



**CLOUD
NATIVE
CON**
Europe 2017



KubeCon
A CNCF EVENT



Kubernetes from Dev to Prod

Subhas Dandapani, CI/CD Platform Lead, GoEuro



**CLOUD
NATIVE
CON**
Europe 2017



KubeCon
A CNCF EVENT



GOEURO



12 countries

30k+ stations

200+ airports

10+ million monthly searches



Subhas Dandapani

CI/CD Platform Lead
medium.com/@rdsbhas
github.com/rdsbhas



**CLOUD
NATIVE
CON**
Europe 2017



KubeCon
A CNCF EVENT



Transition Journey

- 50 engineers, 10 services → 150 engineers, 50 services
- 10 coordinated releases → 150+ independent releases per week
- Legacy → CI/CD/Docker/K8s
Excluding persistent services
- Through all kinds of legacy problems
Integration patterns: REST, File, MQ, DB-driven, Metrics-driven, Compile-time, Ancillary artifacts
Dev hacks for ops issues, Ops hacks for dev issues
- Transition in 4 months across company
(ground work laid before in a pilot project)



CLOUD
NATIVE
CON
Europe 2017



KubeCon
A CNCF EVENT



Development

GOEURO





CLOUD
NATIVE
CON
Europe 2017

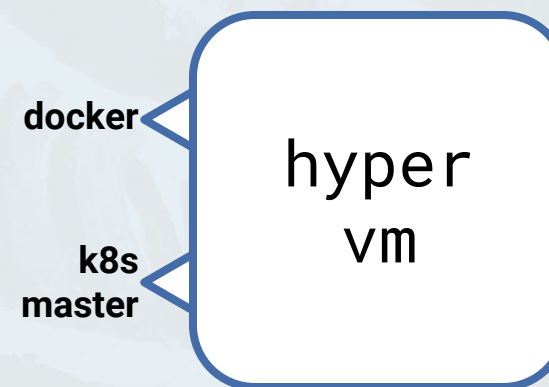


KubeCon
A CNCF EVENT



From the bottom up: hyper-vm

- Single-node k8s+docker VM for development
- docker build & kubectl apply on same node
- Metadata based kubernetes customization (e.g. installing services at bootup)
- Optimized for development & testing in the cloud
docker bind mounting (nested resolv, etc), integrate with internal network, build-arg-activated docker server patches (build caching), etc
- First point to upgrade kubernetes from dev to prod





CLOUD
NATIVE
CON
Europe 2017



KubeCon
A CNCF EVENT



Development

service-repo

└ ops

└ Dockerfile


└ helm chart

└ ci.yaml

└ Makefile

- Every service has "ops" folder
- Self-contain most dev+ops requirements

docker build



docker

k8s
master

hyper
vm



CLOUD
NATIVE
CON
Europe 2017



KubeCon
A CNCF EVENT



y8s (yaml-k8s-sharing)

service-repo

└ **ops**

└ Dockerfile

└ **helm chart**

└ ci.yaml

- Adopt Helm Chart standard/layout/templating
- But Render-as-a-service
- Additional simplification:
 - 1 input: "env", configure everything else out of the box
 - env=dev: use local image
 - env=test: use stubs (provider-driven stubs/mocks)
 - env=qa/preprod/prod: pin revision
 - Optional params (overrides, etc)



CLOUD
NATIVE
CON
Europe 2017



KubeCon
A CNCF EVENT



y8s (yaml-k8s-sharing)

```
kubectl apply -f  
http://y8s.int/repos/any-service-repo  
?env=<test|dev|qa|...>  
&values[foo]=bar&filters[kind]=Service&...
```

- kubectl apply **any repo** →
- DevOps services also applied/deployed same way (router, ingress, core services, database, couchbase, etc)
- Additional REST options (values, filters, validate, dry run, json/yaml content types, etc) and generated metadata (for tracing back to source)

docker

k8s
master

hyper
vm



**CLOUD
NATIVE
CON**
Europe 2017



KubeCon
A CNCF EVENT



Landscape

Develop

ops/chart in repo
self contained repos

y8s
consume any repo/service

hyper-vm
develop in the cloud

Integrate

Deploy



CLOUD
NATIVE
CON
Europe 2017



KubeCon
A CNCF EVENT



Integration

GOEURO





CLOUD
NATIVE
CON
Europe 2017



KubeCon
A CNCF EVENT



CI Contract

service-repo

└ **ops**

└ Dockerfile

└ helm chart

└ **ci.yaml**

- Adopt [GitLab CI.yaml](#) syntax/design
- But parse/implement it ourselves
- Automate all CI pipelines/jobs
- We orchestrate all triggers, do all plumbing
pr checks, manual checkpoints, cron jobs, distributed caching, nested
docker builds, environments, retry/reassignment rules, slack notifications,
emails, release annotations/graphs, etc



CLOUD
NATIVE
CON
Europe 2017



KubeCon
A CNCF EVENT



ci.yaml

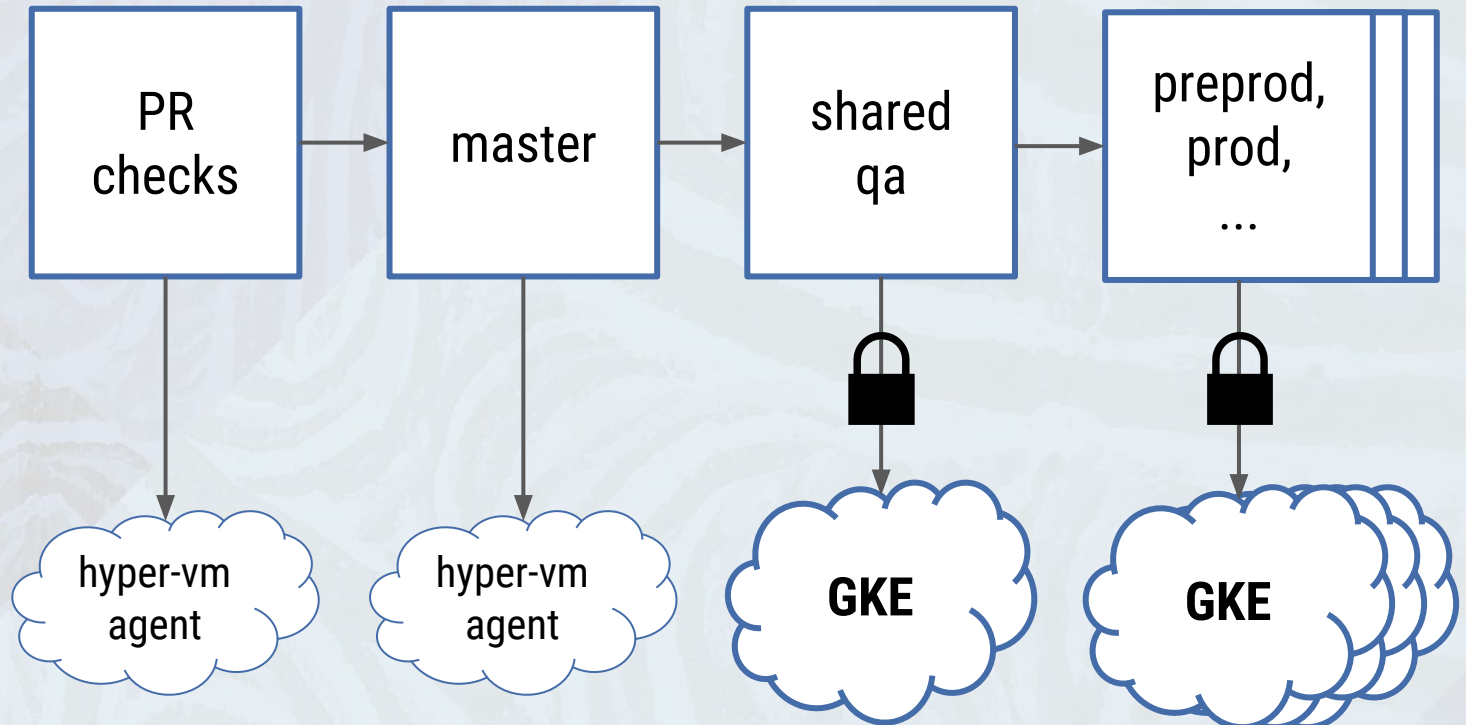
```
image: jenkins-plain:latest
```

```
unit-tests:  
  stage: test  
  script: test.sh
```

```
master:  
  script: release.sh
```

```
deploy-qa:  
  environment:  
    name: qa  
    credentials: "..."  
  script: deploy.sh
```

```
deploy-preprod:  
deploy-prod:  
  ...
```





CLOUD
NATIVE
CON
Europe 2017

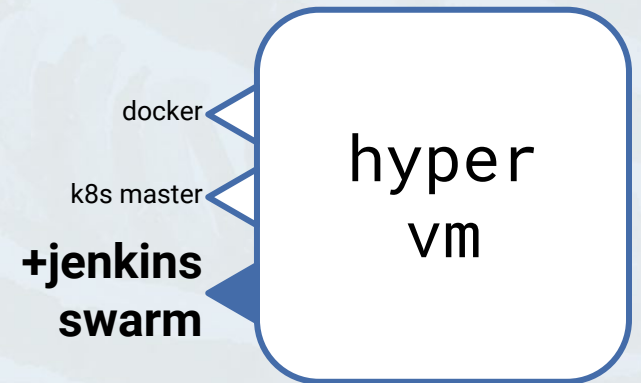


KubeCon
A CNCF EVENT



CI Contract

- Every agent is hyper-vm + jenkins slave
- docker/kubectl readily configured
- Jobs run inside configurable build images
- Recreate from source every 2 hours
with no downtime, no possibility of snowflake CI
- Encourage early integration tests
`kubectl apply -f y8s.g.int/repos/anything` and test
- **Take Jobs all the way to Prod** (safety in upcoming slides)





CLOUD
NATIVE
CON
Europe 2017



KubeCon
A CNCF EVENT



Job Environment - Closer Look #1

```
image: jenkins-plain:latest
```


```
my-job:
```

```
  script: make build
```

- Start using docker/kubectl right away
- Cleaned up after every build
- Early Integration/E2E tests

```
kubectl apply -f "y8s.int/...dependencies..."
```

```
kubectl run / docker run <e2e tests>
```





CLOUD
NATIVE
CON
Europe 2017



KubeCon
A CNCF EVENT



Job Environment - Closer Look #2


my-job:

script: make deploy

environment:

name: preview-24h

- Provision on-the-fly ephemeral hyper-vm
Same build contract, kubectl/docker readily pointed to that VM
- But lives after job completes
- Show demos, connect for development/debugging





**CLOUD
NATIVE
CON**
Europe 2017



KubeCon
A CNCF EVENT



Job Environment - Closer Look #3

my-job:

script: make deploy

environment:

name: qa|preprod|prod-region

credentials: <credentials>

- Apply to Deploy! Including prod regions
- We'll talk about security in later slides





CLOUD
NATIVE
CON
Europe 2017



KubeCon
A CNCF EVENT



Composition

my-job:

stage: pr

when: manual

environment:

name: preview-24h

Run on PRs

when: manual

on pipelines: manual checkpoint
on PRs: comment "trigger my-job"

And boot an ephemeral preview VM

Product demos on demand,
GitHub Flow (deploy from PR to prod), etc



CLOUD
NATIVE
CON
Europe 2017



KubeCon
A CNCF EVENT



Patterns

- Database migrations
 - `kubectl run --overrides=$(curl -H ...json y8s.g.int/...migration)`
- End to end testing
 - Shared E2E tests treated as services, have their own ops folder as well
 - `kubectl apply y8s.g.int/repos/... && kubectl run <e2e-tests>`
- Additional Services
 - helm templates for services like selenium, couchbase, mysql, etc
 - `kubectl apply y8s.g.int/...services...`
 - No need for additional abstraction (like travis-ci services, for example)



**CLOUD
NATIVE
CON**
Europe 2017



KubeCon
A CNCF EVENT



Landscape

Develop

ops/chart in repo
self contained repos

y8s
consume any repo/service

hyper-vm
develop in the cloud

Integrate

ci.yaml
pipelines & patterns

hyper-vm + ci agent
early E2E in CI

ready environments
preview, qa, preprod, prod, ...

Deploy



CLOUD
NATIVE
CON
Europe 2017



KubeCon
A CNCF EVENT



Deployment

GOEURO





CLOUD
NATIVE
CON
Europe 2017




KubeCon
A CNCF EVENT



auth8s

- Kubernetes API gateway/proxy
 - Deployed in every k8s cluster, proxies to kubernetes master
- Security
 - Each namespace gets its own username/password
 - k8s 1.6 RBAC is in "beta", this security part can be offloaded to k8s when its stable
- Compliance
 - Validate and enforce CPU/Memory Limits, Live/Readiness probes, ...
 - Necessary annotations on every resource, whitelisted resource kinds, ...
 - Much more using HTTP interception, Necessary for Governance and Standards
 - Supports most features (except interactive shells)





**CLOUD
NATIVE
CON**
Europe 2017



KubeCon
A CNCF EVENT



Value additions: Cluster services

- Auto-register DNS for independent clusters
- Metrics: heapster, restart counters, etc
- Logging: Ship structured container logs using filebeat
- Routing: nodePort, containerizing current router, etc
 - Moving to: Ingress
 - Automatically instrument external health checks (pingdom-style)



**CLOUD
NATIVE
CON**
Europe 2017



KubeCon
A CNCF EVENT



Value additions: Sidecar and Regional services

- Sidecar container for big artifacts for legacy services
- Region-specific external services
 - All devops-managed external resources (db, logstash, etc) have their own pipelines and yamls
 - Provide containerized/test in dev/qa/preprod/etc
 - Endpoints pointing to persistent resources in prod
- Multi-region focused from the beginning



**CLOUD
NATIVE
CON**
Europe 2017



KubeCon
A CNCF EVENT



Value additions: Annotation driven

- Coming up: More annotation-driven cluster services
 - Capture periodic thread-dumps by declaring annotations
 - Autoscale by declaring metric queries in annotations
 - Automated external uptime checks and status pages for ingresses
 - etc



**CLOUD
NATIVE
CON**
Europe 2017



KubeCon
A CNCF EVENT



Landscape

Develop

ops/chart in repo
self contained repos

y8s
consume any repo/service

hyper-vm
develop in the cloud

Integrate

ci.yaml
generated pipelines & jobs

hyper-vm + ci agent
early E2E in CI

ready environments
preview, qa, preprod, prod, ...

Deploy

gke
best kubernetes experience

auth8s
security, compliance

value added services
routing, logs, metrics, dns, ...



CLOUD
NATIVE
CON
Europe 2017



KubeCon
A CNCF EVENT



Learnings

GOEURO





**CLOUD
NATIVE
CON**
Europe 2017



KubeCon
A CNCF EVENT



What we learned?

- Containers are a cultural change
 - Lots of presentations, PoCs, consulting, lobbying, “we needed this yesterday”
 - We did at same time: CI, CD, Unversioning, Dockerizing, Kubernetesing
 - Couldn't have done without support from teams, managers, product, business
 - Don't try too much: Changing monitoring, routing, building, processes, etc.
- Migrate as-it-is
 - Services closer to 12factor, With small dep-graph, Hard legacy challenges
 - 5 minute bootup, 400 MB images, Dev vs Ops hacks, etc
 - Run with 0% traffic, Shadow comparison, QA/Preprod, Traffic switching
 - Establish the pipe, Optimize later



**CLOUD
NATIVE
CON**
Europe 2017



KubeCon
A CNCF EVENT



What Kubernetes Enabled Us

- Transition from Cloud resources => Container resources
 - From Cloud: Disk Image + InstanceGroup/ASG + LoadBalancer
 - To Container: Docker Image + K8s Deployment + K8s Service/Ingress
 - From: packer build, terraform apply
 - To: docker build, kubectl apply
- Allow teams to deliver from Dev to Prod
 - Balance safety, compliance, automation and ease-of-use
- Transition legacy and modern services
 - Integrate and iterate fast, next 1000 releases per week



**CLOUD
NATIVE
CON**
Europe 2017

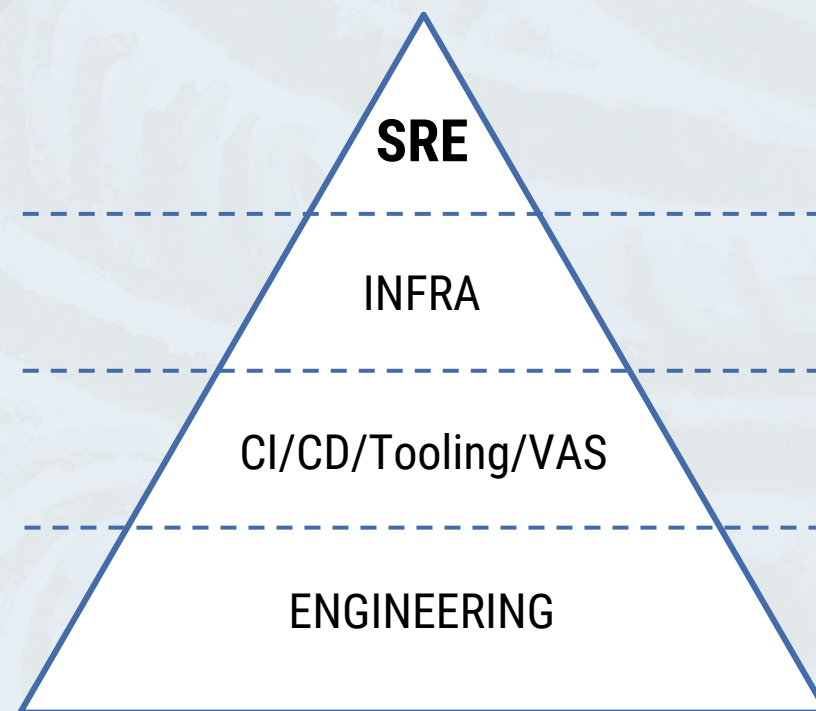


KubeCon
A CNCF EVENT



Stepping Forward

- Site Reliability Engineering
 - Teams deliver up to prod-x%
 - Rollout to rest automatically based on metrics
- Better State Management
- Secret Management
 - Multiple possibilities, ci.yaml, auth8s, etc
- Ingress, Federation
- Namespace Quotas
 - Virtual data center, Team budgets, etc





CLOUD
NATIVE
CON
Europe 2017



KubeCon
A CNCF EVENT



Questions

GOEURO

