



Kubernetes in production: avoiding the pitfalls



Horacio Gonzalez

@LostInBrittany



Who are we?


Introducing ourselves and
introducing OVHcloud




Horacio Gonzalez

@LostInBrittany

Spaniard lost in Brittany,
developer, dreamer and
all-around geek



@LostInBrittany 





OVHcloud: A Global Leader

200k Private cloud
VMs running



**Dedicated
IaaS
Europe**



30 Datacenters


Own
20Tbps
Network
with
35 PoPs

Hosting capacity :
**1.3M Physical
Servers**

**360k
Servers already
deployed**

> **1.4M Customers in 138 Countries**

OVHcloud: 4 Universes of Products



Domain / Email ▼

- Domain names, DNS, SSL, Redirect
- Email, Open-Xchange, Exchange
- Collaborative Tools, NextCloud

PaaS for Web ▼

- Mutu, CloudWeb
- Plesk, CPANEL
- PaaS with Platform.sh

Virtual servers ▼

- VPS, Dedicated Server

SaaS ▼

- Wordpress, Magento, Prestashop
- CRM, Billing, Payment, Stats
- MarketPlace

Support, Managed ▼

- Support Basic
- Support thought Partners
- Managed services

Standalone, Cluster ▼

- General Purpose
- SuperPlan
- Game T2 >20e
- Virtualization T3 >80e
- Storage T4 >300e
- Database T5 >600e
- HCI 12KVA /32KVA
- AI
- VDI Cloud Game
- Network

VPS aaS ▼

- pCC DC
- Virtuozzo Cloud

Wholesales ▼

- IT Integrators, Cloud Storage,
- CDN, Database, ISV, WebHosting
- High Intensive CPU/GPU,

Encrypt ▼

- KMS, HSM
- Encrypt (SGX, Network, Storage)

Compute ▼

- VM K8S, IA IaaS
- Baremetal PaaS for DevOps

Storage ▼

- File, Block, Object, Archive

Databases ▼

- SQL, noSQL, Messaging,
- Dashboard

Network ▼

- IP FO, NAT, LB, VPN, Router,
- DNS, DHCP, TCP/SSL Offload

Security ▼

- IAM, MFA, Encrypt, KMS

IA, DL ▼

- Standard Tools for AI, AI Studio,
- IA IaaS, Hosting API AI

Bigdata, ML, Analytics

- Datalake, ML, Dashboard

Hosted Private Cloud ▼

VMware

- SDDC, vSAN 1AZ / 2AZ
- vCD, Tanzu, Horizon, DBaaS,
- DRaaS

Nutanix

- HCI 1AZ / 2AZ, Databases,
- DRaaS, VDI

OpenStack

- IAM, Compute (VM, K8S)
- Storage, Network, Databases

Storage

- Ontap Select, Nutanix File
- OpenIO, MinIO, CEPH
- Zerto, Veeam, Atempo

AI

- ElementAI, HuggingFace,
- Deepomatic, Systran,
- EarthCube

Bigdata / Analytics / ML

- Cloudera over S3, Dataiku,
- Saagje, Tableau,

Hybrid Cloud ▼


- vRack Connect, Edge-DC, Private DC
- Dell, HP, Cisco, OCP, MultiCloud

Secured Cloud ▼

- GOV, FinTech, Retail, HealthCare

Orchestrating containers


Like herding cats... but in hard mode!




From bare metal to containers




Another paradigm shift



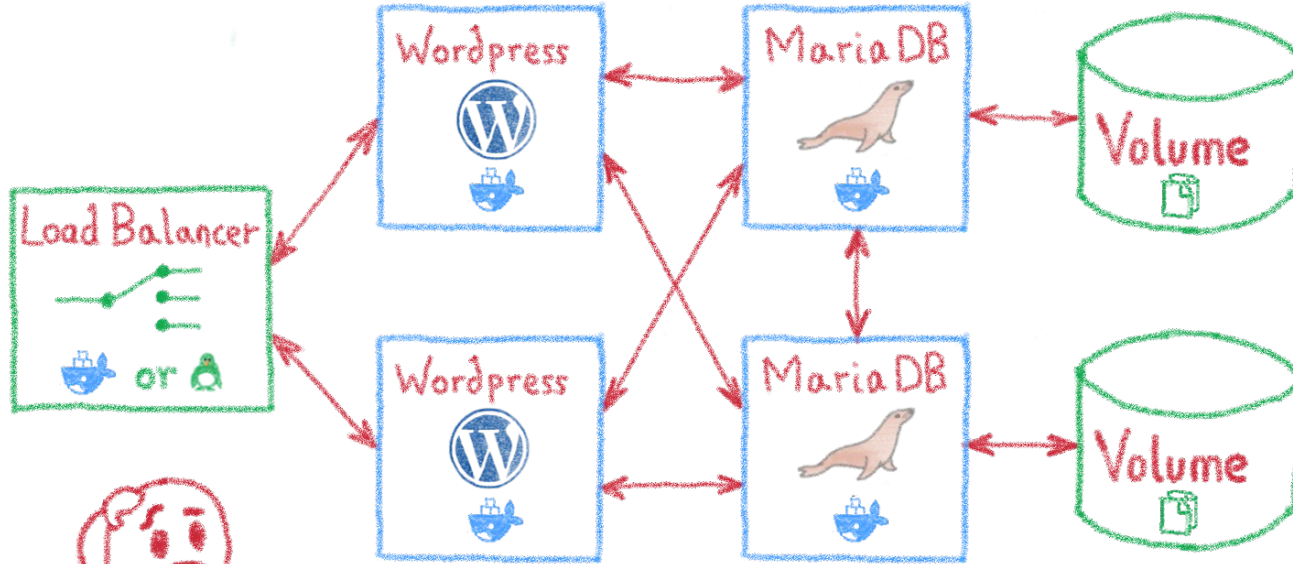
Containers are easy...



For developers




Less simple if you must operate them




Like in a production context


And what about microservices?



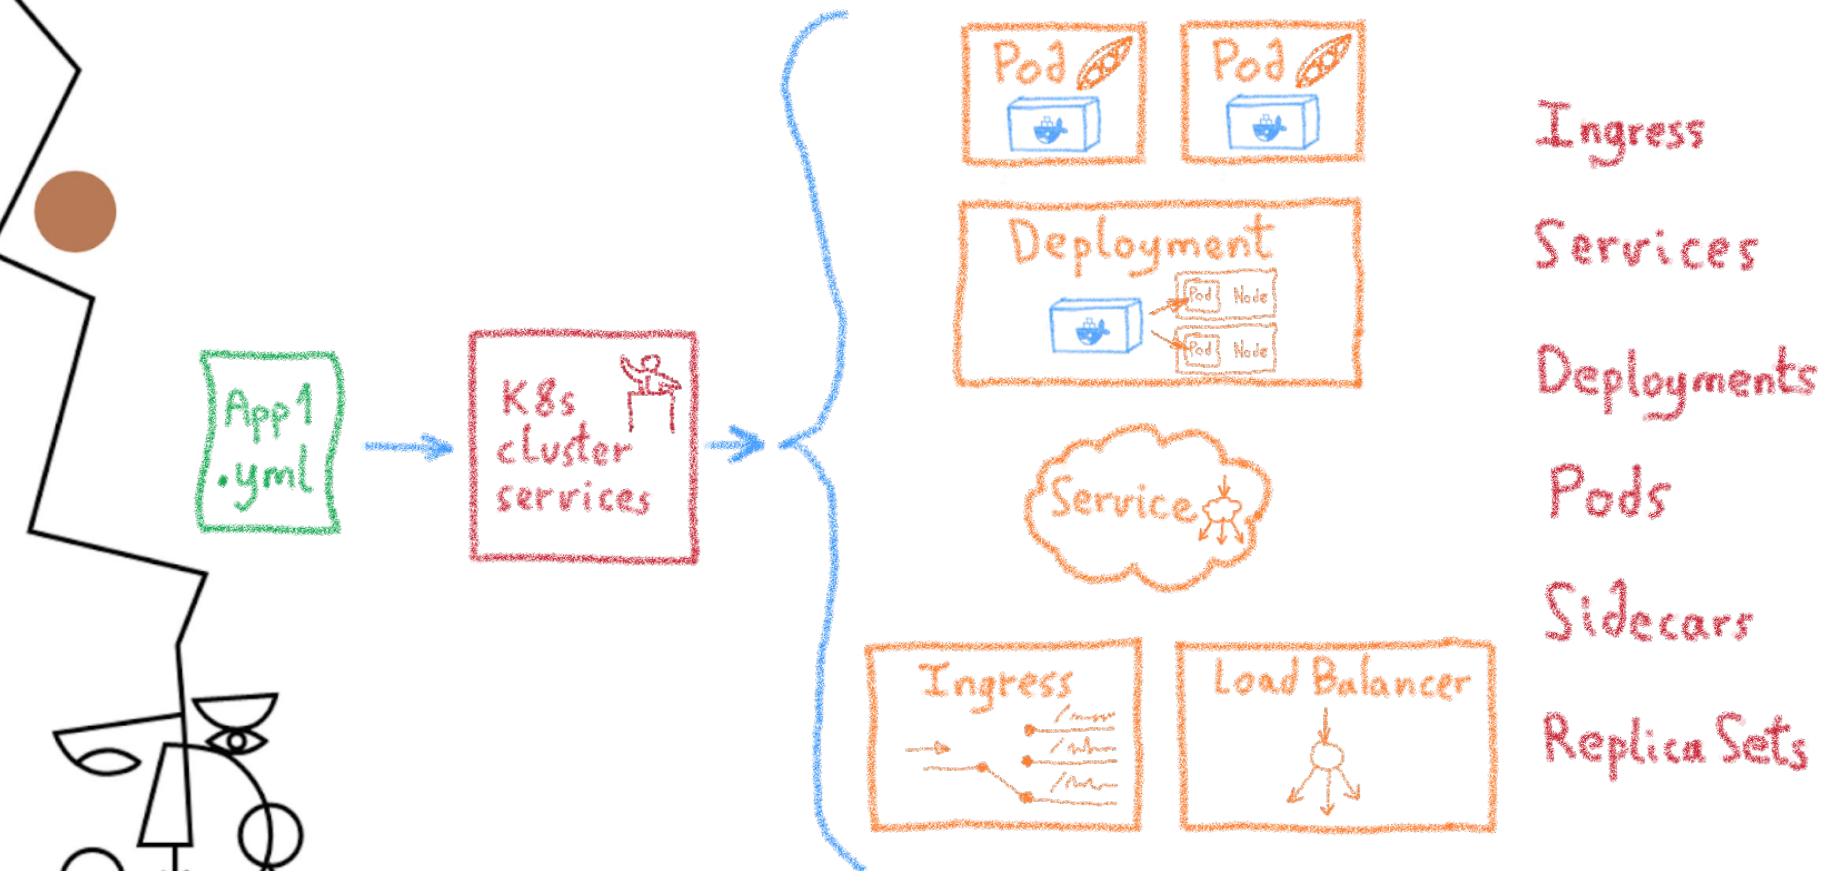
Are you sure you want to
operate them by hand?




Taming microservices with Kubernetes



Desired State Management



Having identical, software defined environments



Dev envs
Staging
Multi-cluster
Multi-cloud




I have deployed on Minikube, woah!

A great fastlane into Kubernetes




Running a full K8s in your laptop



A great learning tool


@LostInBrittany 

Your laptop isn't a true cluster




Don't expect real performances


Beyond the first deployment



So I've deployed my distributed architecture on K8s, everything is good now, isn't it?



Minikube is only the beginning




From Minikube to prod

A journey not for the faint of heart





@LostInBrittany 

Kubernetes can be wonderful


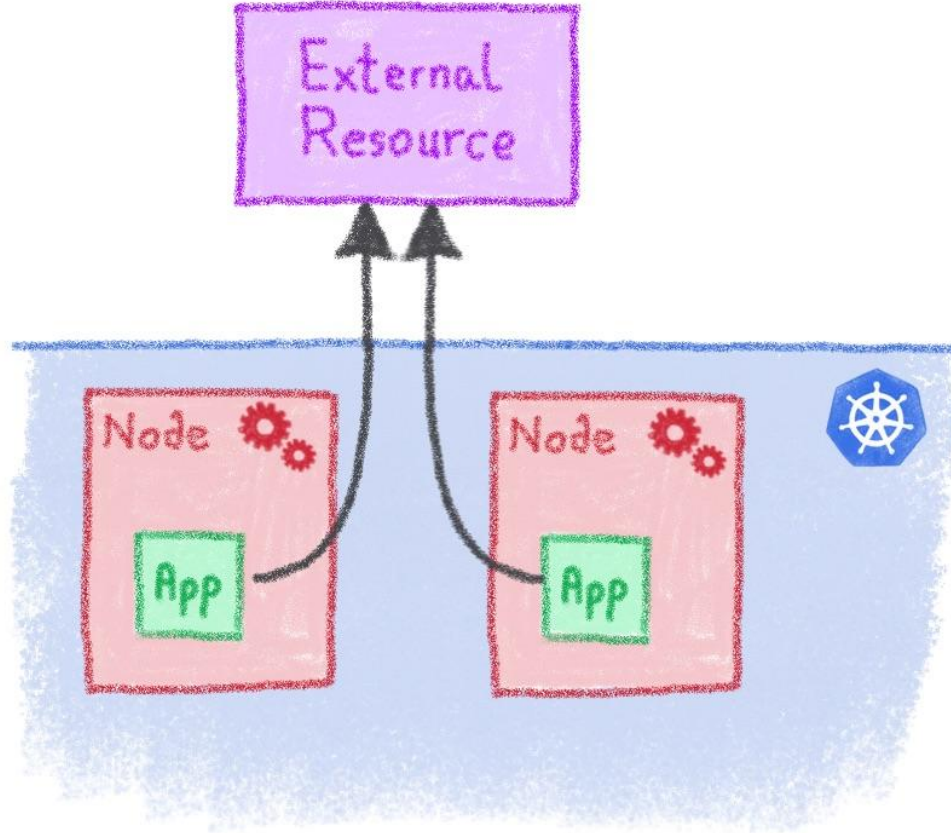


For both developers and devops


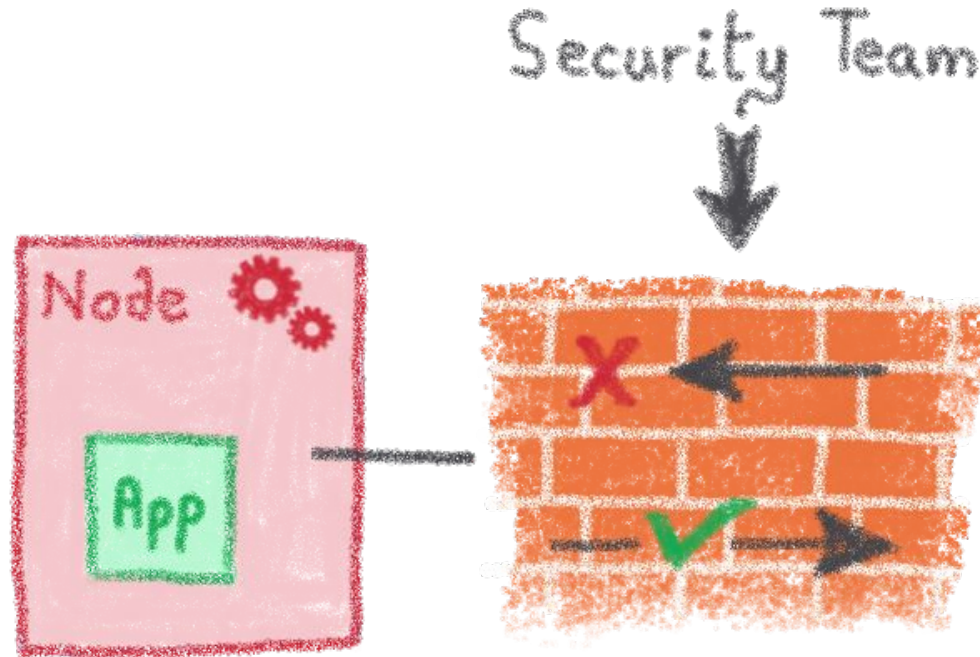
But it comes with a price...





An example among many others





An example among many others





An example among many others



An example among many others



An example among many others



The truth is somewhere inside...


What you see

Abstractions

The truth



A network example: KubeProxy



KubeProxy:
3 proxy modes

- Userspace
- IPTables
- IPVS




A network example: KubeProxy

--proxy-mode ProxyMode


Which proxy mode to use: 'userspace' (older) or 'iptables' (faster) or 'ipvs'. If blank, use the best-available proxy (currently **iptables**). If the iptables proxy is selected, regardless of how, **but the** system's kernel or iptables versions are insufficient, this always falls back to the userspace proxy.

iptables by default




A network example: KubeProxy


IPtables are based on rule chains



More rules → More time to insert or evaluate them



A network example: KubeProxy




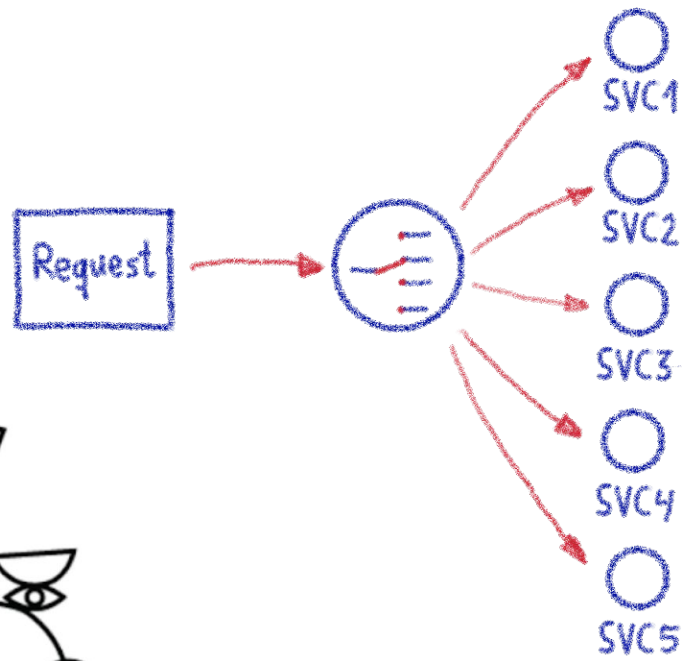
Cluster networking will be slower and slower



A network example: KubeProxy


IPVS is based on hash tables

Constant time to insert or evaluate them



IPVS to the rescue!

Kubernetes networking is complex...



All this traffic...
is it normal?



Network plugins (Flannel, Calico, Weave...)

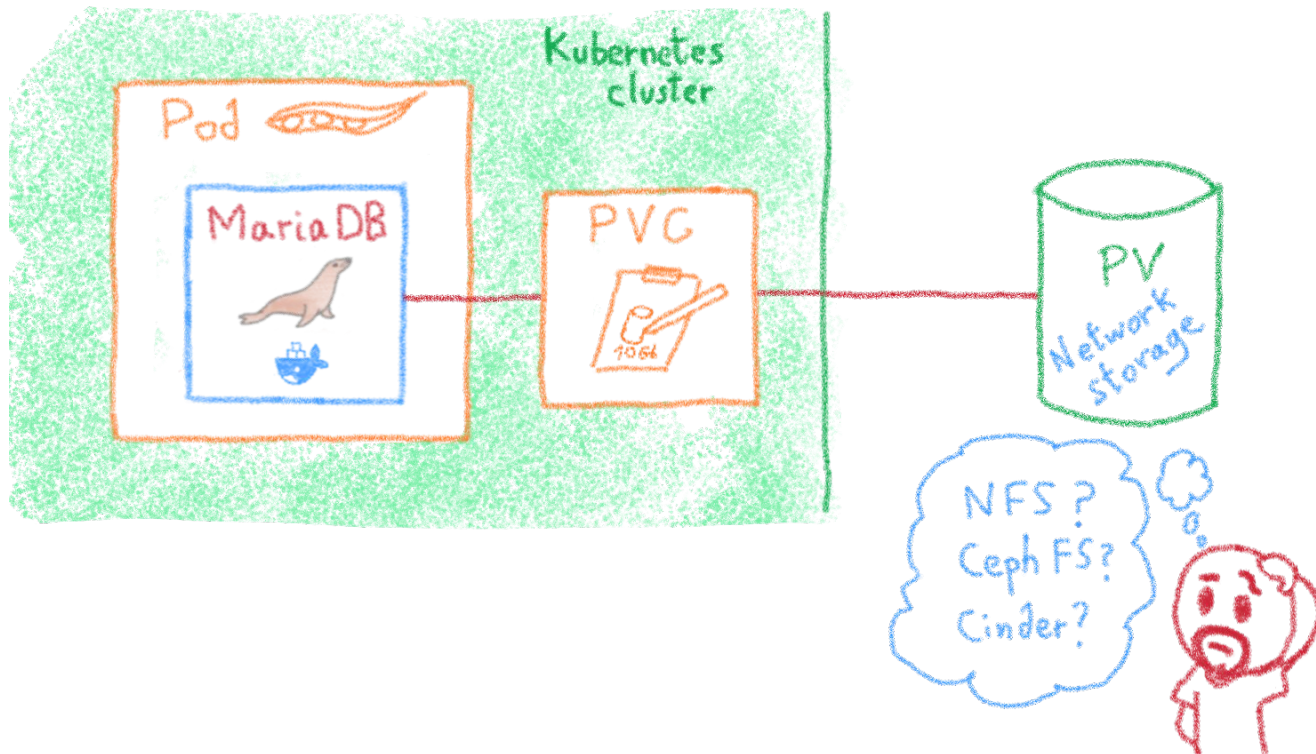
- IPAM - iptables
- routing - crossnode networking


Cluster IP, NodePort, Ingress

Service Meshes, Istio



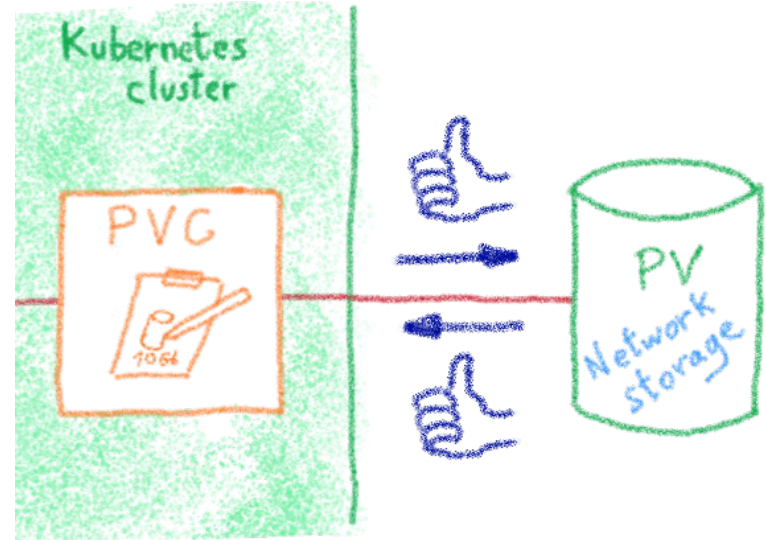
The storage dilemma

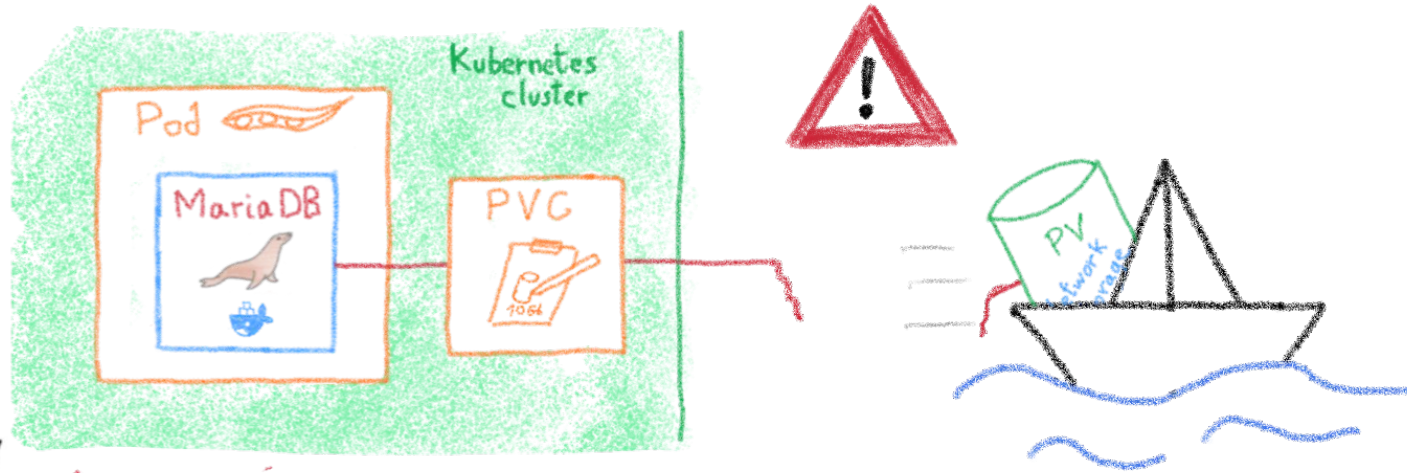




Volumes are handle through CSI
CSI provide an interface between
Kubernetes and storage technologie

Most CSI assume
perfect sync
between Kubernetes
and the storage
backend

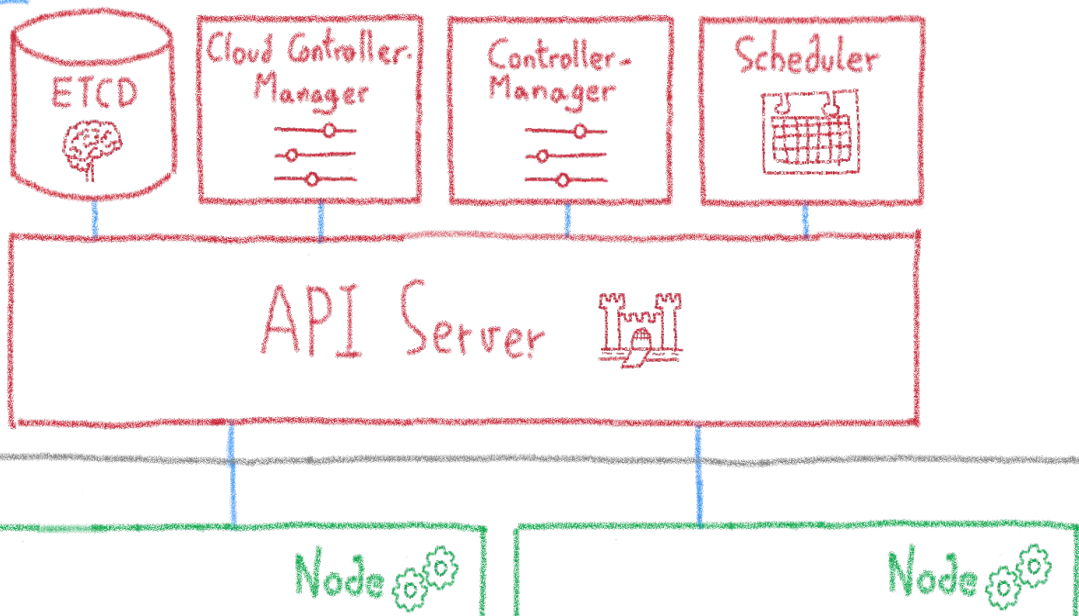




Storage backend are subject to errors or maintenance
Potential state shifts between storage and Kubernetes



The ETCD vulnerability

A single instance ETCD?
Are you sure?



Security

Hardening your Kubernetes



The security journey

Your security journey

Maturity

- Set up a cluster**
 - Restrict access to kubectl
 - Use RBAC
 - Use a Network Policy
 - Use namespaces
 - Bootstrap TLS
- Prevent known attacks**
 - Disable dashboard
 - Disable default service account token
 - Protect node metadata
 - Scan images for known vulnerabilities
- Follow security hygiene**
 - Keep Kubernetes updated
 - Use a minimal OS
 - Use minimal IAM roles
 - Use private IPs on your nodes
 - Monitor access with audit logging
 - Verify binaries that are deployed
- Prevent/limit impact of microservice compromise**
 - Set a Pod Security Policy
 - Protect secrets
 - Consider sandboxing
 - Limit the identity used by pods
 - Use a service mesh for authentication & encryption

Mattias Gees
@MattiasGees

Your security journey with Kubernetes by @MayaKaczorowski
#GoogleNext18


319 12:59 PM - Oct 11, 2018

Open ports (e.g. etcd 2379/TCP)
Kubernetes API (e.g. Tesla hacking)
Exploits (lots of CVEs)
RBAC (e.g. badly defined roles)


Are you kidding me?




Kubernetes is insecure by design*




It's a feature, not a bug.
Up to K8s admin to secure it according to needs



Not everybody has the same security needs



Kubernetes allows to enforce security practices as needed



Listing some good practices

- Close open access
- Define and implement RBAC
- Define and implement Network Policies
- Isolate sensitive workloads




Close open access




Close all by default, open only the needed ports
Follow the least privileged principle

Define and implement RBAC


RBAC: Role-Based Access Control





According to your needs




Define and implement network policies



Use RBAC and Network Policies to isolate your sensitive workload



Always keep up to date



Both Kubernetes and plugins

Because Kubernetes is a big target

Kubernetes » Kubernetes : Vulnerability Statistics

[Vulnerabilities \(22\)](#) [CVSS Scores Report](#) [Browse all versions](#) [Possible matches for this product](#) [Related Metasploit Modules](#)

[Related OVAL Definitions](#) : [Vulnerabilities \(0\)](#) [Patches \(0\)](#) [Inventory Definitions \(0\)](#) [Compliance Definitions \(0\)](#)


[Vulnerability Feeds & Widgets](#)

Vulnerability Trends Over Time

Year	# of Vulnerabilities	DoS	Code Execution	Overflow	Memory Corruption	Sql Injection	XSS	Directory Traversal	Http Response Splitting	Bypass something	Gain Information	Gain Privileges	CSRF	File Inclusion	# of exploits
2016	3										1	1			
2017	2										1				
2018	6									1					
2019	11	2									1				
Total	22	2								1	3	1			
% Of All		9.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	4.5	13.6	4.5	0.0	0.0	

Warning : Vulnerabilities with publish dates before 1999 are not included in this table and chart. (Because there are not many of them and they make the page look bad; and they may years.)

Vulnerabilities By Year



And remember, even the best can get hacked





One of Tesla's cluster got hacked
via an unprotected K8s API endpoint,
and was used to mine cryptocurrency...

Remain attentive, don't get too confident




Extensibility

Enhance your Kubernetes



Kubernetes is modular



Fully extensible


- Kubernetes API
- Cluster demons
- Controllers
- Custom resources
- ...

Operators


Let's see how some of those plugins can help you

Helm


A package management for K8s




Complex deployments



- Ingress
- Services
- Deployments
- Pods
- Sidecars
- Replica Sets
- Stateful Sets



Using static YAML files

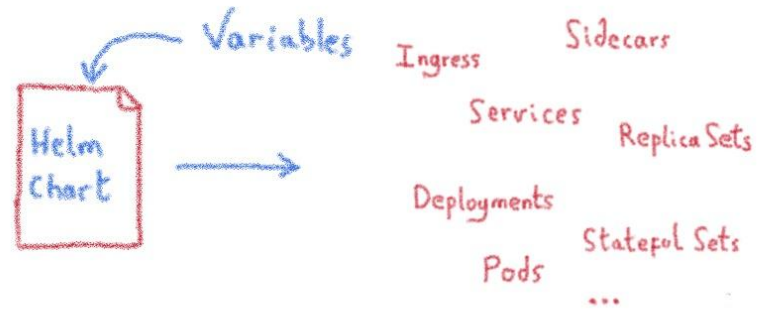


But if I need to customize things?





Complex deployments

A package manager for Kubernetes





- Manage complexity 
- Simple sharing 

- Easy upgrades 
- Easy rollbacks 







Istio

**A service mesh for Kubernetes...
and much more!**



Istio: A service mesh... but not only



Rolling upgrades

A/B Testing


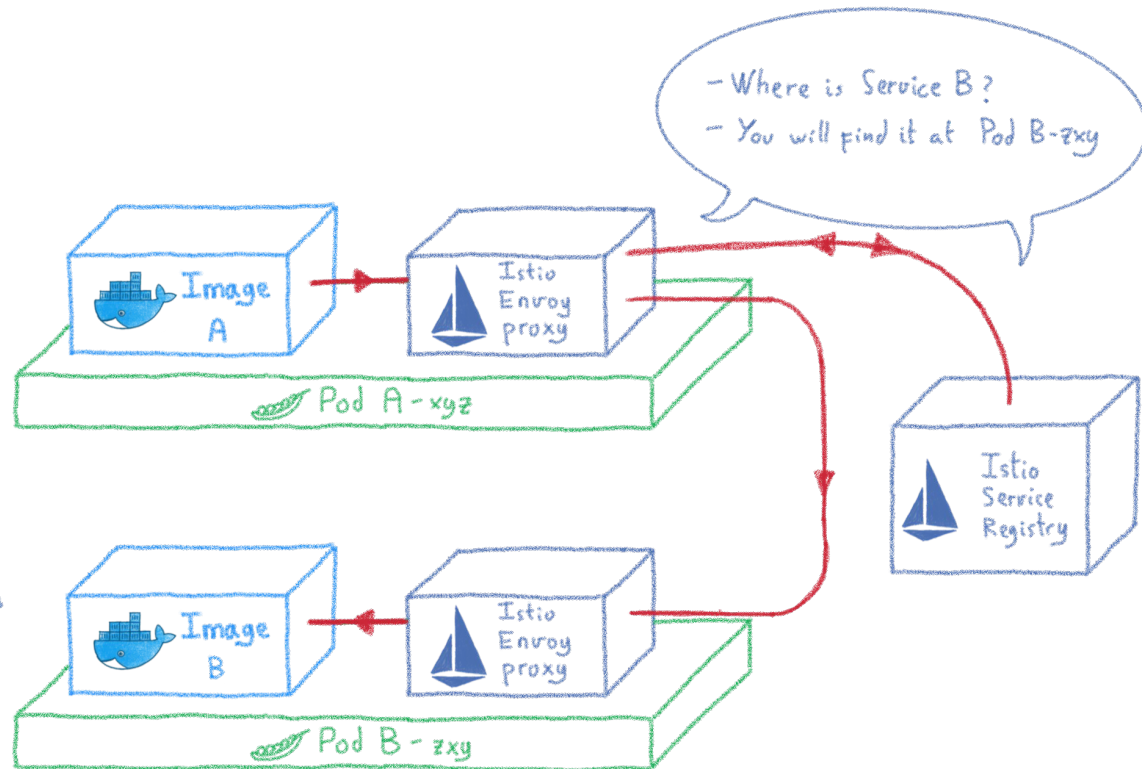
Canary Testing

Edge traffic management

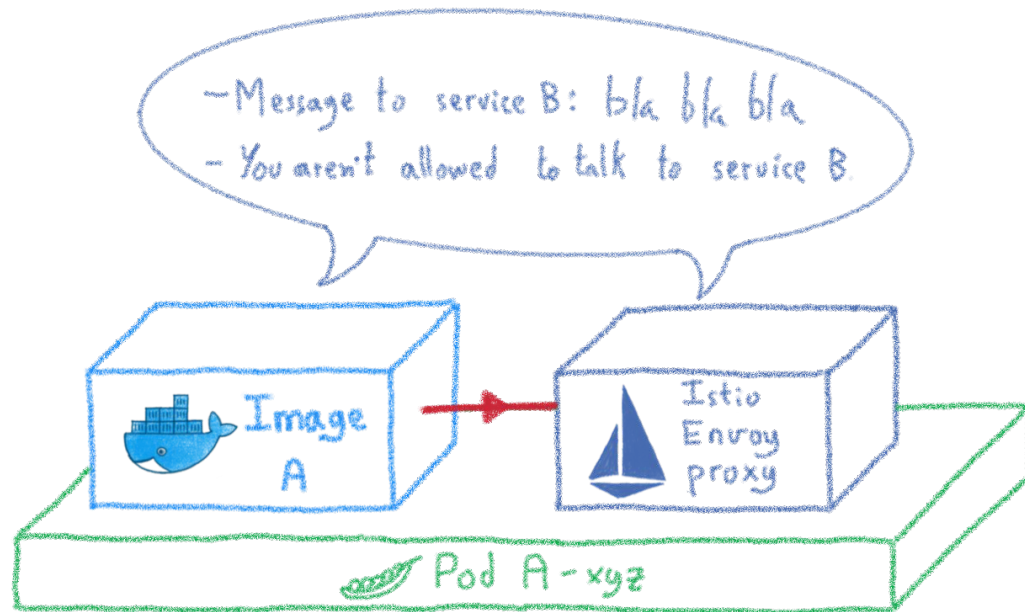
Multiclustert service mesh



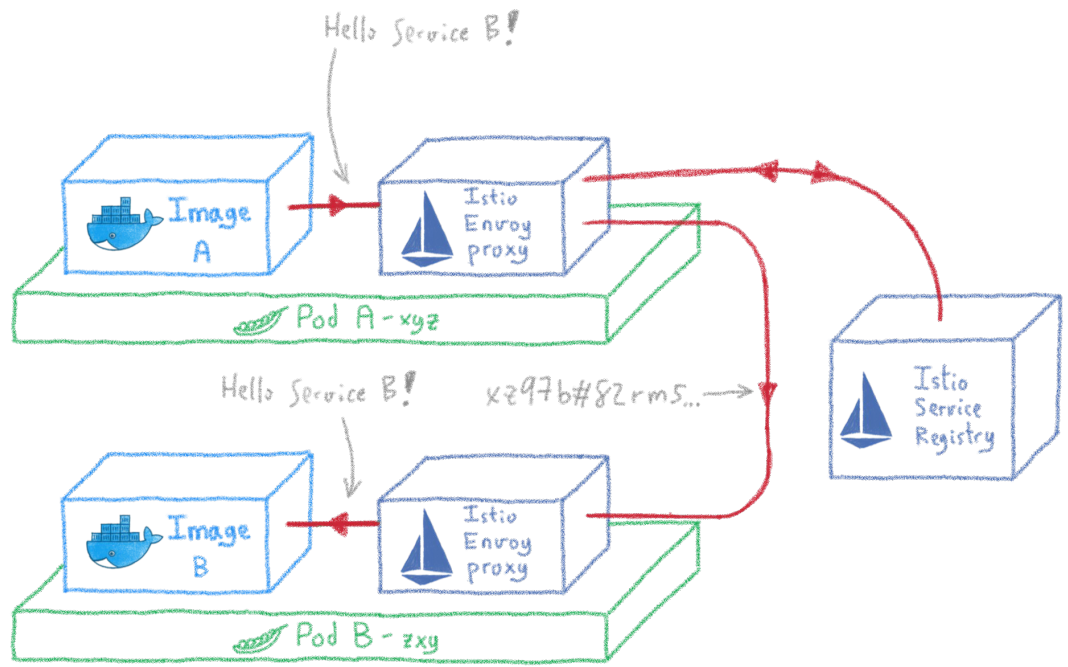


Service discovery



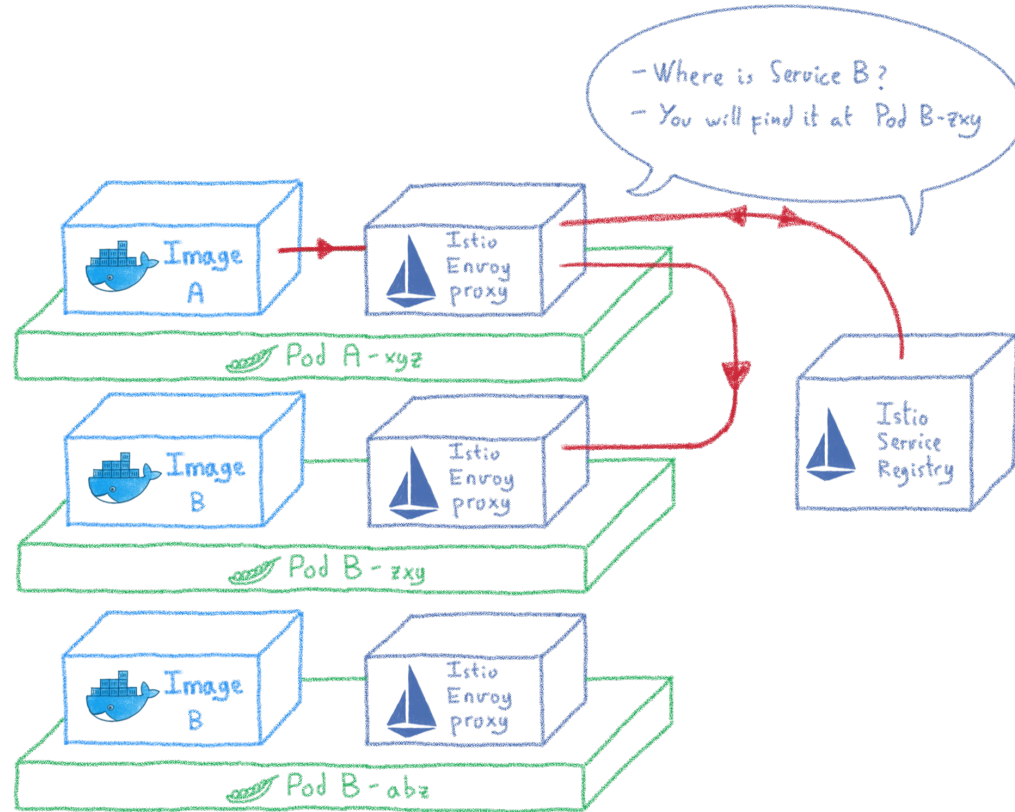

Traffic control






Encrypting internal communications






Routing and load balancing






Rolling upgrades





Rolling upgrades




A/B testing



Monitoring your cluster




- Metrics
 - Logs
 - Tracing
- } at {
- Envoy level
 - Control plane level



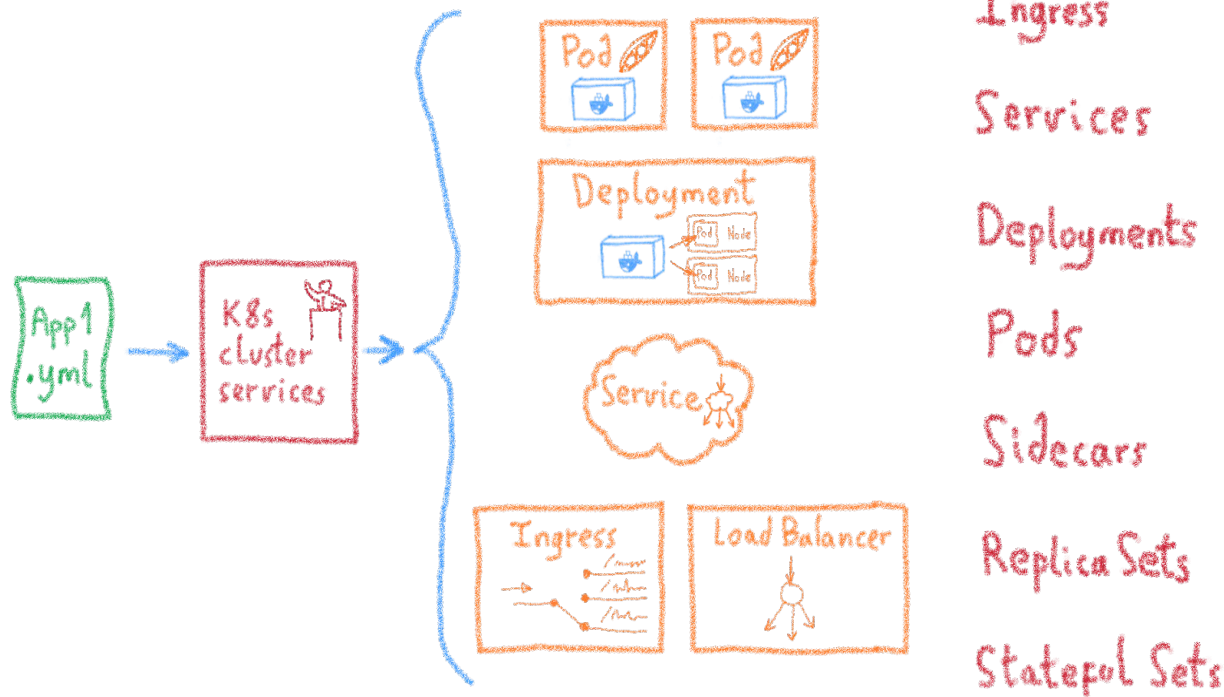
Dashboards

Velero


Backing up your Kubernetes



Kubernetes: Desired State Management



YAML files allows to clone a cluster



Dev envs



Staging

Multi-cluster

Multi-cloud



But what about the data?




Velero




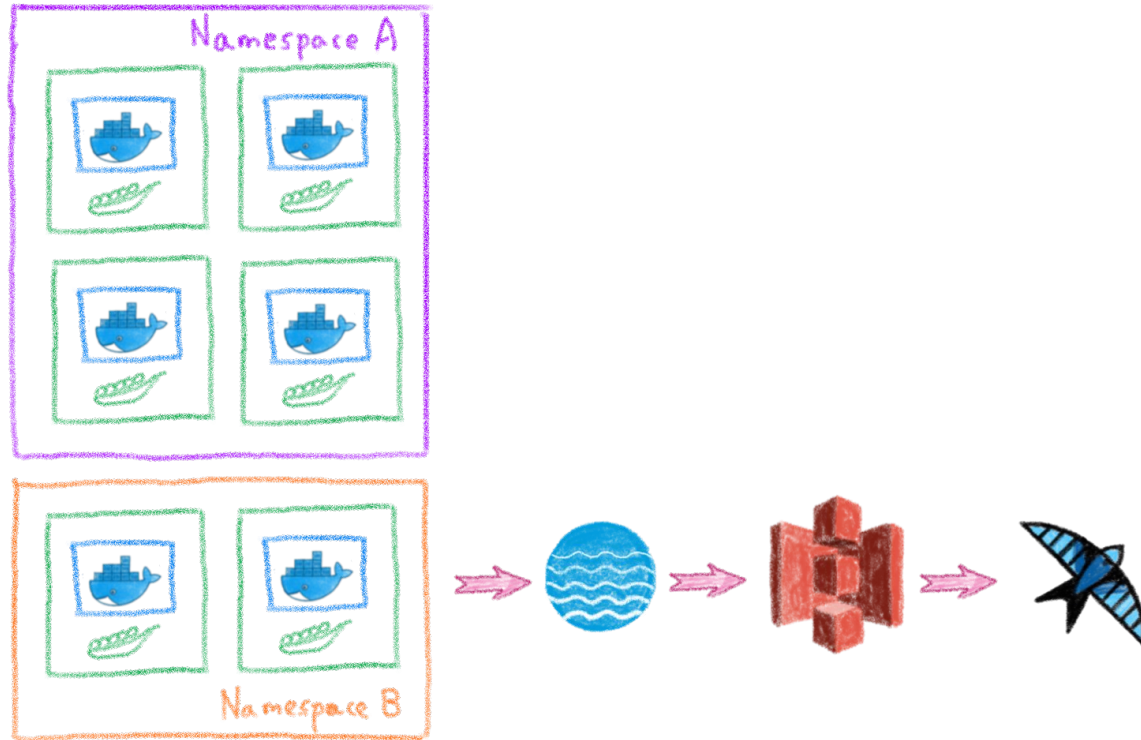
Backup and migrate Kubernetes applications
and their persistent volumes

S3 based backup





On any S3 protocol compatible store



Backup all or part of a cluster



Schedule backups




Backups hooks




Conclusion

And one more thing...




Kubernetes is easy to begin with




Minikube, K3s...

@LostInBrittany 


Kubernetes is powerful



It can make Developers' and DevOps' lives easier



But there is a price: operating it



Lot of things to think about

We have seen some of them



Security



Deployment



Monitoring




Backups





Different roles



Each role asks for very different knowledge and skill sets



Operating a Kubernetes cluster is hard




But we have a good news...


Most companies don't need to do it!




As they don't build and rack their own servers!



If you don't need to build it, choose a certified managed solution



You get the cluster, the operator
get the problems




Like our OVH Managed Kubernetes



Made with  by the Platform team

@LostInBrittany 



CLOUD
CONNECT 2021
by  CIONET



Thank you for listening!

