



KubeCon



CloudNativeCon

Europe 2020




Virtual



KEEP CLOUD NATIVE

CONNECTED





KubeCon



CloudNativeCon

Europe 2020

Virtual

Using KubeVirt to Run VMs at Scale

Marcus Sorensen, Apple Inc

Fabian Deutsch, Red Hat



KubeCon



CloudNativeCon

Europe 2020

Virtual

- Fabian
 - Intro
 - Architecture
 - ~~Classical VM use case with OKD UI~~
- Marcus
 - Running KubeVirt in production
 - Lessons learned and pain point



KubeCon



CloudNativeCon

Europe 2020

Virtual

The future is bright



KubeCon



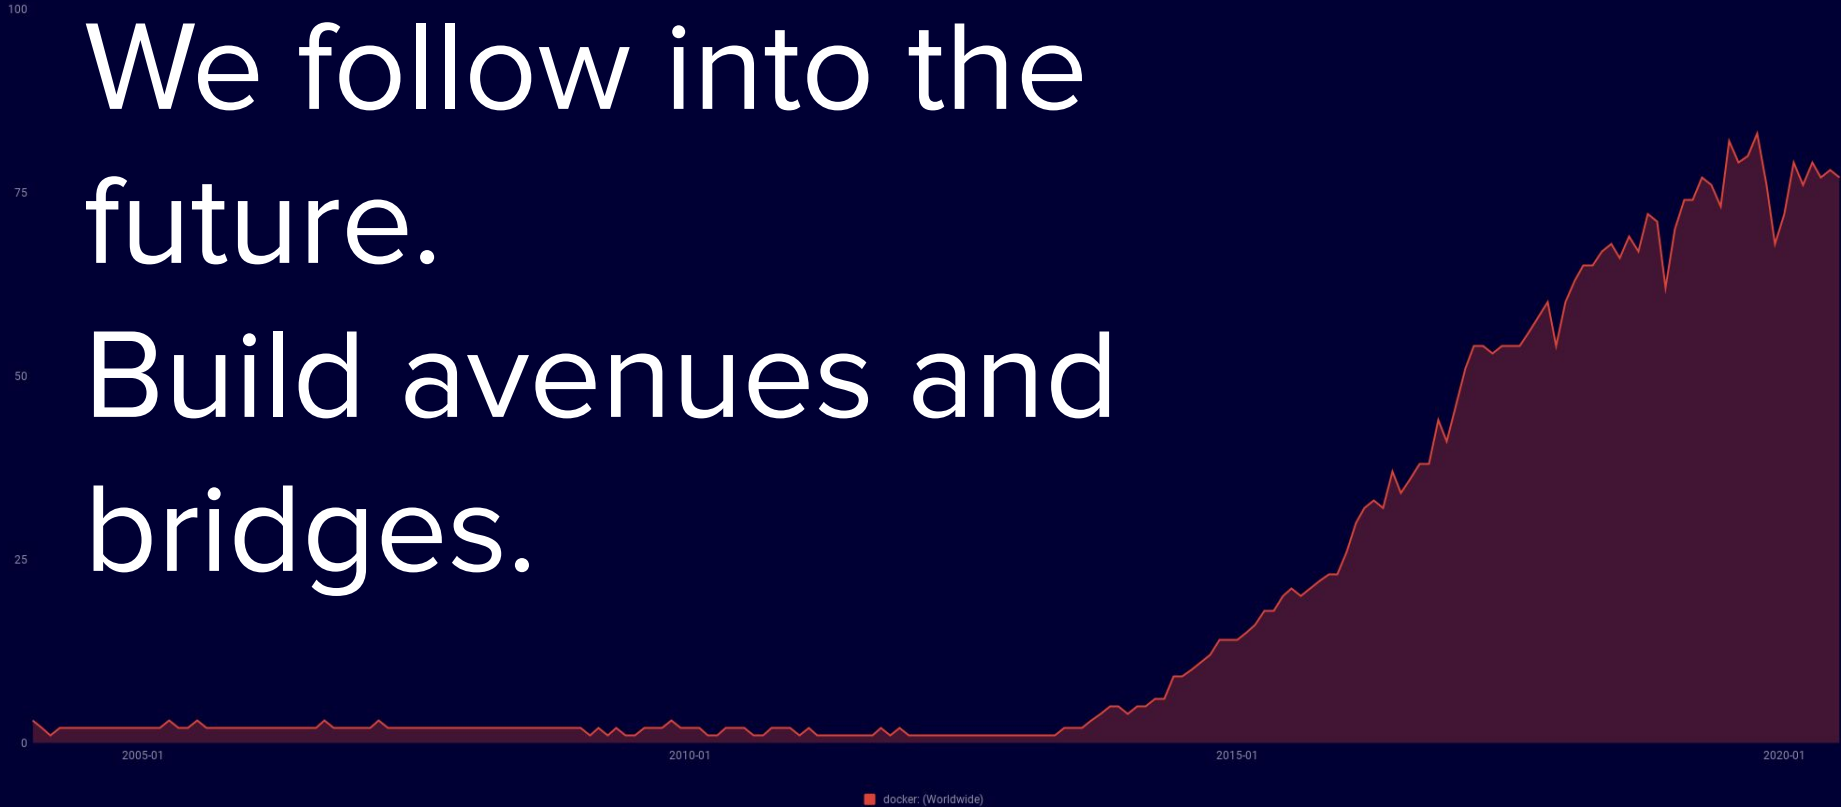
CloudNativeCon

Europe 2020

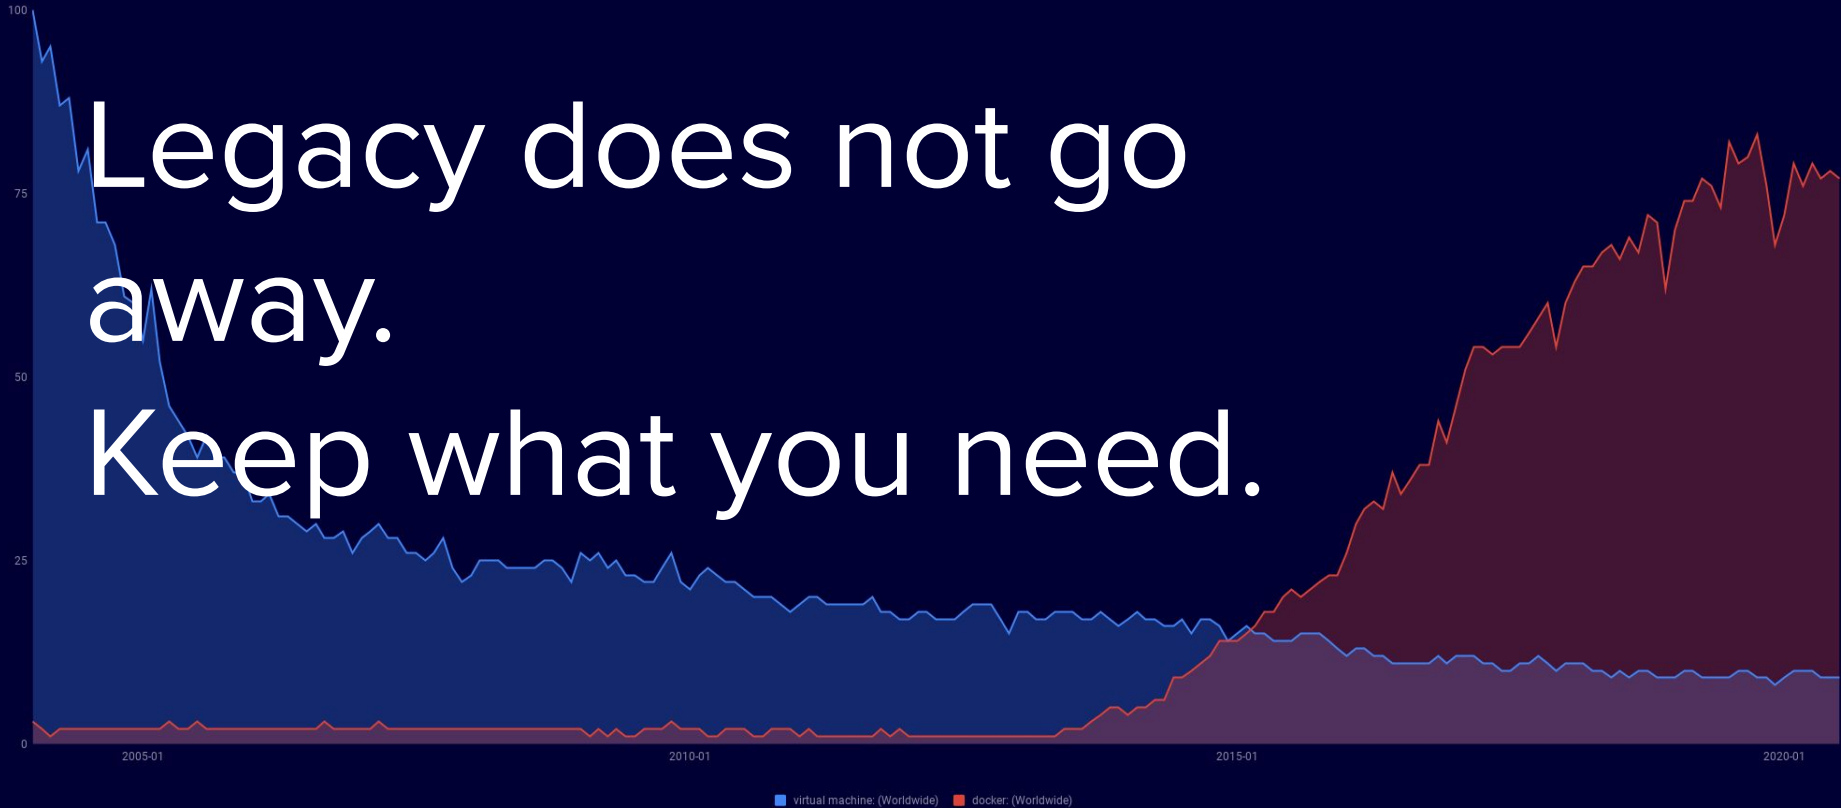
Virtual

The future is brighter *er* *with*
Kubernetes

We follow into the
future.
Build avenues and
bridges.



Legacy does not go
away.
Keep what you need.





 **CLOUD NATIVE**
COMPUTING FOUNDATION

CNCF Sandbox Project



KubeCon



CloudNativeCon

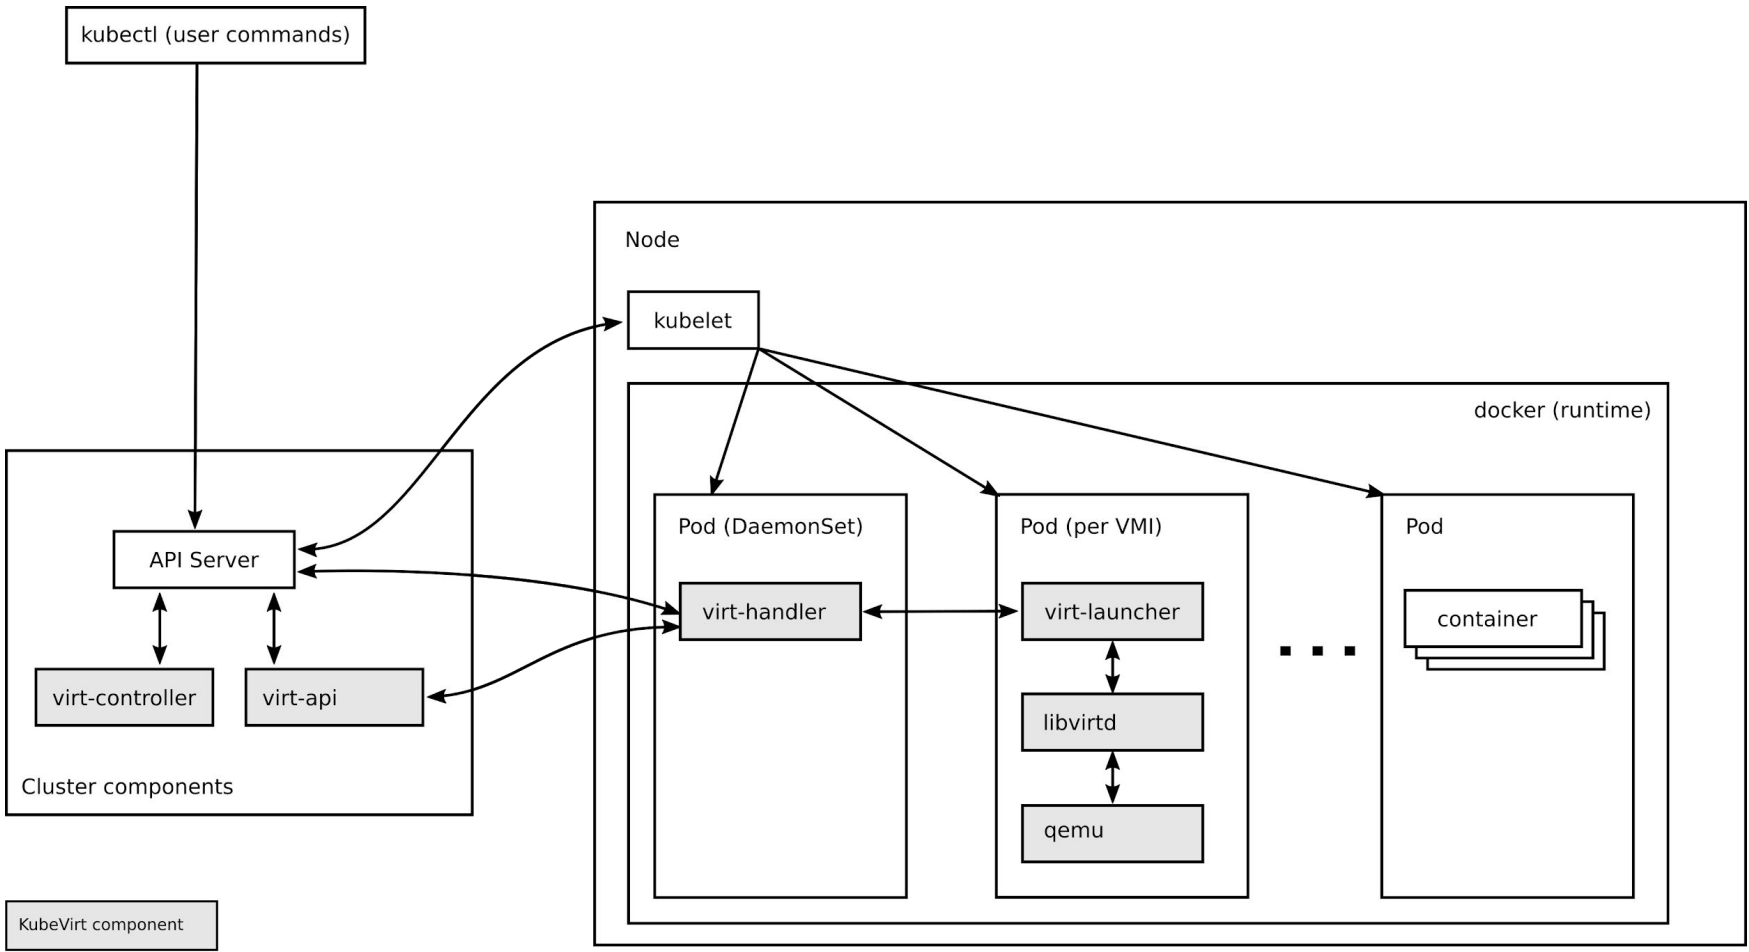
Europe 2020

Virtual

Marry

Virtual Machines Legacy and Kubernetes Future*

Using common Kubernetes ingredients: Operators, CRDs, and Pods



KubeVirt component

Kubernetes component



KubeCon



CloudNativeCon

Europe 2020

Virtual

```
$ kubectl ...  
# to rule them all
```

Start, Stop,
Live Migration,
Multi Network, SR-IOV multus*
Block Storage, Persistent Volumes
Snapshots, Volume Snapshots
...



KubeCon



CloudNativeCon

Europe 2020

Virtual

Develop on and use One platform

Enable developer self-service, simplify and unify access to resources.



Virtual

→ Kube-on-Kube

Kubernetes hosted on KubeVirt. Nothing else.



KubeCon



CloudNativeCon

Europe 2020

Virtual

Operate

One platform

Concentrate an organization's knowledge on
operating Kubernetes



KubeCon



CloudNativeCon

Europe 2020

Virtual

Scale

One platform

Using common Kubernetes ingredients:
Operators, CRDs, and Pods



KubeCon



CloudNativeCon

Europe 2020

Virtual

Marcus Sorensen
Cloud Developer, Apple Inc

Why KubeVirt?

- Enables VM workloads to leverage Kubernetes
- Enables single compute platform for containers and VMs
- Works with standard upstream Kubernetes



KubeCon



CloudNativeCon

Europe 2020

Virtual

Why KubeVirt?

- Allows VM workloads to leverage Kubernetes



KubeCon



CloudNativeCon

Europe 2020

Virtual

Why KubeVirt?

- Allows VM workloads to leverage Kubernetes
- Works with standard upstream Kubernetes

Why KubeVirt?

- Allows VM workloads to leverage Kubernetes
- Works with standard upstream Kubernetes
- Enables single compute platform for containers and VMs



KubeCon



CloudNativeCon

Europe 2020

Virtual

Challenges



Challenges

- Opinionated virt-operator deployment model

Challenges

- Opinionated virt-operator deployment model
- Self-signed certificates and cert authority

Challenges

- Opinionated virt-operator deployment model
- Self-signed certificates and cert authority
- No support for a custom kubelet root path



KubeCon



CloudNativeCon

Europe 2020

Virtual

Added --kubelet-pods-dir to virt-handler

Challenges

- Opinionated virt-operator deployment model
- Usage of self-signed certificates and CA
- No support for a custom kubelet root path
- Implicit requirement for NET_RAW capability, fixed

Challenges

- Opinionated virt-operator deployment model
- Usage of self-signed certificates and CA
- No support for a custom kubelet root path
- Implicit requirement for NET_RAW capability, fixed
- Found and fixed Containerd incompatibility



KubeCon



CloudNativeCon

Europe 2020

Virtual

Infrastructure Developer Use Cases



KubeCon



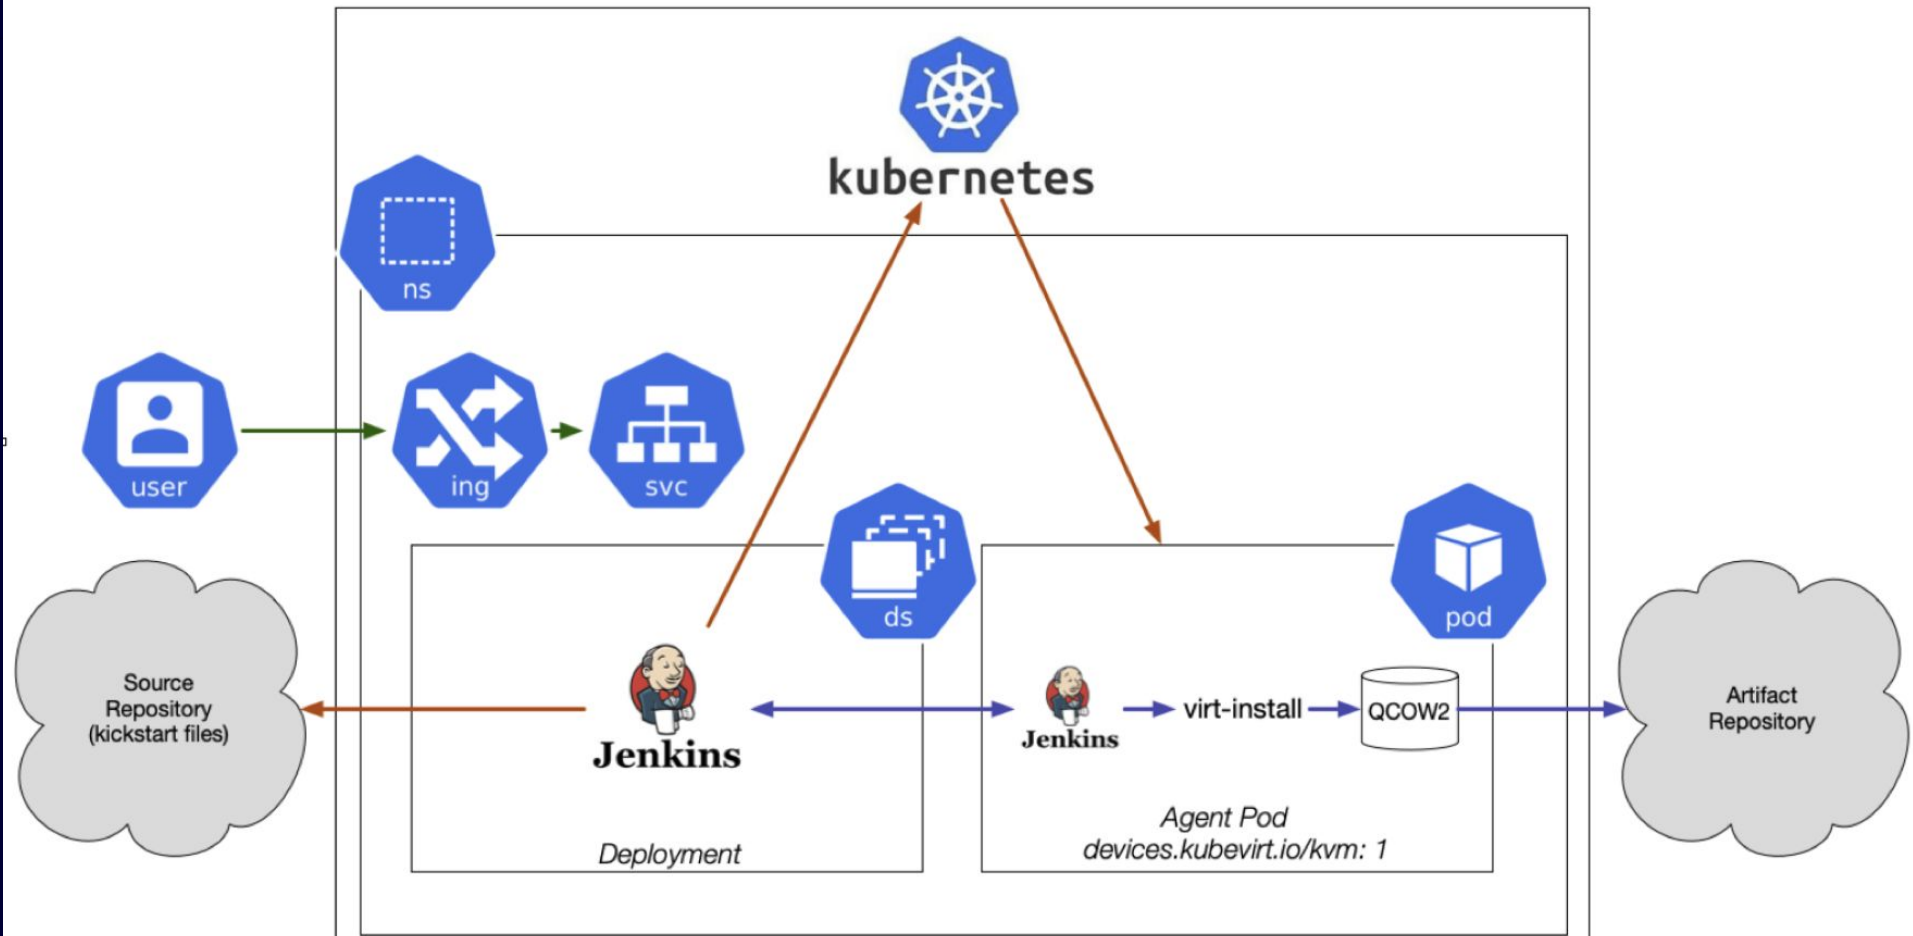
CloudNativeCon

Europe 2020

Virtual

VM Template Builder

VM Template Build Pipeline



Kubernetes Pod Template

Name

Namespace

Labels

Usage

Pod template to inherit from

Max number of instances

Pod Retention

Time in minutes to retain agent when idle

Time in seconds for Pod deadline

Timeout in seconds for Jenkins connection

Annotations

List of annotations to set in slave pod

Raw yaml for the Pod

```
resources:
  limits:
    cpu: "2"
    devices.kubevirt.io/kvm: "1"
    memory: 4Gi
  requests:
    cpu: "1"
    devices.kubevirt.io/kvm: "1"
    memory: 2Gi
```

Virtual



KubeCon

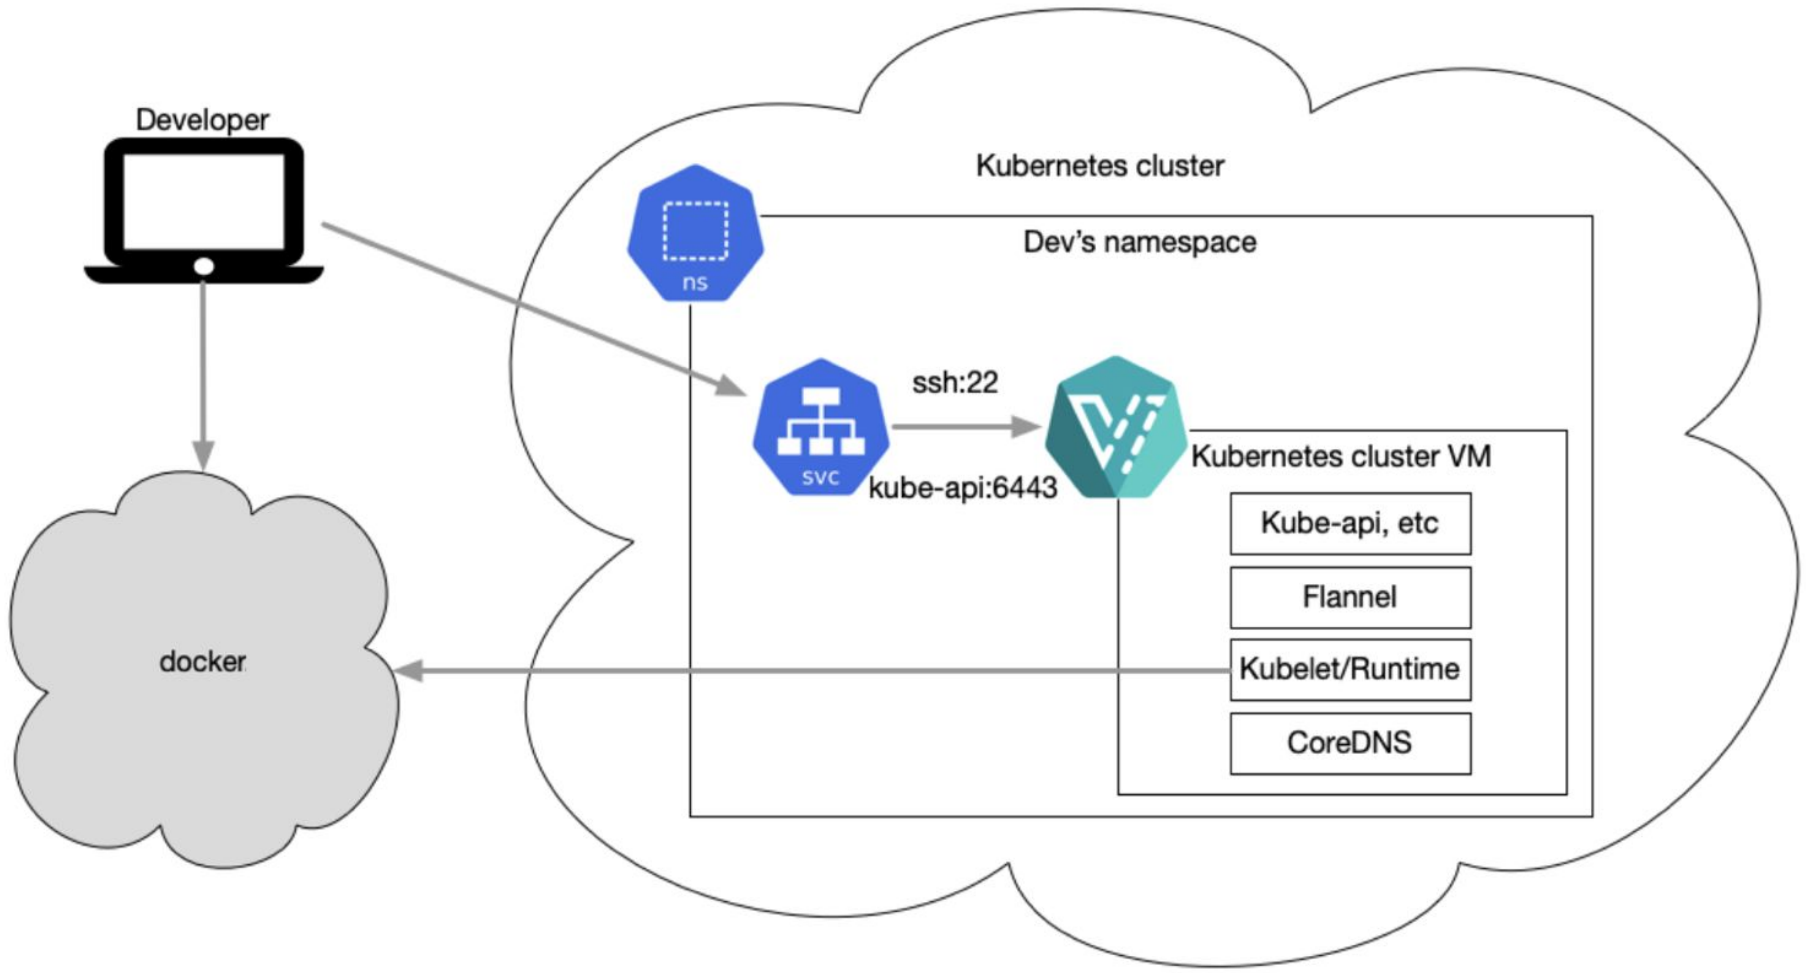


CloudNativeCon

Europe 2020

Virtual

Kubernetes Clusters





KubeCon



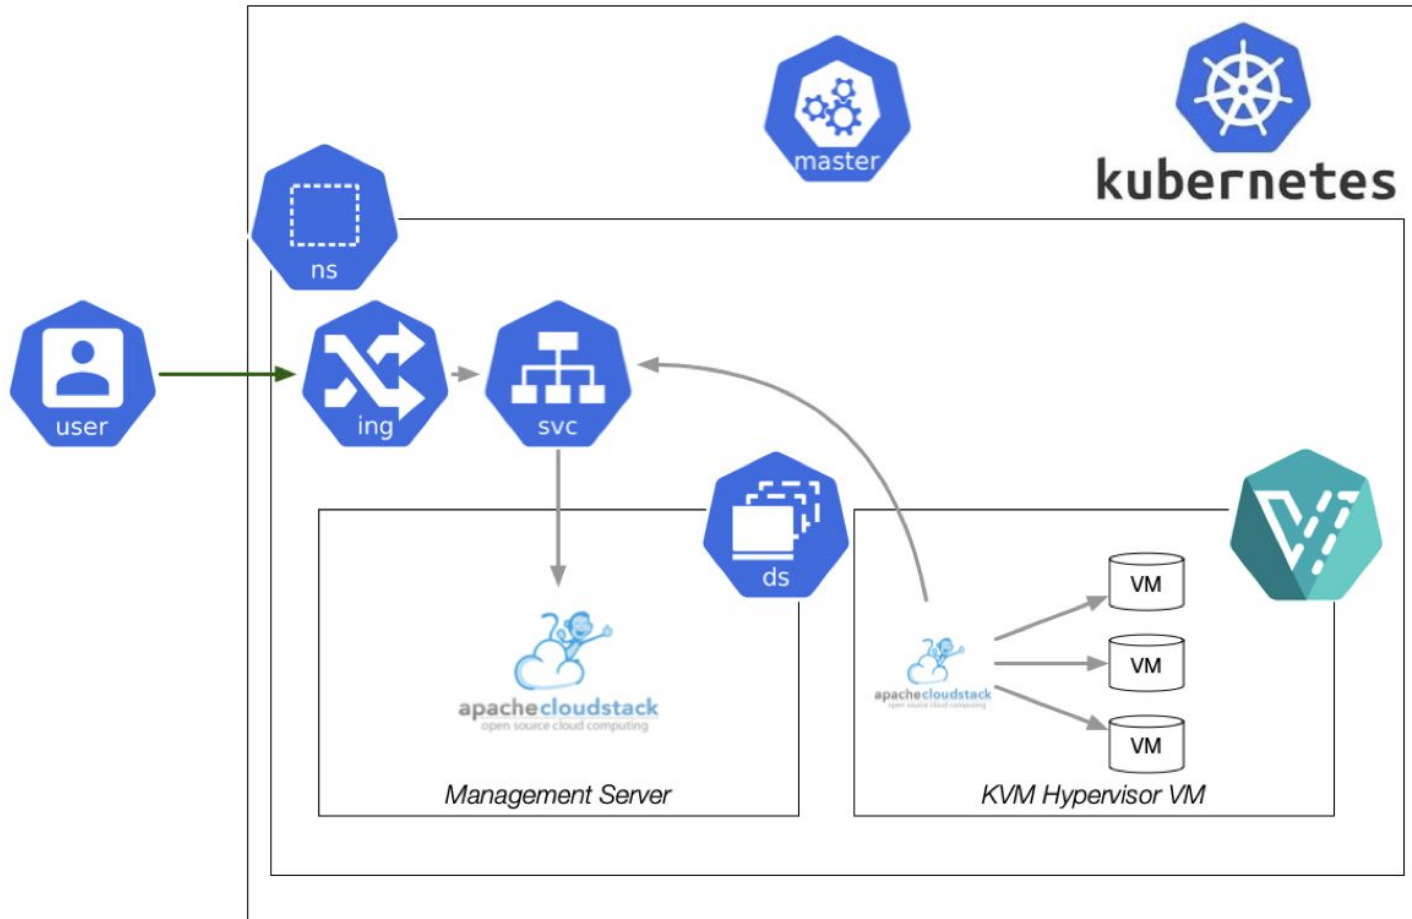
CloudNativeCon

Europe 2020

Virtual

Hypervisor Development

Cloud Development/Test Environment





KubeCon



CloudNativeCon

Europe 2020

Virtual

Kubernetes Node Nested KVM Support



KubeCon



CloudNativeCon

Europe 2020

Virtual

```
# cat /etc/modprobe.d/kvm.conf  
options kvm_intel nested=1  
options kvm_amd nested=1
```



KubeCon



CloudNativeCon

Europe 2020

Virtual

Other Use Cases

- Appliances
- Other OS
- Bare metal requests



KubeCon



CloudNativeCon

Europe 2020

Virtual

Tips

- KubeVirt tests can be run as a cron job



KubeCon



CloudNativeCon

Europe 2020

Virtual

Tips

- KubeVirt tests can be run as a cron job
- Containerized Data Importer **DataVolume**



KubeCon



CloudNativeCon

Europe 2020

Virtual

Tips

- KubeVirt tests can be run as a cron job
- Containerized Data Importer DataVolume
- Watch quotas and IP use with virt-handler



Tips

- KubeVirt tests can be run as a cron job
- Containerized Data Importer DataVolume
- Watch quotas and IP use with virt-handler
- /dev/kvm can be used standalone



KubeCon



CloudNativeCon

Europe 2020

Virtual

Thank You!



KubeCon



CloudNativeCon

Europe 2020

Virtual

Q&A