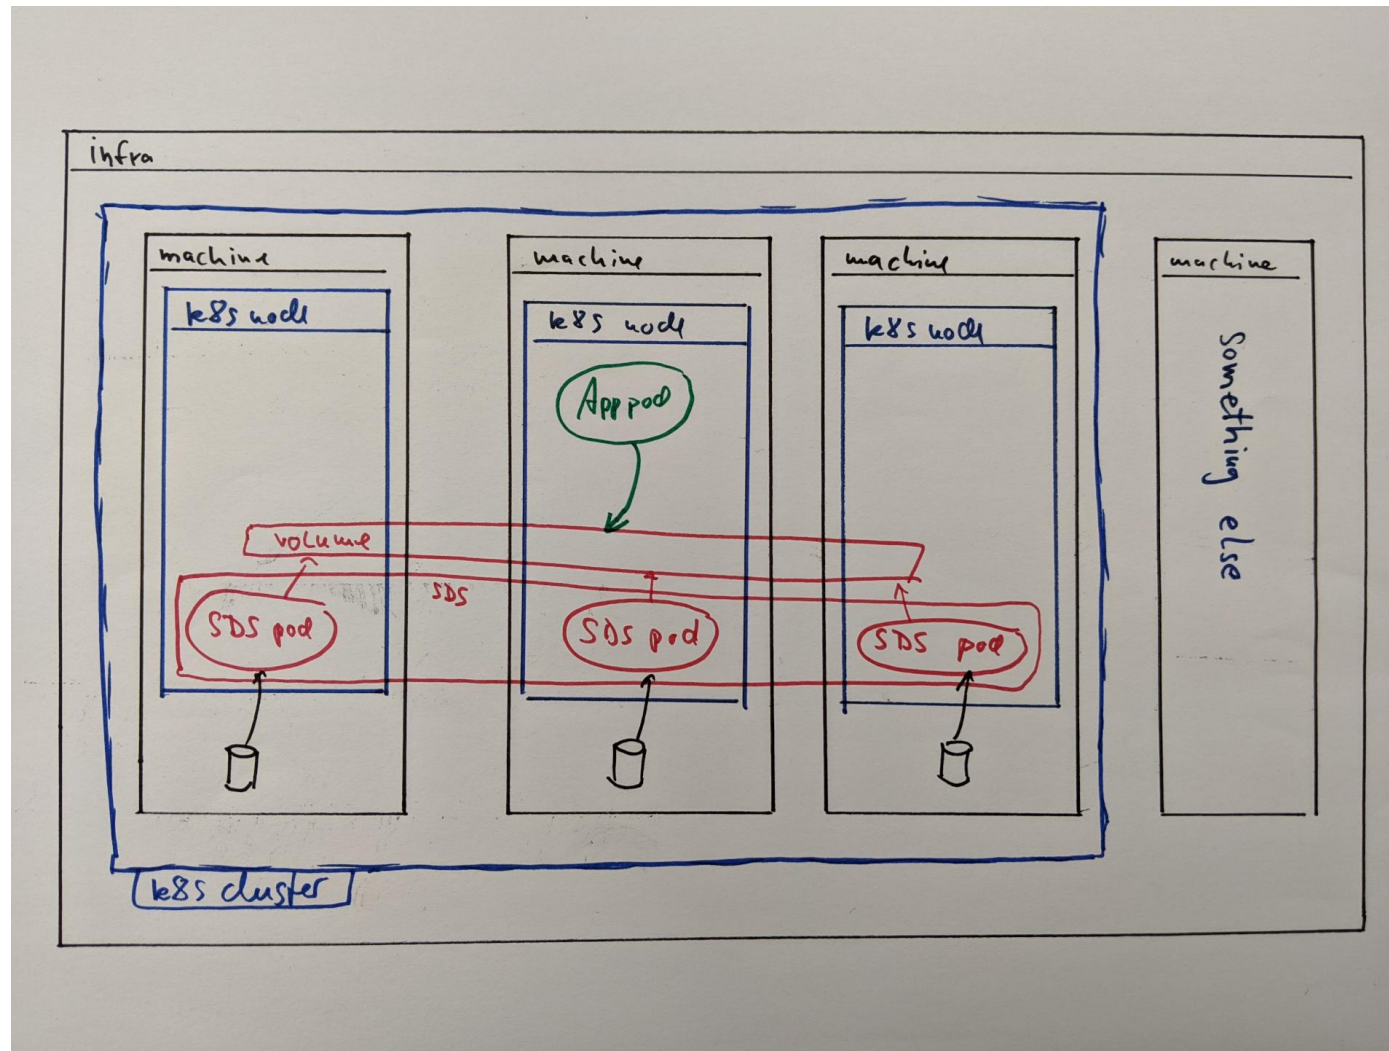


OCS

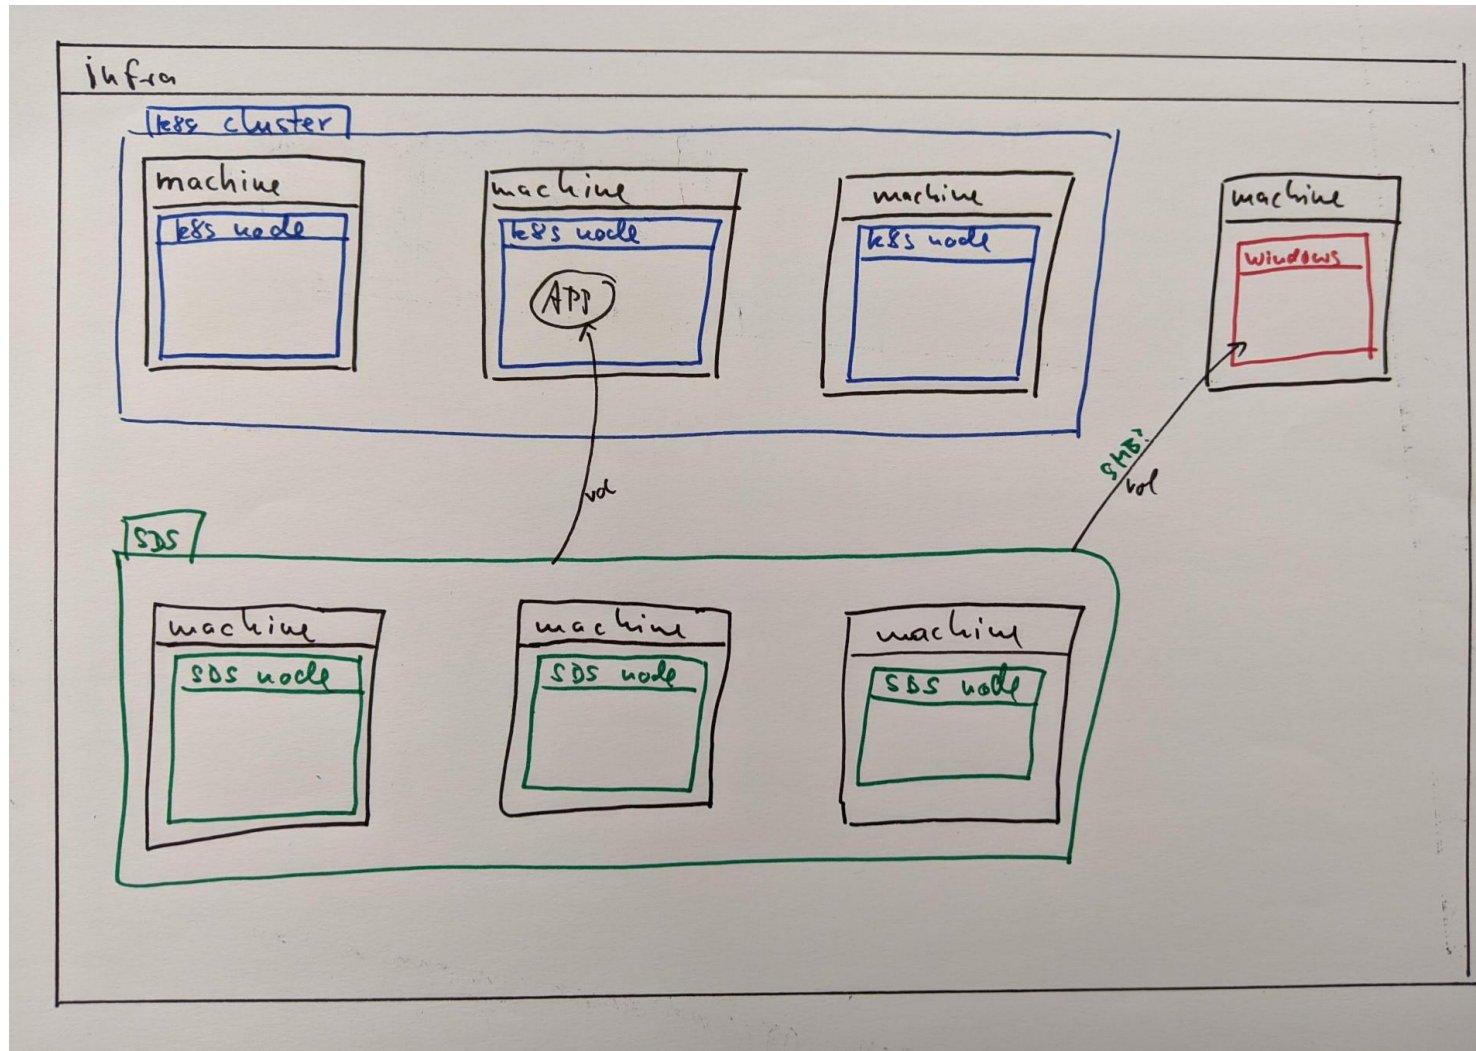
Architecture



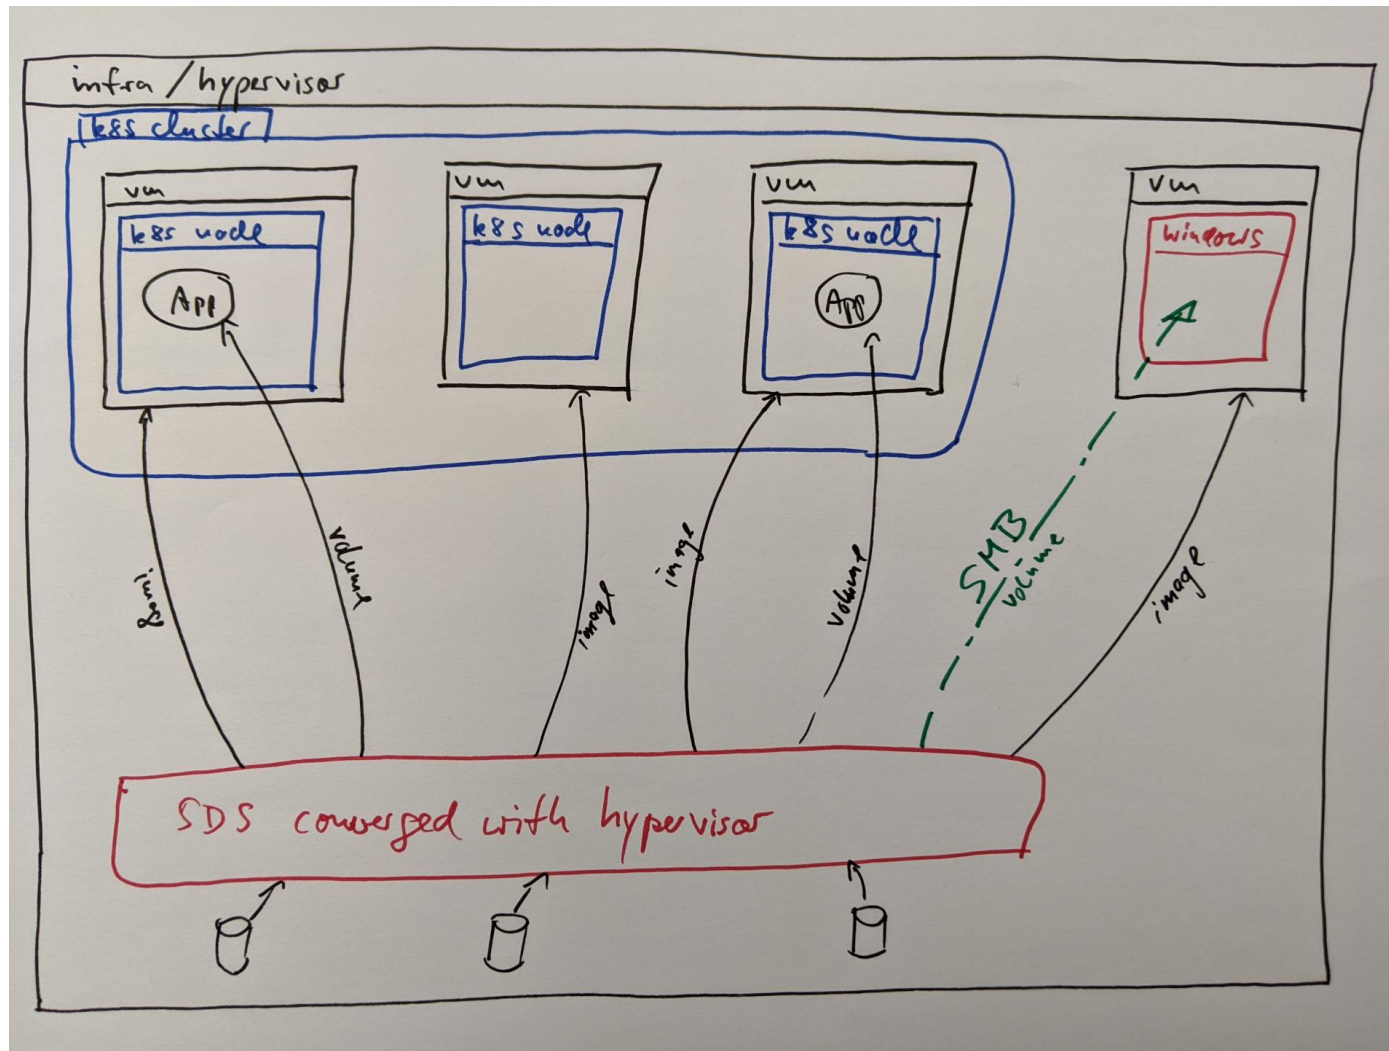
Storage converged with the platform (k8s)




Storage on top of the infrastructure



Storage converged with infrastructure



Reverse the layers: Platform is Infrastructure



kubevirt

Running VMs in containers



- ▶ Expand k8s with APIs to run and manage VMs
- ▶ Run VMs alongside container apps
- ▶ Consume the same storage

- ▶ <https://github.com/kubevirt/kubevirt>
- ▶ <https://kubevirt.io/>

Enter SAMBA

- ▶ Use Cases
- ▶ Container
- ▶ Operator
- ▶ Provisioner?



Use cases for Samba in k8s

Windows VMs running in kubevirt as k8s containers.

- ▶ Need storage to store their application data
- ▶ Need to exchange data between VMs
- ▶ Prefer SMB over NFS...

Mixed k8s clusters: Windows worker nodes


- ▶ How can Windows workers use PVCs from Linux nodes?
- ▶ SMB for file
- ▶ iSCSI for block

Extra thought:

- ▶ Provide an AD service for projects with Samba!

The basic idea

Extremely simple Samba service
container concept



SMB Service on PVCs

Establish notion of an SMB Service per PVC. Even an SMB-PVC. (1:1 correspondence.)

Managed by an operator.



Very simple Samba container

No complicated clustered Samba.

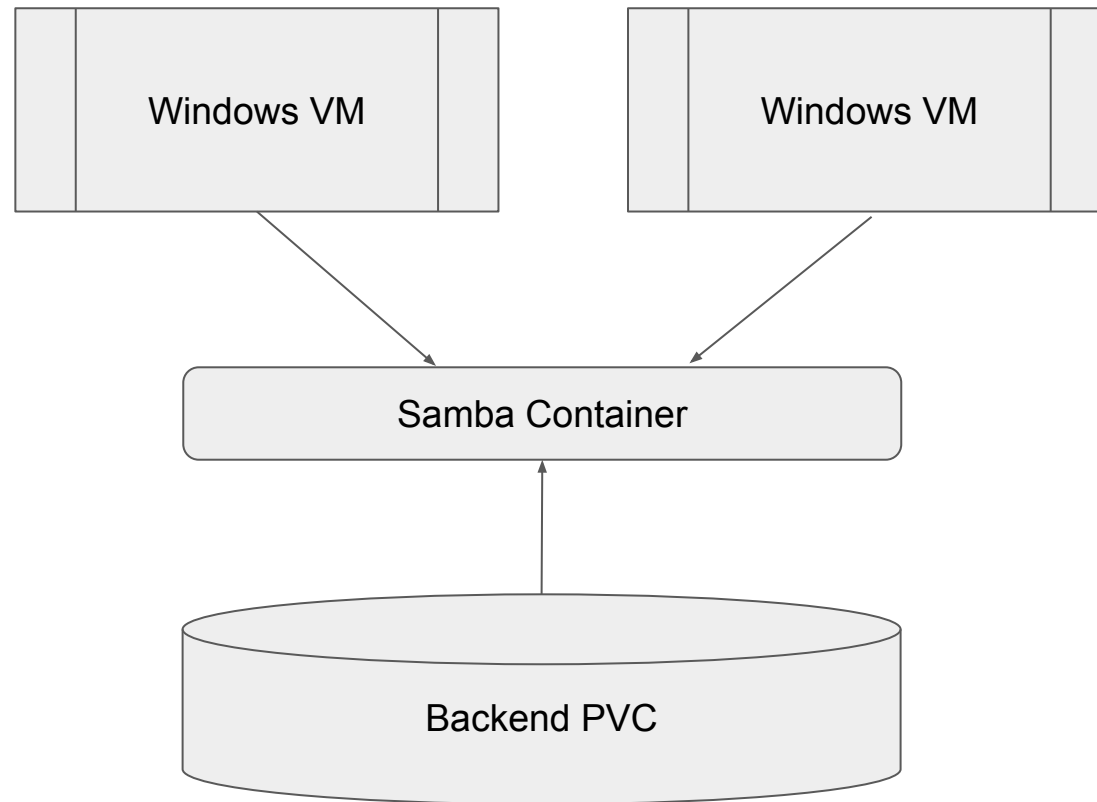
Just re-exporting a directory from the container.



Possibly: Provisioner for SMB-access PVCs

A CSI driver for mounting PVCs via SMB already exists.





Simple Samba Container

<https://github.com/obnoxxx/samba-container>

<https://quay.io/repository/obnox/samba-centos8>

- ▶ Runs only smbd
- ▶ Serves /share out as [share]
- ▶ Currently only supports one user

TODO:

- ▶ Multiple users
- ▶ AD integration

Samba Operator

<https://github.com/obnoxxx/samba-operator>

<https://quay.io/repository/obnox/samba-operator>

Two CRDs:

- ▶ SmbService : use existing PVC
- ▶ SmbPvc: create new PVC and SmbService

Demo!



Thank you!

obnox@samba.org

obnox@redhat.com

Questions?