

please

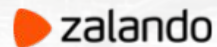
Ask questions  
through the app



*Rate Session*

Thank you!





# WHY I LOVE KUBERNETES FAILURE STORIES

---

GOTO BERLIN

2019-10-24

HENNING JACOBS

@try\_except\_



# ROLLING OUT KUBERNETES?

**Y** **Hacker News** [new](#) | [threads](#) | [past](#) | [comments](#) | [ask](#) | [show](#) | [jobs](#) | [submit](#)

▲ Ask HN: Do's/don'ts of working with Kubernetes you learned through experience?

33 points by fiddlerINT 1 day ago | [flag](#) | [hide](#) | [past](#) | [web](#) | [favorite](#) | 26 comments

*"We are rolling out Kubernetes to production next month and I'm interested to hear from people who made that step already."*

# DON'T USE IT !!!!!

**Y** **Hacker News** [new](#) | [threads](#) | [past](#) | [comments](#) | [ask](#) | [show](#) | [jobs](#) | [submit](#)

▲ Ask HN: Do's/don'ts of working with Kubernetes you learned through experience?

33 points by fiddlerINT 1 day ago | [flag](#) | [hide](#) | [past](#) | [web](#) | [favorite](#) | 26 comments

We are rolling out Kubernetes to production next month and I'm interested to hear from people who made that step already.

▲ iamnothere123 5 hours ago [-]

**DON'T USE IT !!!!!**

[reply](#)

▲ anon284271 19 hours ago [-]

**Don't use Kubernetes.**

[reply](#)

# KUBERNETES FAILURE STORIES



François Zaninotto

@francoisz



Kubernetes Failure Stories. The fact that this list has a dedicated website is a serious symptom of the complexity problem. [#Docker](#) [#DevOps](#) [k8s.af](#)



[Tweet übersetzen](#)

## Kubernetes Failure Stories

A compiled list of links to public failure stories related to Kubernetes. Most recent publications on top.

- 10 Ways to Shoot Yourself in the Foot with Kubernetes, #9 Will Surprise You - Datadog - KubeCon Barcelona 2019
  - involved: CoreDNS, `ndots:5`, IPVS conntrack, `imagePullPolicy: Always`, DaemonSet, NAT instances, `latest` tag, API server `OOMKill`, kube2iam, cluster-autoscaler, PodPriority, audit logs, `spec.replicas`, AWS ASG rebalance, CronJob, Pod toleration, zombies, `readinessProbe.exec`, cgroup freeze, kubectl
  - impact: unknown, API server outage, pending pods, slow deployments
- How Spotify Accidentally Deleted All its Kube Clusters with No User Impact - Spotify - KubeCon Barcelona 2019
  - involved: GKE, cluster deletion, browser tabs, Terraform, global state file, git PRs, GCP permissions
  - impact: no impact on end users
- Kubernetes Failure Stories - Zalando - KubeCon Barcelona 2019
  - involved: Skipper-Ingress, AWS, `OOMKill`, high latency, CronJob, CoreDNS, `ndots:5`, etcd, CPU throttling
  - impact: multiple production outages
- Oh Sh\*t! The Config Changed! - Pusher - KubeCon Barcelona 2019
  - involved: AWS, nginx, ConfigMap change
  - impact: production outage
- Misunderstanding the behaviour of one templating line - Skyscanner - blog post 2019
  - involved: HAProxy-Ingress, Service VIPs, Golang templating

# ZALANDO AT A GLANCE

**~ 5.4** billion EUR  
revenue 2018

**> 300**  
million

visits  
per  
month

**> 400,000**  
product choices

**~ 14,000**  
employees in  
Europe

**> 80%**  
of visits via  
mobile devices

**> 28**  
million  
active customers

**> 2,000**  
brands

**17**  
countries

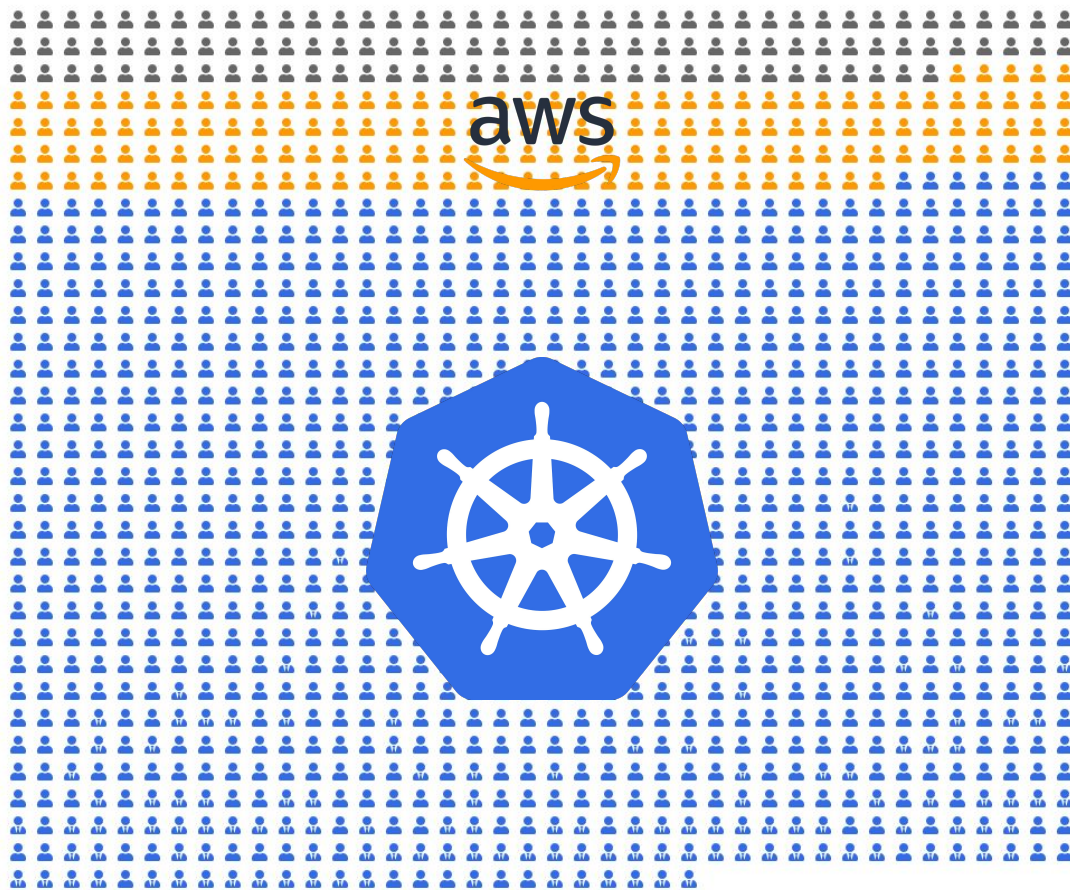
## 2019: SCALE

396 Accounts



140 Clusters

# 2019: DEVELOPERS USING KUBERNETES



zalando-incubator / kubernetes-on-aws			Unwatch 33	Unstar 250	Fork 60
Code	Issues 13	Pull requests 6	Actions	Insights	Settings
Branch: dev - kubernetes-on-aws / cluster / manifests /			Create new file	Upload files	Find file
Merge pull request #2084 from zalando-incubator/update/ingress-td			Latest commit 16d5ee 9 days ago		
..					
01-platformcredentialsset	PCS: validate application name	17 days ago			
01-vertical-pod-autoscaler	Updated VPA to version 0.4.0 and associated objects	a month ago			
01-visibility	ZMON: use a user-defined priority class	9 months ago			
admission-control	Update admission-controller & proxy	4 days ago			
audittrail-adapter	Hostnetwork will take resolv.conf from host	2 months ago			
cadvisor	ndots for kube-system	2 months ago			
cluster-lifecycle-controller	Update CLC to master-4	2 months ago			
coredns-local	Update CoreDNS to v1.4.0	2 months ago			
cron	add cron namespace to all cluster, such that we can introduce best pr...	2 years ago			
dashboard	ndots for kube-system	2 months ago			
default-limits	Use the correct feature flag in default-limits	3 months ago			
efs-provisioner	ndots for kube-system	2 months ago			
emergency-access-service	Update EAS	a month ago			
etcd-backup	ndots for kube-system	2 months ago			
external-dns	Updated the VPAs to v1beta2	24 days ago			
fannel	Update fannel-awaiter	a month ago			
heapster	Updated the VPAs to v1beta2	24 days ago			
infrastructure-secrets	Add secret with cluster-inf secrets to default ns	a year ago			
ingress-controller	update to hofix release and remove quiet flag	a day ago			
ingress-template-controller	Put ingress-template-controller behind feature toggle to gradually de...	a month ago			
kube-cluster-autoscaler	Add support for customizable AZs	16 days ago			
kube-dns-metrics	ndots for kube-system	2 months ago			
kube-downscaler	kube-downscaler v0.12	a month ago			
kube-jenitor	kube-jenitor v0.7	19 days ago			
kube-job-cleaner	Add a feature toggle for disabling kube-job-cleaner	8 days ago			
kube-metrics-adapter	Update to master-31	7 days ago			
kube-node-ready	Allow updating kube-node-ready/kube-proxy	5 days ago			
kube-proxy	Allow updating kube-node-ready/kube-proxy	5 days ago			
kube-state-metrics	Add VPA to kube-state-metrics	7 days ago			
kube-static-egress-controller	remove debug logging to reduce logs and fix restart problem caused by...	2 days ago			
kube-system-system	Replace secretary with a 'static' docker config,	2 years ago			
kube2iam	ndots for kube-system	2 months ago			
kubernetes-lifecycle-metrics	Updated the VPAs to v1beta2	24 days ago			
logging-agent	Fluentd config; turn off S3 bucket checks	23 days ago			
metrics-server	Update metrics-server to v0.3.2, use RBAC	22 days ago			
nvidia	Increase nvidia-driver-installer yet again	25 days ago			
pdb-controller	ndots for kube-system	2 months ago			
prometheus-node-exporter	Hostnetwork will take resolv.conf from host	2 months ago			
prometheus	Update to Prometheus v2.9.2	9 days ago			
ppp	disallow privilege escalation for restricted policy	5 months ago			
roles	Squash all commits from rbac branch	5 months ago			
skipper	add 4x the current buffer size which is less than 1M in total	23 days ago			
storageclass	Update zones in 'standard' storage class	7 months ago			
zmon-agent	Update zmon-agent	10 days ago			
zmon-aws-agent	zmon-aws-agent: fix spurious describe_images calls	11 days ago			
zmon-redis	ndots for kube-system	2 months ago			
zmon-scheduler	ndots for kube-system	2 months ago			
zmon-worker	Upgrade ZMON worker	16 days ago			
deletions.yaml	Drop pod quotas in 'default' and 'kube-system'	5 days ago			

47+ cluster components



**INCIDENT**

**#1**

# INCIDENT #1: INGRESS ERRORS

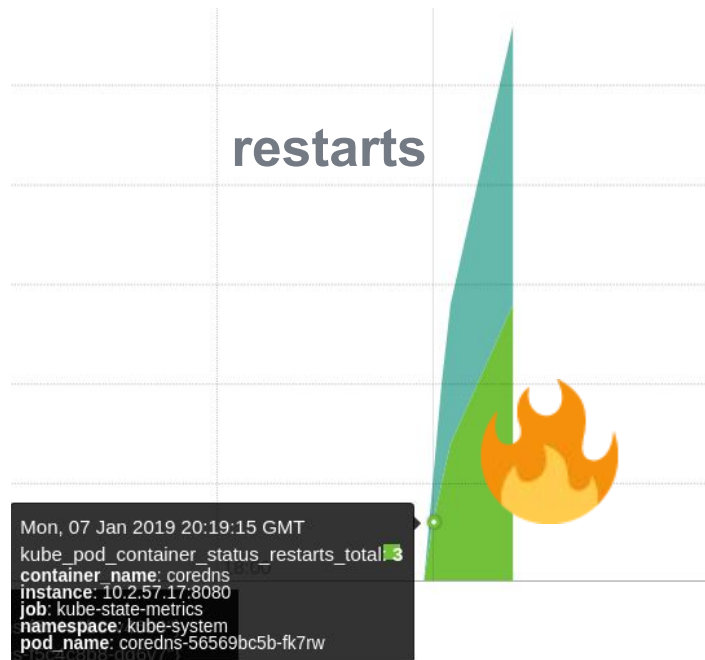


# INCIDENT #1: COREDNS OOMKILL

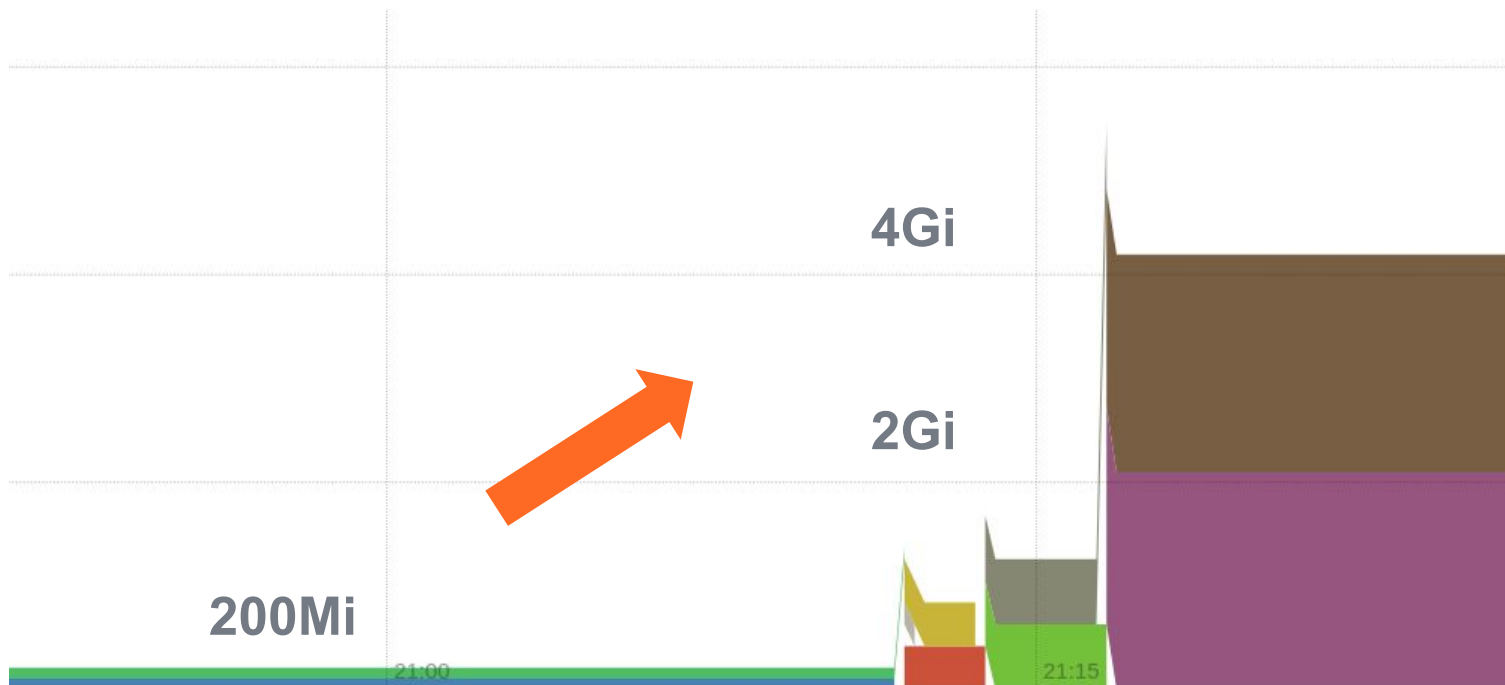
```
coredns invoked oom-killer:  
gfp_mask=0x14000c0(GFP_KERNEL),  
nodemask=(null), order=0, oom_score_adj=994
```

```
Memory cgroup out of memory: Kill process 6428  
(coredns) score 2050 or sacrifice child
```

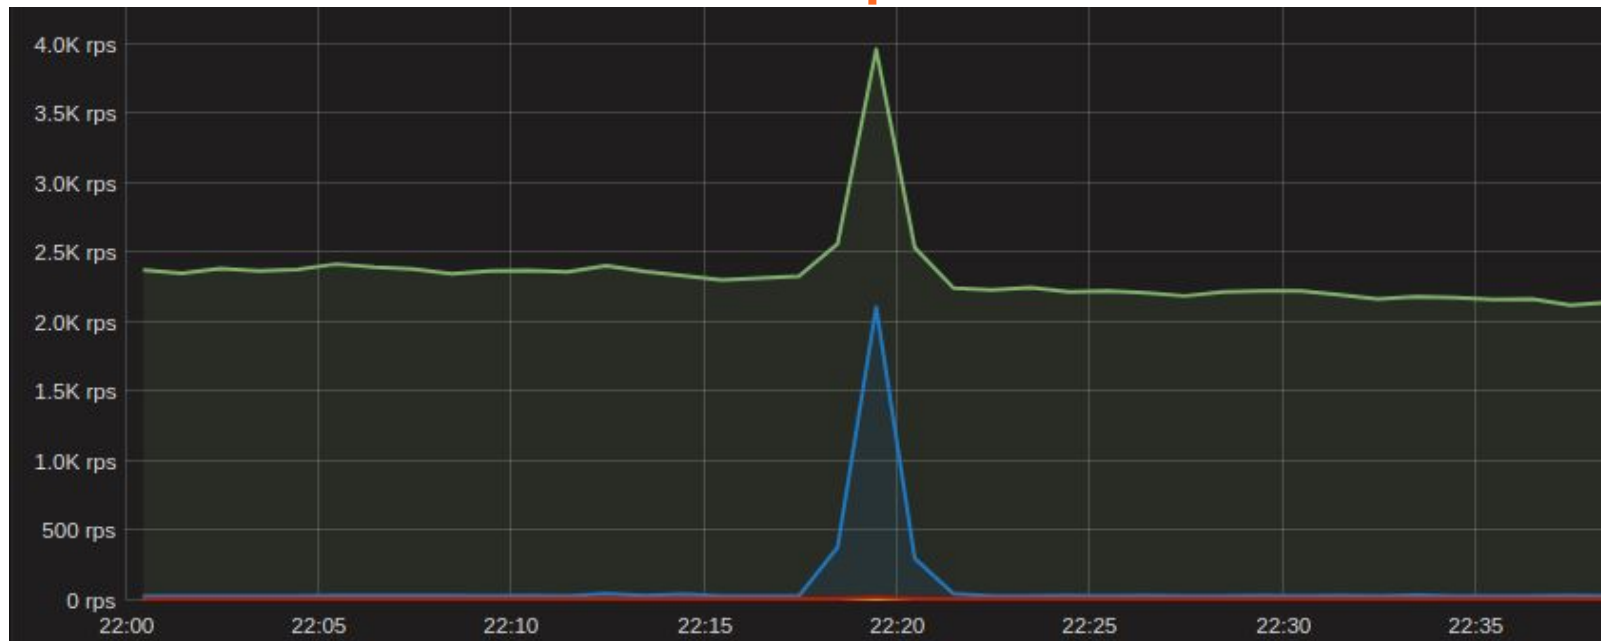
```
oom_reaper: reaped process 6428 (coredns),  
now anon-rss:0kB, file-rss:0kB, shmem-rss:0kB
```



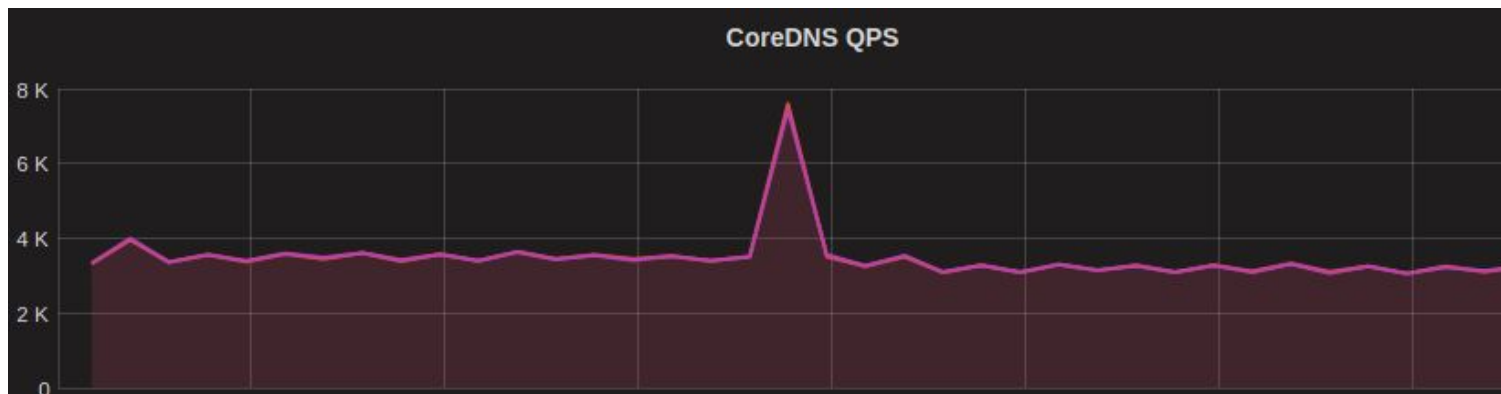
# STOP THE BLEEDING: INCREASE MEMORY LIMIT



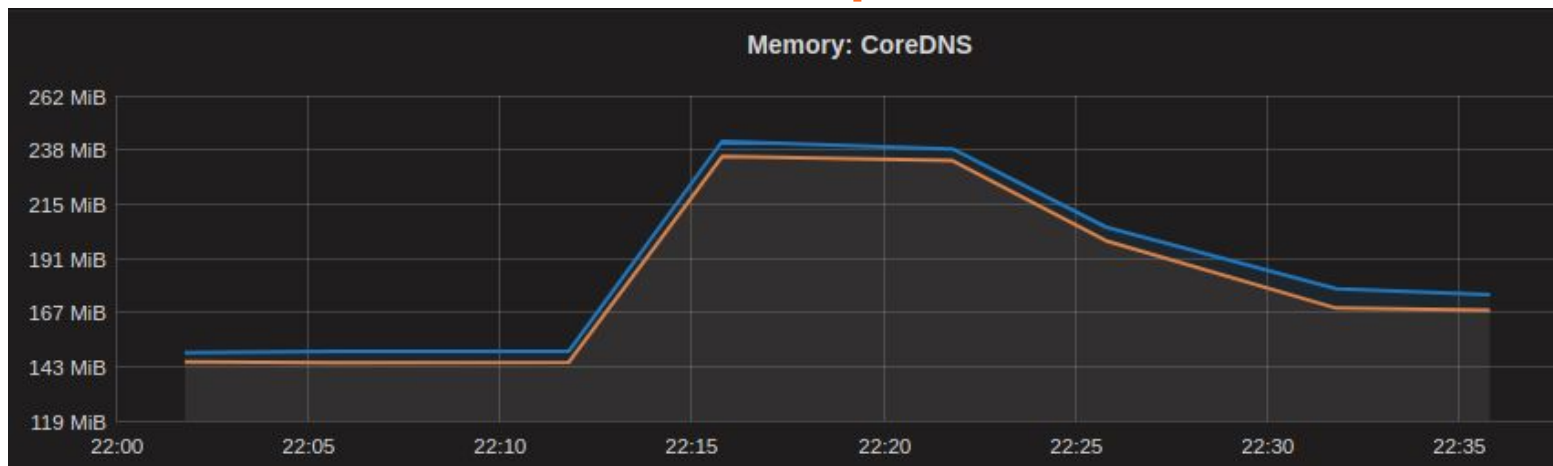
# SPIKE IN HTTP REQUESTS



# SPIKE IN DNS QUERIES



# INCREASE IN MEMORY USAGE



# INCIDENT #1: CONTRIBUTING FACTORS

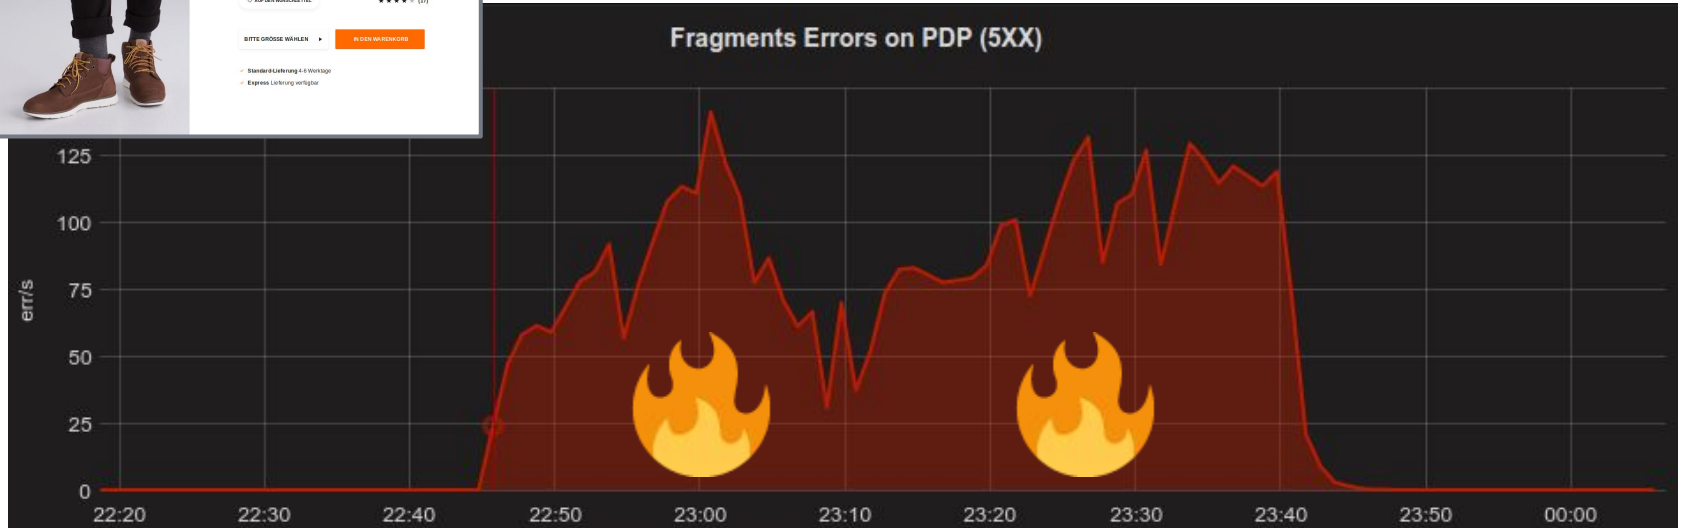
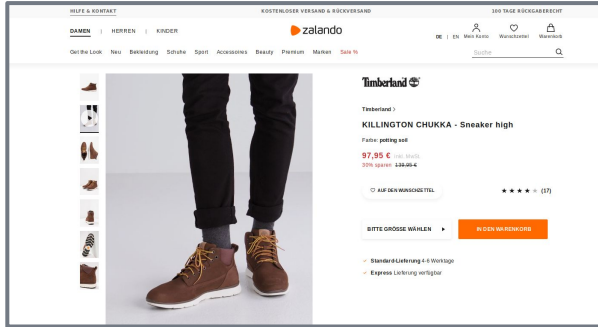


- HTTP retries
- No DNS caching
- Kubernetes ndots:5 problem
- Short maximum lifetime of HTTP connections
- Fixed memory limit for CoreDNS
- Monitoring affected by DNS outage

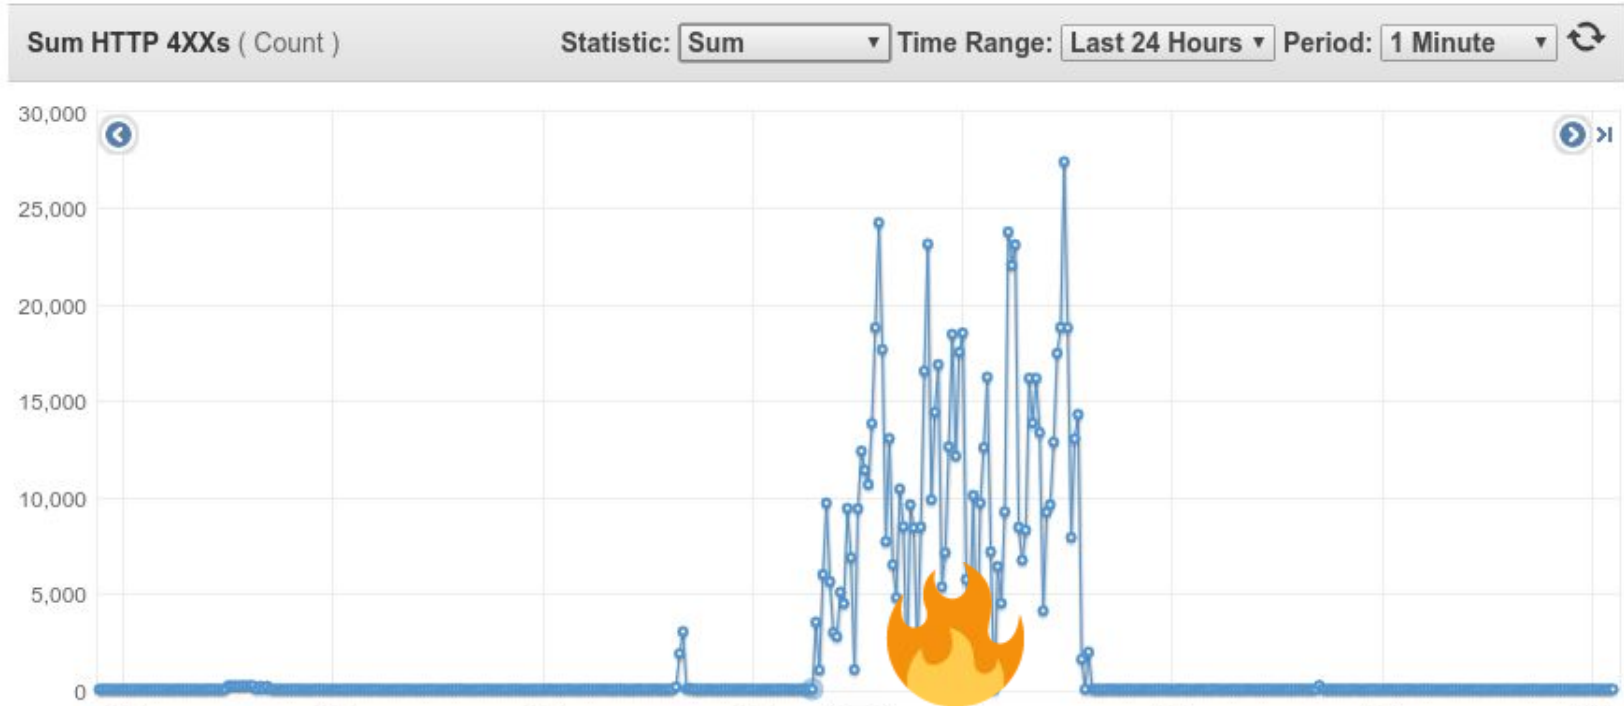
**INCIDENT**

**#2**

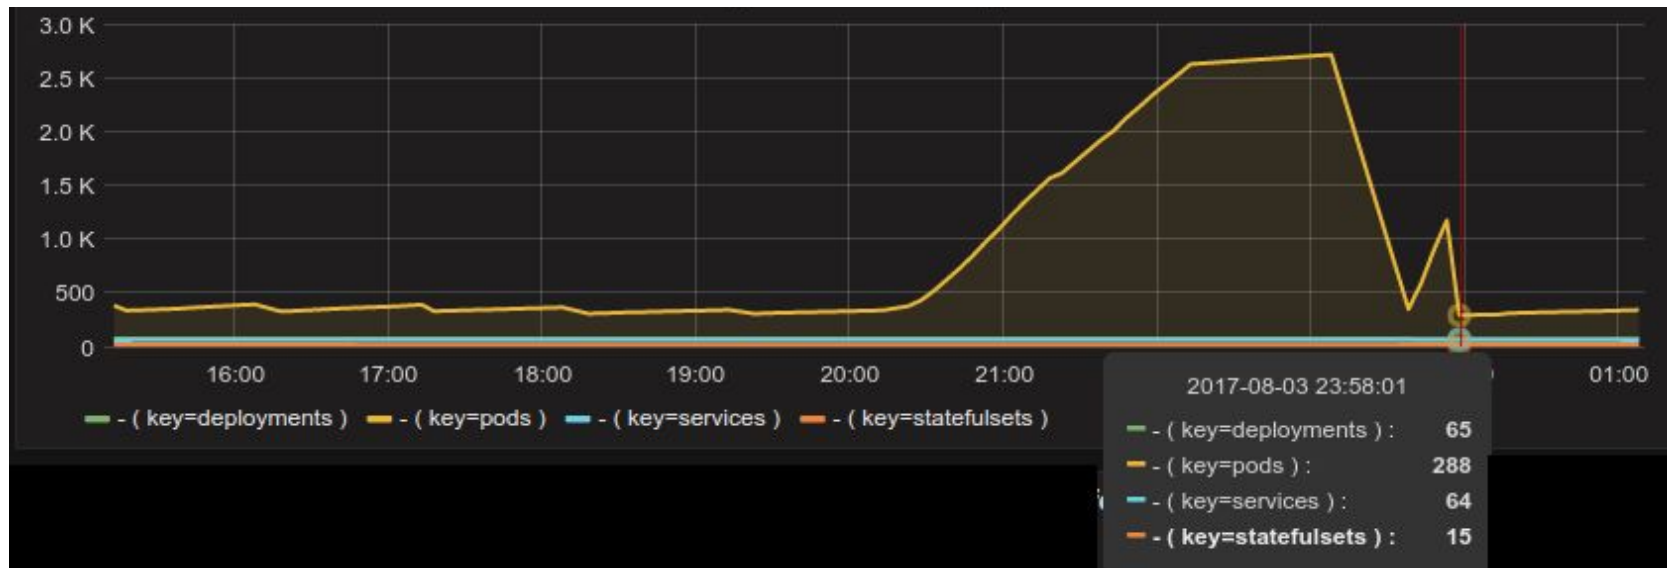
# INCIDENT #2: CUSTOMER IMPACT



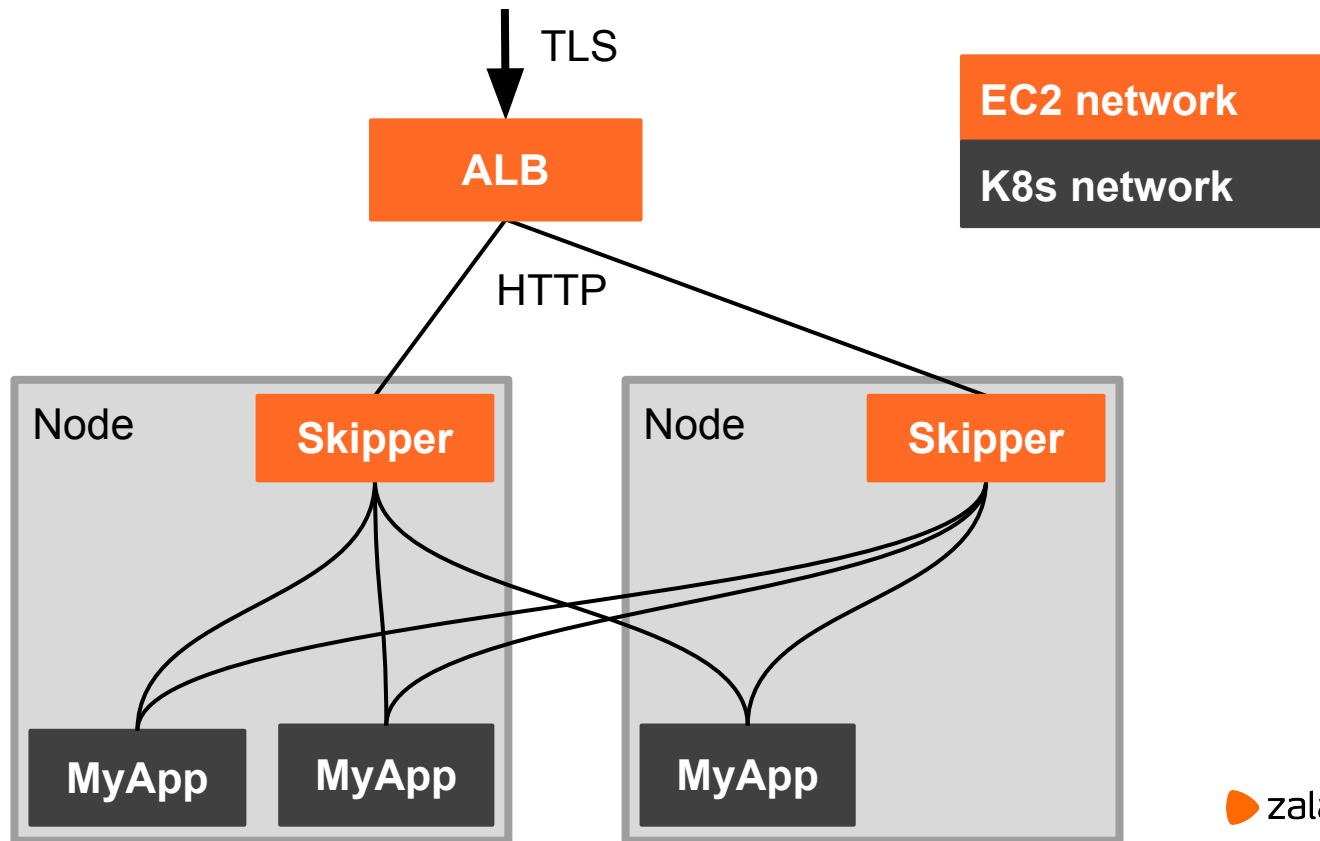
# INCIDENT #2: IAM RETURNING 404



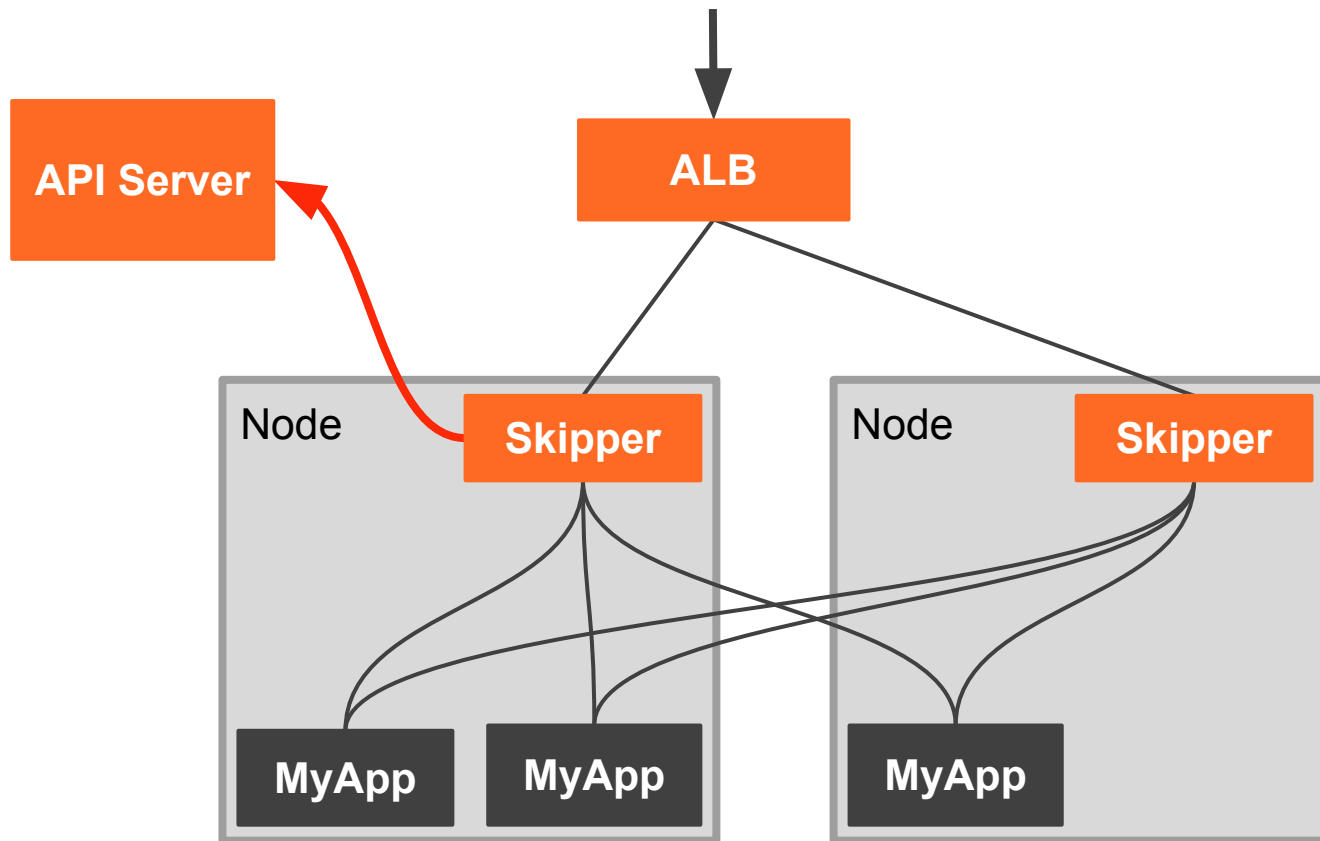
## INCIDENT #2: NUMBER OF PODS



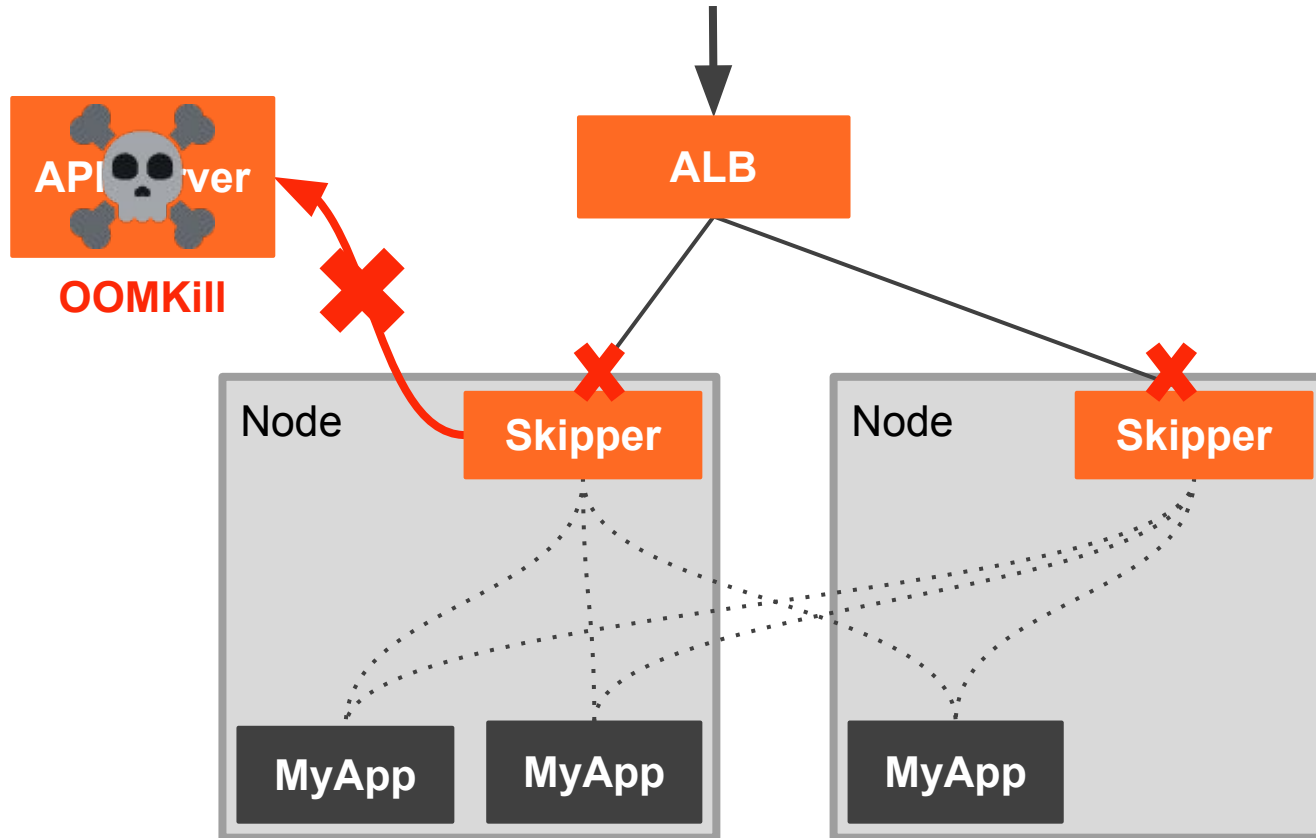
# LIFE OF A REQUEST (INGRESS)



# ROUTES FROM API SERVER



# API SERVER DOWN



## INCIDENT #2: INNOCENT MANIFEST

```
apiVersion: batch/v2alpha1
kind: CronJob
metadata:
  name: "foobar"
spec:
  schedule: "*/15 9-19 * * Mon-Fri"
  jobTemplate:
    spec:
      template:
        spec:
          restartPolicy: Never
          concurrencyPolicy: Forbid
          successfulJobsHistoryLimit: 1
          failedJobsHistoryLimit: 1
          containers:
            ...
```

## INCIDENT #2: FIXED CRON JOB

```
apiVersion: batch/v2alpha1
kind: CronJob
metadata:
  name: "foobar"
spec:
  schedule: "7 8-18 * * Mon-Fri"
  concurrencyPolicy: Forbid
  successfulJobsHistoryLimit: 1
  failedJobsHistoryLimit: 1
  jobTemplate:
    spec:
      activeDeadlineSeconds: 120
      template:
        spec:
          restartPolicy: Never
          containers:
```

## INCIDENT #2: LESSONS LEARNED

- Fix Ingress to stay “healthy” during API server problems
- Fix Ingress to retain last known set of routes
- Use quota for number of pods



```
apiVersion: v1
kind: ResourceQuota
metadata:
  name: compute-resources
spec:
  hard:
    pods: "1500"
```

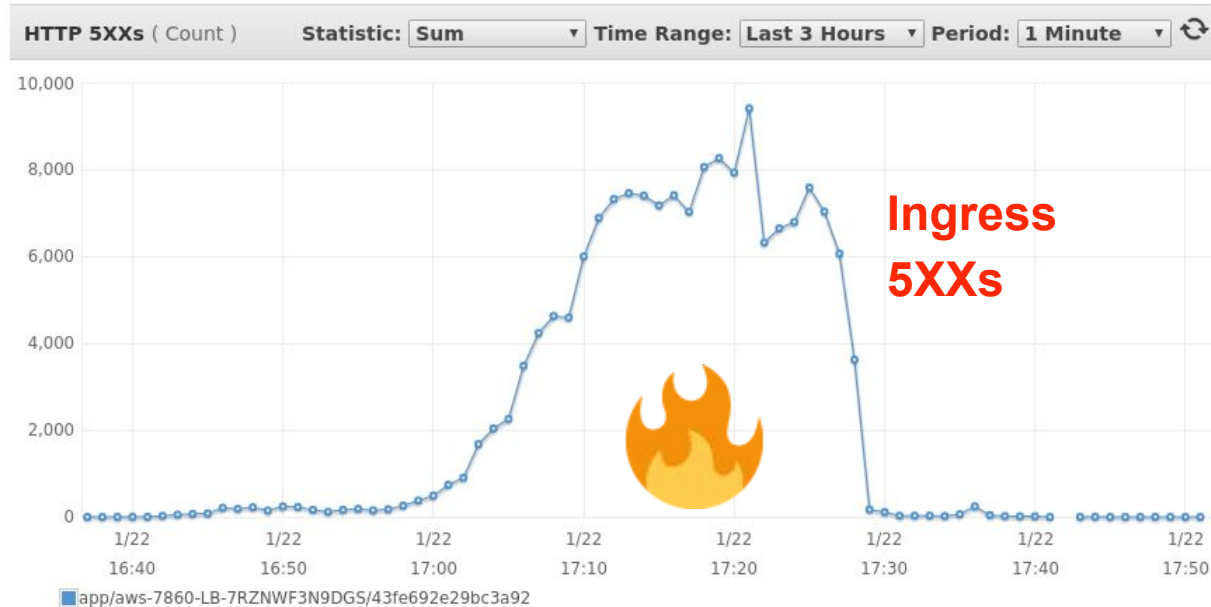
NOTE: we dropped quotas recently  
[github.com/zalando-incubator/kubernetes-on-aws/pull/2059](https://github.com/zalando-incubator/kubernetes-on-aws/pull/2059)

**INCIDENT**

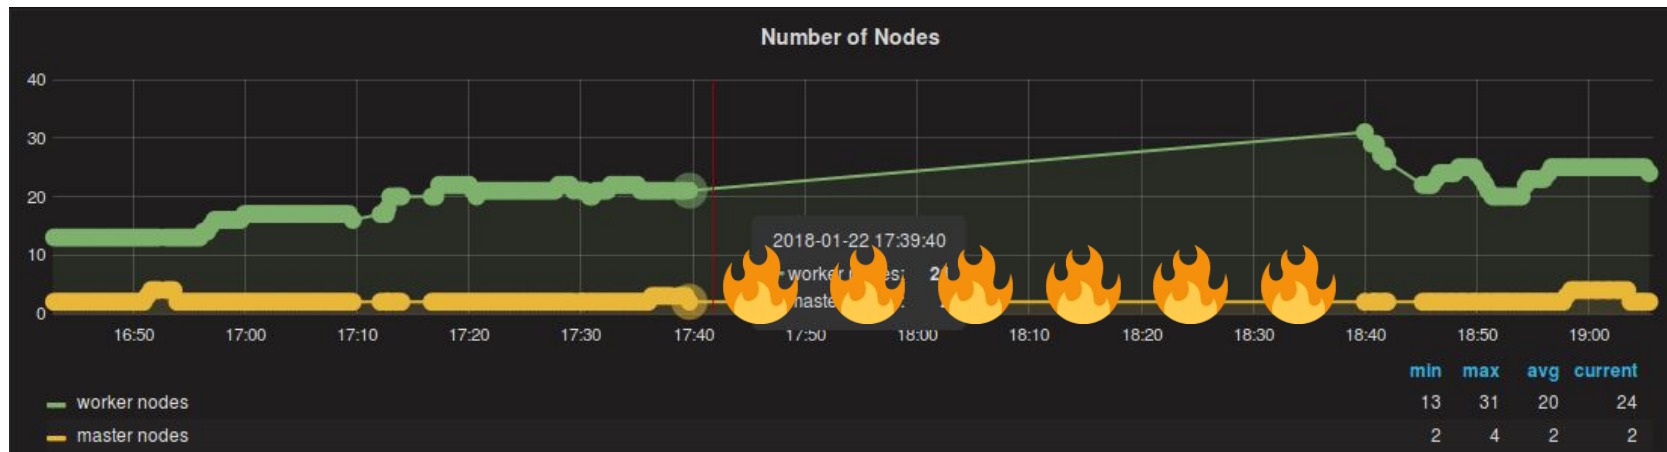
**#3**

# INCIDENT #3: IMPACT

## CloudWatch Monitoring Details



# INCIDENT #3: CLUSTER DOWN?



# INCIDENT #3: THE TRIGGER


zalando-incubator / **kubernetes-on-aws** Unwatch 25 Unstar 147 Fork 32

[Code](#) [Issues 10](#) [Pull requests 10](#) [Insights](#) [Settings](#)

## dev -> alpha #810 Edit

**Merged** aermakov-zalando merged 19 commits into `alpha` from `dev-to-alpha` 29 days ago

[Conversation 6](#) [Commits 19](#) [Files changed 12](#) +289 -59







**aermakov-zalando** commented 29 days ago Member

Highlights:

- k8s as flannel backend
- Disable image-policy simulation mode
- Fix source IP filtering in skipper
- HA kube-dns

**mikkeloscar** and others added some commits on Nov 18, 2016

-   Use k8s as flannel backend Verified 1e444fa
-   Move flannel to manifests folder Verified 4f35dde

**Reviewers** linkl

**Assignees** No one—assign yourself

**Labels** ready-to-test

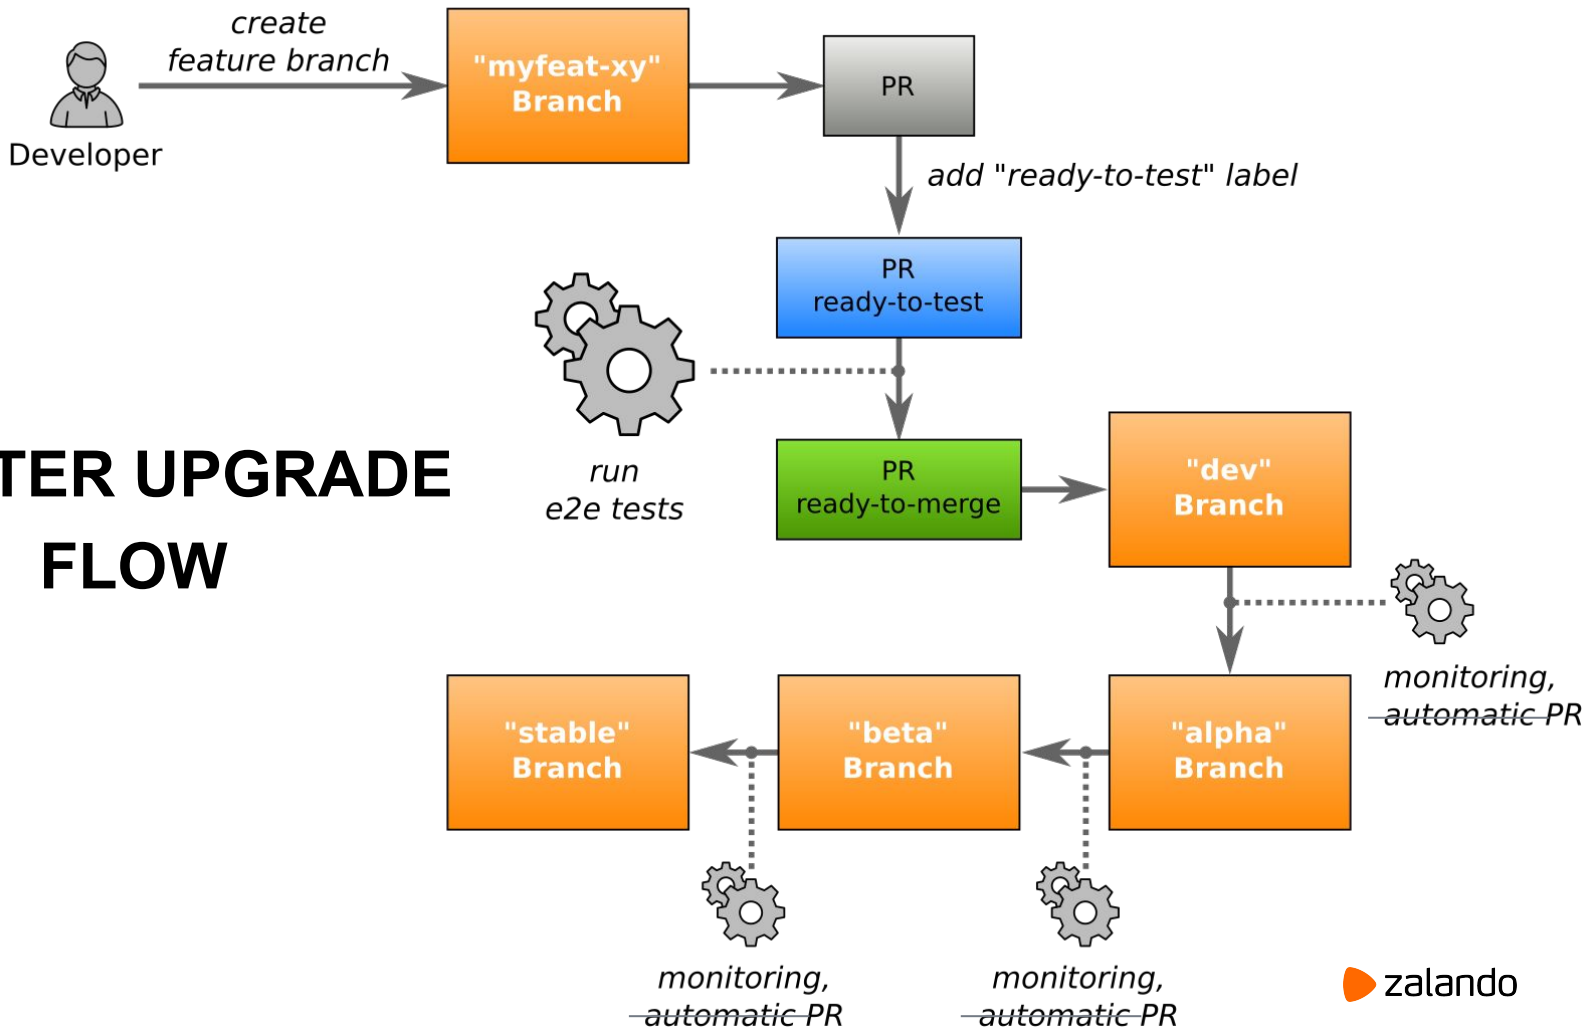
**Projects** None yet

What We Believe

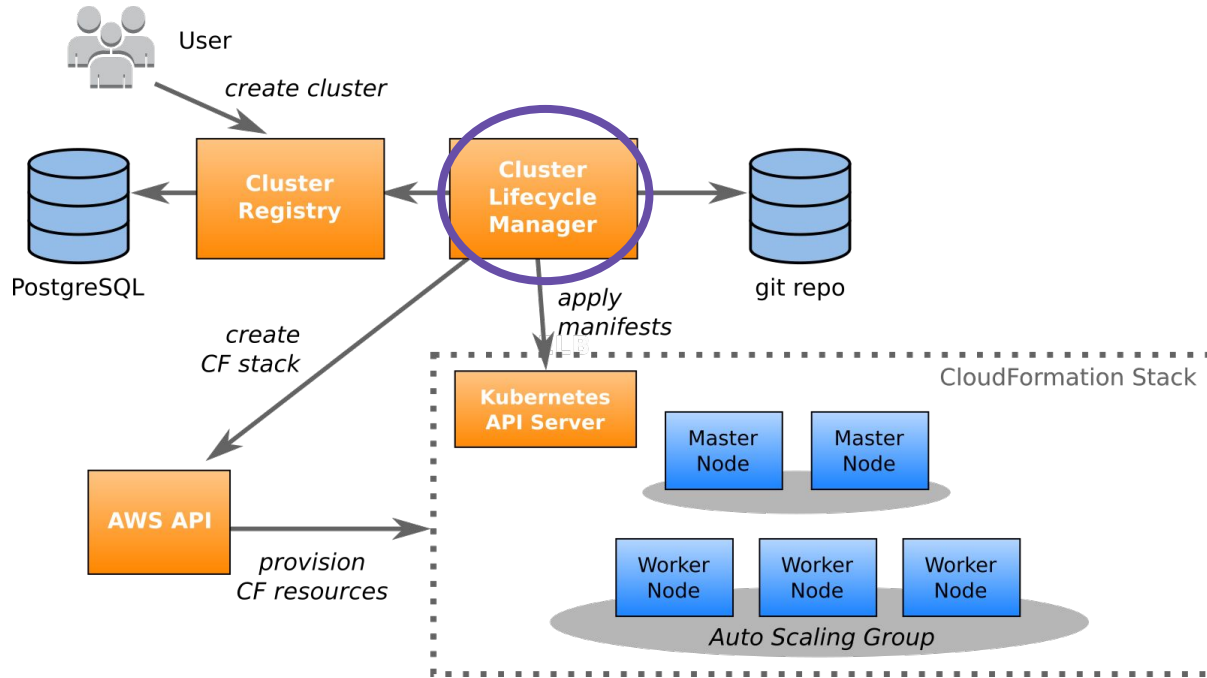
**VOL. 1 ISSUE 6**

# Human Error is NEVER the Root Cause

# CLUSTER UPGRADE FLOW



# CLUSTER LIFECYCLE MANAGER (CLM)





[github.com/zalando-incubator/cluster-lifecycle-manager](https://github.com/zalando-incubator/cluster-lifecycle-manager)

# CLUSTER CHANNELS


Channel	Description	Clusters
dev	Development and playground clusters.	3
alpha	Main infrastructure clusters ( <b>important to us</b> ).	2
beta	Product clusters for the rest of the organization (non-prod).	60+
stable	Product clusters for the rest of the organization (prod).	60+


# E2E TESTS ON EVERY PR

🔑  Update pdb-controller Verified ● c37c725

🏷️  mikkeloscar added the **ready-to-test** label an hour ago






Add more commits by pushing to the **pdb-controller-update** branch on **zalando-incubator/kubernetes-on-aws**.






**Some checks haven't completed yet**  
2 pending and 5 successful checks

[Hide all checks](#)

●	 <b>ci/jenkins/statefulsets-tests</b>	Pending — Running e2e tests.	<a href="#">Details</a>
●	 <b>zappr</b>	Pending — This PR needs 2 more approvals (0/2 given).	<b>Required</b>
✓	 <b>ci/jenkins/conformance-tests</b>	— Build finished.	<a href="#">Details</a>
✓	 <b>ci/jenkins/zalando-tests</b>	— Build finished.	<a href="#">Details</a>
✓	 <b>continuous-integration/travis-ci/pr</b>	— The Travis CI build passed	<a href="#">Details</a>

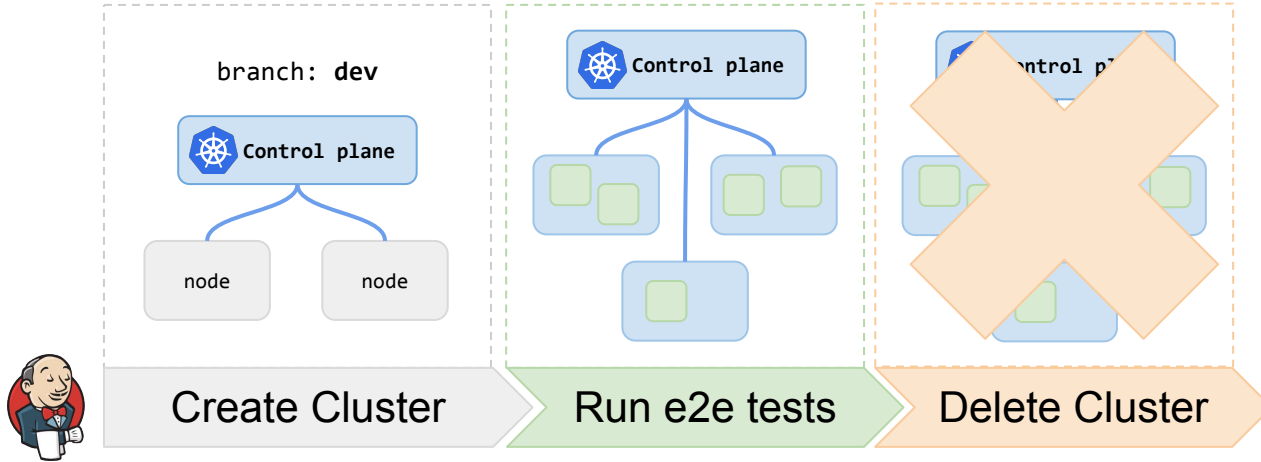


**Required statuses must pass before merging**  
All required [status checks](#) on this pull request must run successfully to enable automatic merging.

[github.com/zalando-incubator/kubernetes-on-aws](https://github.com/zalando-incubator/kubernetes-on-aws)

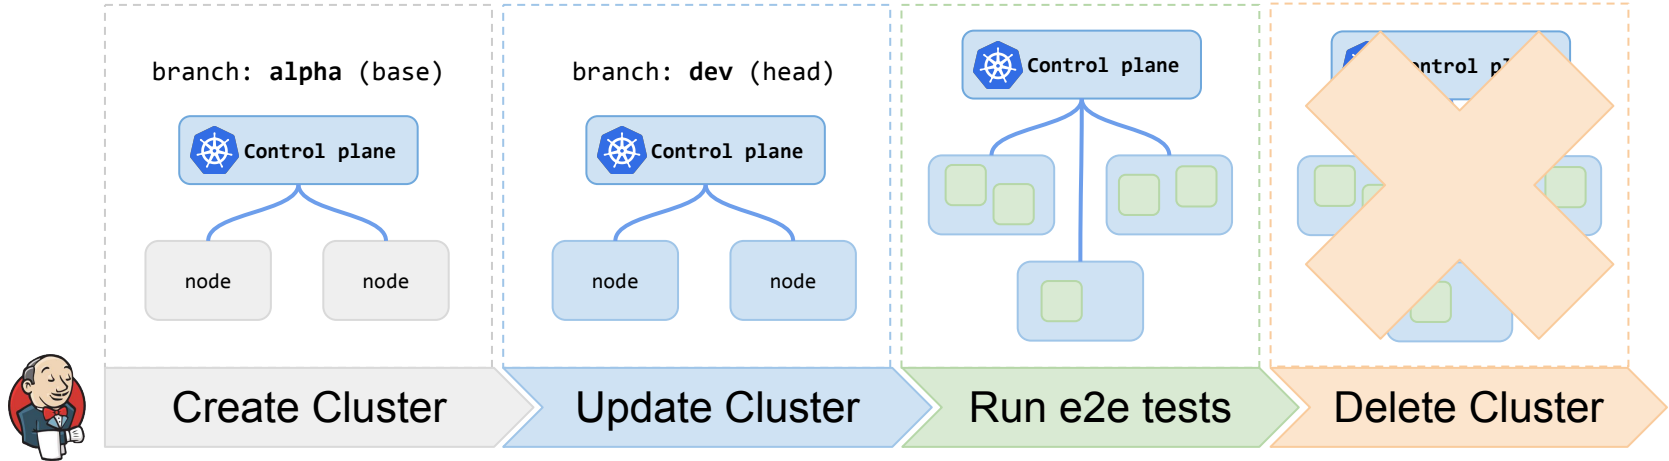
# RUNNING E2E TESTS (BEFORE)

Testing **dev** to **alpha** upgrade



# RUNNING E2E TESTS (NOW)

Testing **dev** to **alpha** upgrade



## INCIDENT #3: LESSONS LEARNED

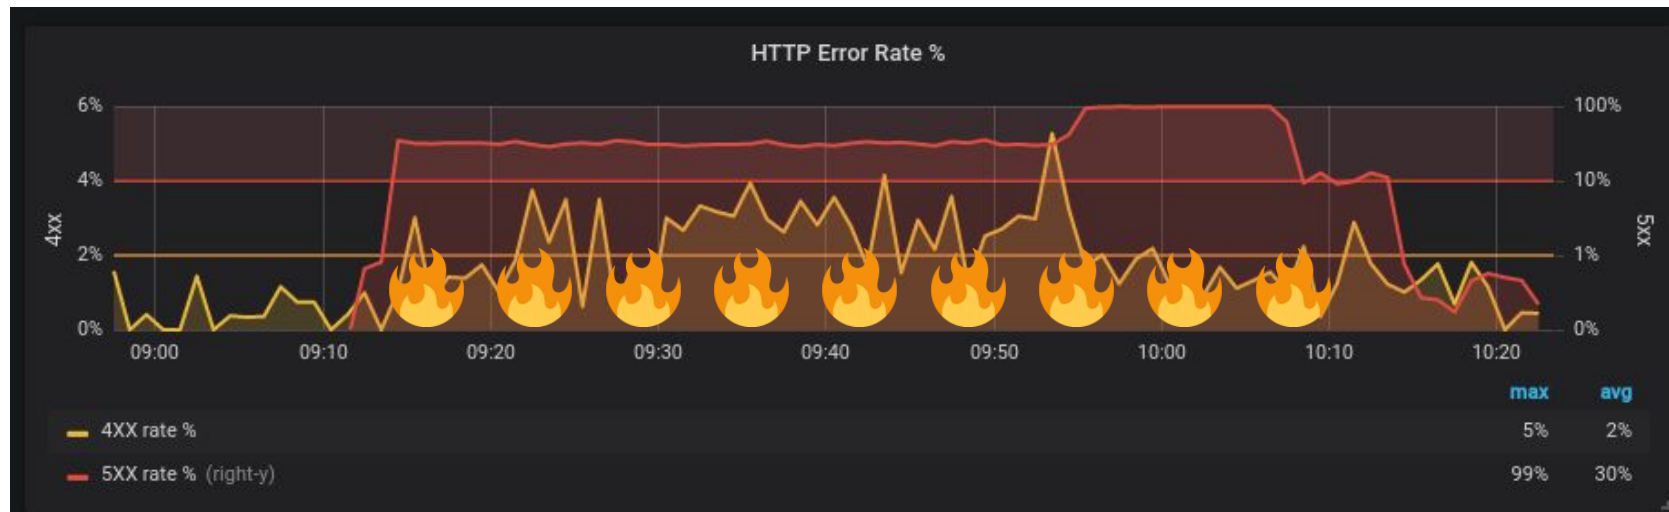
- Automated e2e tests are pretty good, but not enough
- Test the diff/migration automatically
  - Bootstrap new cluster with previous configuration
  - Apply new configuration
  - Run end-to-end & conformance tests



**INCIDENT**

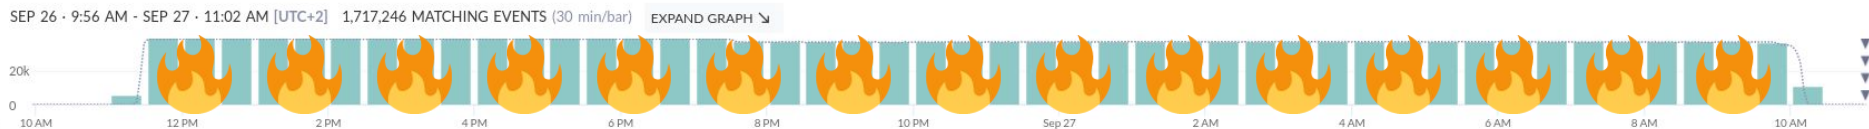
**#4**

# INCIDENT #4: IMPACT



# INCIDENT #4: FLANNEL ERRORS

Failed to list \*v1.Node: Unauthorized



# INCIDENT #4: RBAC CHANGES

17 cluster/manifests/deletions.yaml

Viewed ...

@@ -44,3 +44,20 @@ post\_apply:

```
44 - name: zmon-scheduler
45   kind: VerticalPodAutoscaler
46   namespace: visibility
```

```
44 - name: zmon-scheduler
45   kind: VerticalPodAutoscaler
46   namespace: visibility
```

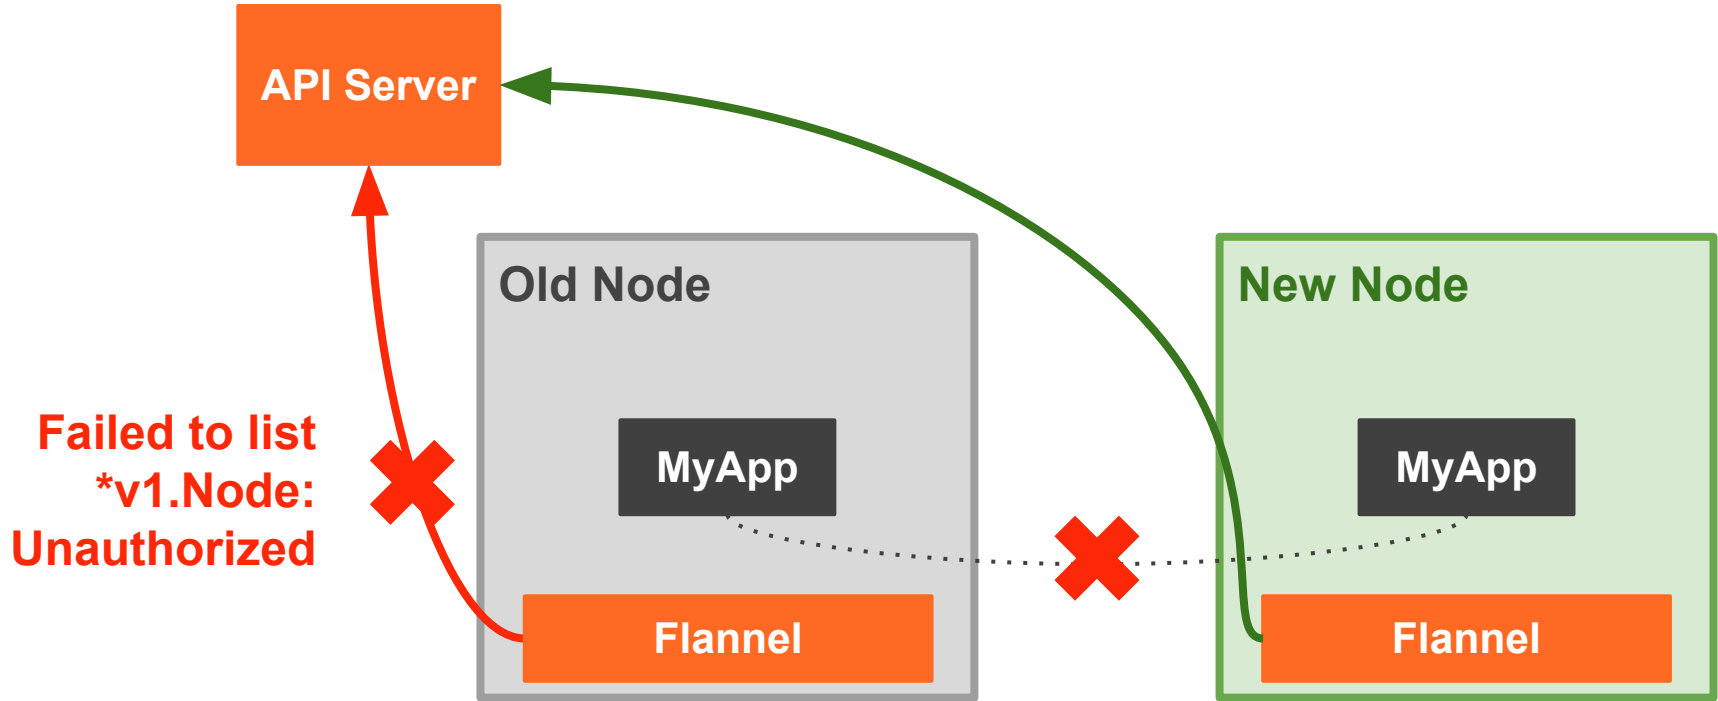
```
47 + - name: kubernetes-dashboard
48 +   kind: RoleBinding
49 +   namespace: kube-system
```

```
50 + - name: system
51 +   namespace: kube-system
52 +   kind: ServiceAccount
```

```
53 + - name: privileged-psp
54 +   namespace: kube-system
55 +   kind: RoleBinding
```

```
56 + - name: cdp-deployer
57 +   kind: ClusterRoleBinding
```

## INCIDENT #4: NETWORK SPLIT



**INCIDENT**

**#5**

## INCIDENT #5: IMPACT

Error during Pod creation:

```
MountVolume.SetUp failed for volume  
"outfit-delivery-api-credentials" :  
secrets "outfit-delivery-api-credentials" not found
```

⇒ All new Kubernetes deployments fail

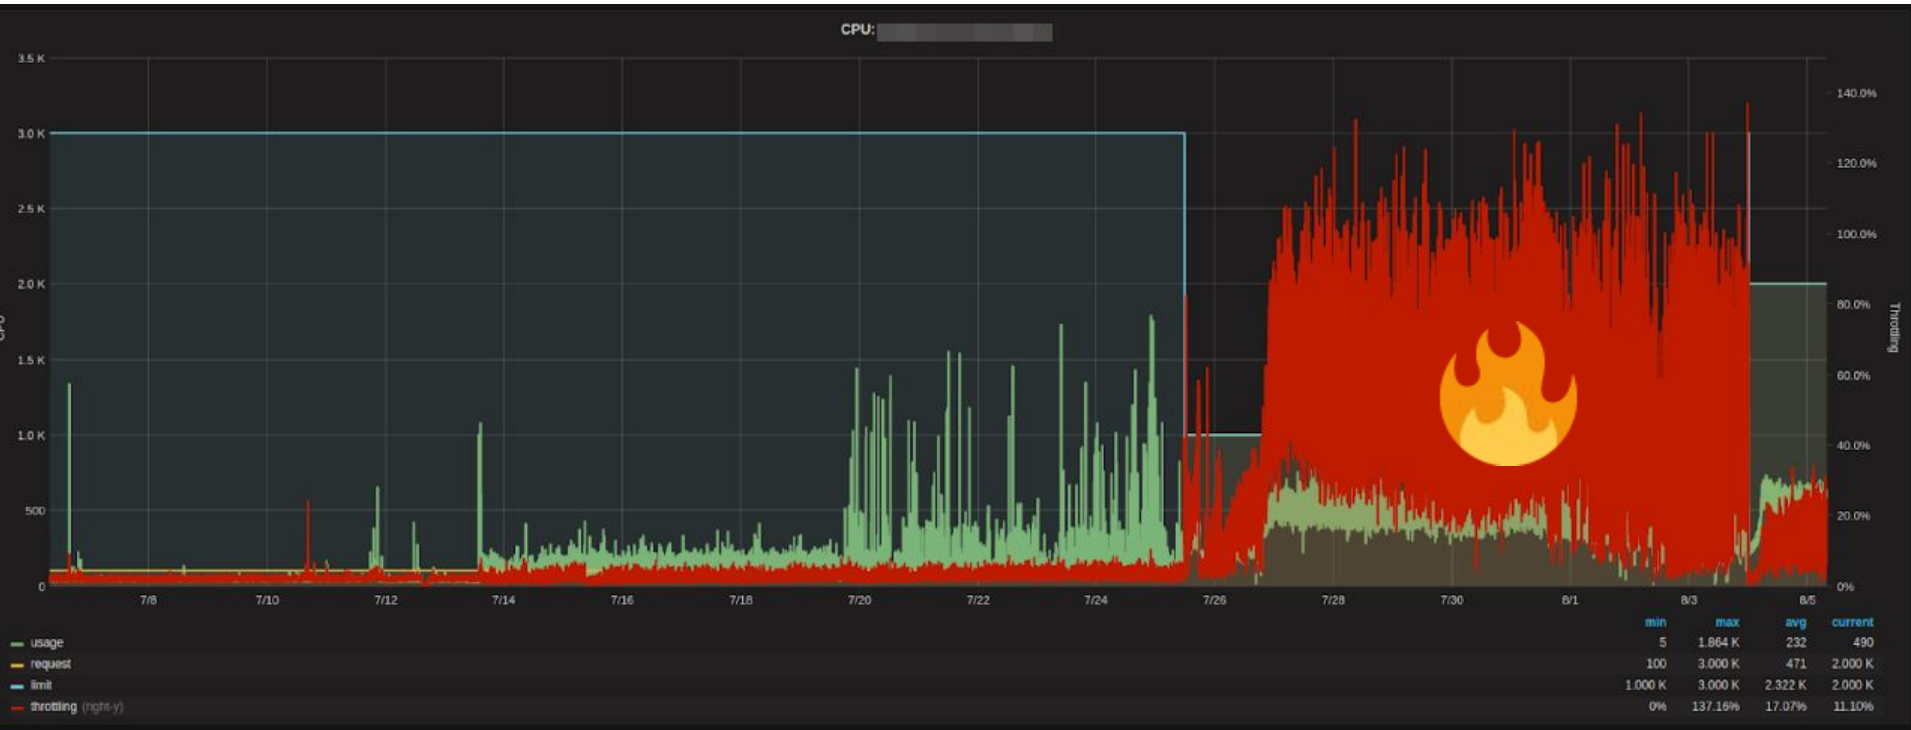


## INCIDENT #5: CREDENTIALS QUEUE

```
17:30:07 | [pool-6-thread-1 ] | Current queue size: 7115, current number of active workers: 20
17:31:07 | [pool-6-thread-1 ] | Current queue size: 7505, current number of active workers: 20
17:32:07 | [pool-6-thread-1 ] | Current queue size: 7886, current number of active workers: 20
..
17:37:07 | [pool-6-thread-1 ] | Current queue size: 9686, current number of active workers: 20
..
17:44:07 | [pool-6-thread-1 ] | Current queue size: 11976, current number of active workers: 20
..
19:16:07 | [pool-6-thread-1 ] | Current queue size: 58381, current number of active workers: 20
```



# INCIDENT #5: CPU THROTTLING



## INCIDENT #5: WHAT HAPPENED

Scaled down IAM provider  
to reduce **Slack**

+ Number of deployments increased

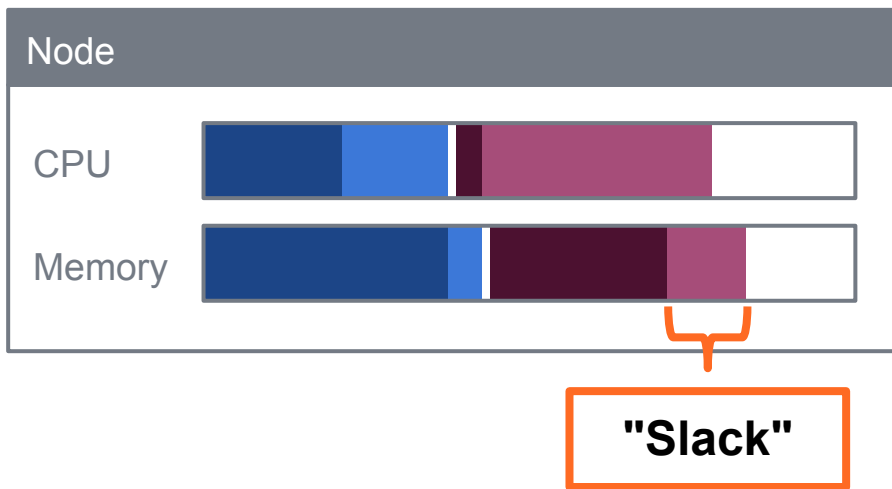


⇒ Process could not process credentials fast enough

# SLACK

CPU/memory requests "block" resources on nodes.

Difference between actual usage and requests → **Slack**



# DISABLING CPU THROTTLING

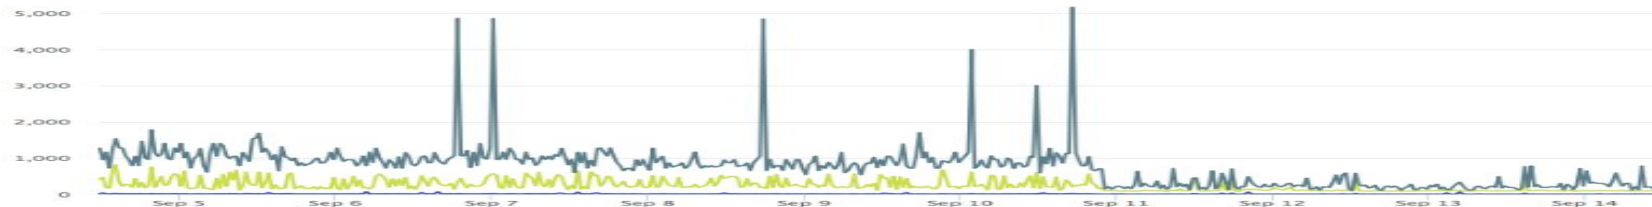
```
kubelet ... --cpu-cfs-quota=false
```

## [Announcement] CPU limits will be disabled

TLDR: to **improve performance** and efficiency we will disable CPU limits in Kubernetes clusters. Please **revise your resource requests** if necessary.

We're going to disable CPU limits in the Kubernetes clusters. According to our experiments, this should improve the latencies for your applications and allow us to use the nodes more efficiently. To ensure that your applications get their fair share of CPU, please update your deployments' [resource requests](#) so they match the actual usage. You can use the [Application Dashboard](#) to find out how much CPU your applications use.

## ⇒ Ingress Latency Improvements



# DISABLING CPU THROTTLING

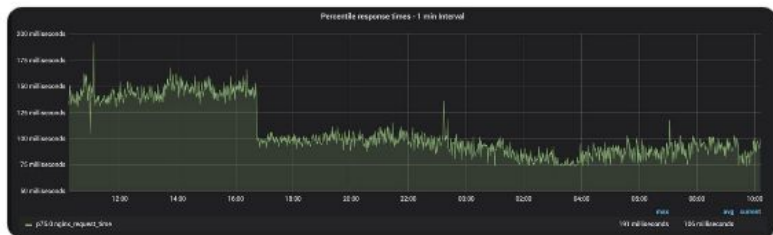


Thomas Peitz

@it\_supertramp

We have reduced 75 percentile response time over all apps from 150ms to 90ms after disabling CFS quota (CPU limits) on one of our **#kubernetes** cluster - **#KubeCon** learning by **@try\_except\_**

[Tweet übersetzen](#)



10:16 - 29. Mai 2019

46 Retweets 162 „Gefällt mir“-Angaben



8

46

162



Tim Hockin @thockin · 30. Mai

Antwort an **@it\_supertramp** **@try\_except\_**

This is why I always advise:

- 1) Always set memory limit == request
- 2) Never set CPU limit

(for locally adjusted values of "always" and "never")

[Tweet übersetzen](#)

5

15

55



# MORE TOPICS

- 🔥 **Graceful Pod shutdown** and race conditions (endpoints, Ingress)
- 🔥 **Incompatible Kubernetes changes**
- 🔥 CoreOS **ContainerLinux** "stable" won't boot
- 🔥 Kubernetes **EBS volume handling**
- 🔥 **Docker**



# RACE CONDITIONS..

```
21 priorityClassName: system-node-critical
22 serviceAccountName: system
23 containers:
24   - name: delayed-install-cni
25     image: registry.opensource.zalan.do/teapot/flannel:v0.10.0-8
26     command:
27     - /bin/sh
28     args:
29     - c
30     - "sleep 120 && cp -f /etc/kube-flannel/cni-conf.json /etc/cni/net.d/10-flannel.conf && cat"
31     stdin: true
32     volumeMounts:
33     - name: cni
34       mountPath: /etc/cni/net.d
35     - name: flannel-cfg
36       mountPath: /etc/kube-flannel/
```

# DOCKER.. (ON GKE)

```
25 # We simply kill the process when there is a failure. Another systemd service will
26 # automatically restart the process.
27 function docker_monitoring {
28     while [ 1 ]; do
29         if ! timeout 10 docker ps > /dev/null; then
30             echo "Docker daemon failed!"
31             pkill docker
32             # Wait for a while, as we don't want to kill it again before it is really up.
33             sleep 30
34         else
35             sleep "${SLEEP_SECONDS}"
36         fi
37     done
38 }
```





## COMMON PITFALLS

# COMMON PITFALLS

- Insufficient e2e tests
- Readiness & Liveness Probes
- Resource Requests & Limits
- DNS



# READINESS & LIVENESS PROBES



Sandor Szücs

@sszuecs



Most people that are new to [#kubernetes](#) do the same mistakes:

- no readinessprobe
- wrong readinessprobe
- livenessprobe = readinessprobe
- non graceful shutdown
- graceful shutdown which is not graceful enough, best use lifecycle hook [opensource.zalando.com/skipper/kubern...](#)
- pre fork mode



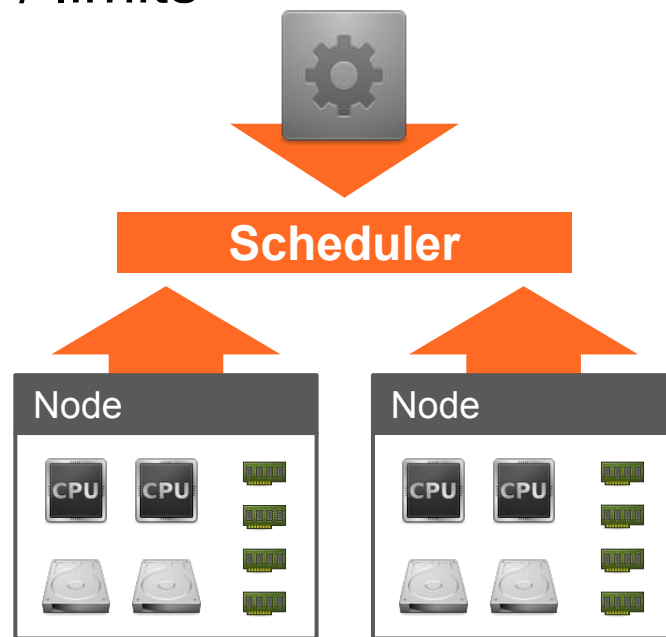
161 1:52 PM - Sep 21, 2019



[srcco.de/posts/kubernetes-liveness-probes-are-dangerous.html](https://srcco.de/posts/kubernetes-liveness-probes-are-dangerous.html)

# RESOURCE REQUESTS & LIMITS

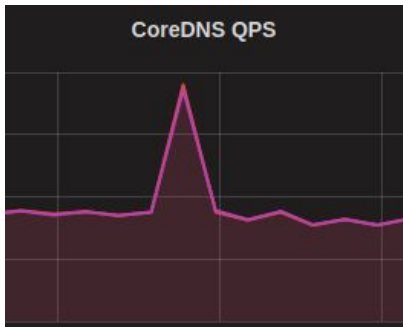
- No resource requests / limits
- QoS
- OOM
- Overcommit
- CPU throttling



# DNS

- ndots: 5
- musl, conntrack, UDP
- overload

*It's ~~never~~ always DNS.*



## resolv.conf

```
search    <namespace>.svc.cluster.local
          svc.cluster.local
          cluster.local
          ec2.internal
```

3+ search domains  
ndots : 5

```
options   ndots:5
```

www.google.com?

```
1: www.google.com.<namespace>.svc.cluster.local
2: www.google.com.svc.cluster.local
3: www.google.com.cluster.local
4: www.google.com.google.internal
5: www.google.com
```

A? / AAAA?	NXDOMAIN
A? / AAAA?	NXDOMAIN
A? / AAAA?	NXDOMAIN
A? / AAAA?	NXDOMAIN
A? / AAAA?	NOERROR

[youtube.com/watch?v=QKI-JRs2RIE](https://youtube.com/watch?v=QKI-JRs2RIE)

# AWS EKS IN PRODUCTION

## DNS lookup scaling

Out of the box, AWS provides a `kube-dns` deployment containing a single pod of scale `1`. After a week or so in production, I was skimming our logs and came across this beauty. This reinforced something I had seen in our exception handling system.

```
dnsmasq[14]: Maximum number of concurrent DNS queries reached (max: 150)
```

[kubedex.com/90-days-of-aws-eks-in-production/](https://kubedex.com/90-days-of-aws-eks-in-production/)



## TOOLS & PRACTICES TO HELP YOU


# HELPFUL


- Automated e2e tests
- Monitoring
- OpenTracing
- Kubernetes Web View
- Emergency access
- Kubernetes Failure Stories



# AUTOMATED E2E TESTS

KUBERNETES-ON-AWS


 BUILD

 E2E-TESTS


GENERAL


DEPLOYMENT STATUS

LOGS

 Ran for 37 m, 46 s

2:44:44 PM time="2019-10-21T12:44:44Z" level=debug msg="Stack 'nodepool-default-worker-splitaz-aws -

 README.md



## E2E tests for Kubernetes on AWS

This directory contains e2e tests for the kubernetes-on-aws configuration.

It is based on the test framework used by Kubernetes itself. See [Kubernetes e2e tests](#) for examples of how to write the tests or checkout the files already defined e.g. `external_dns.go`.

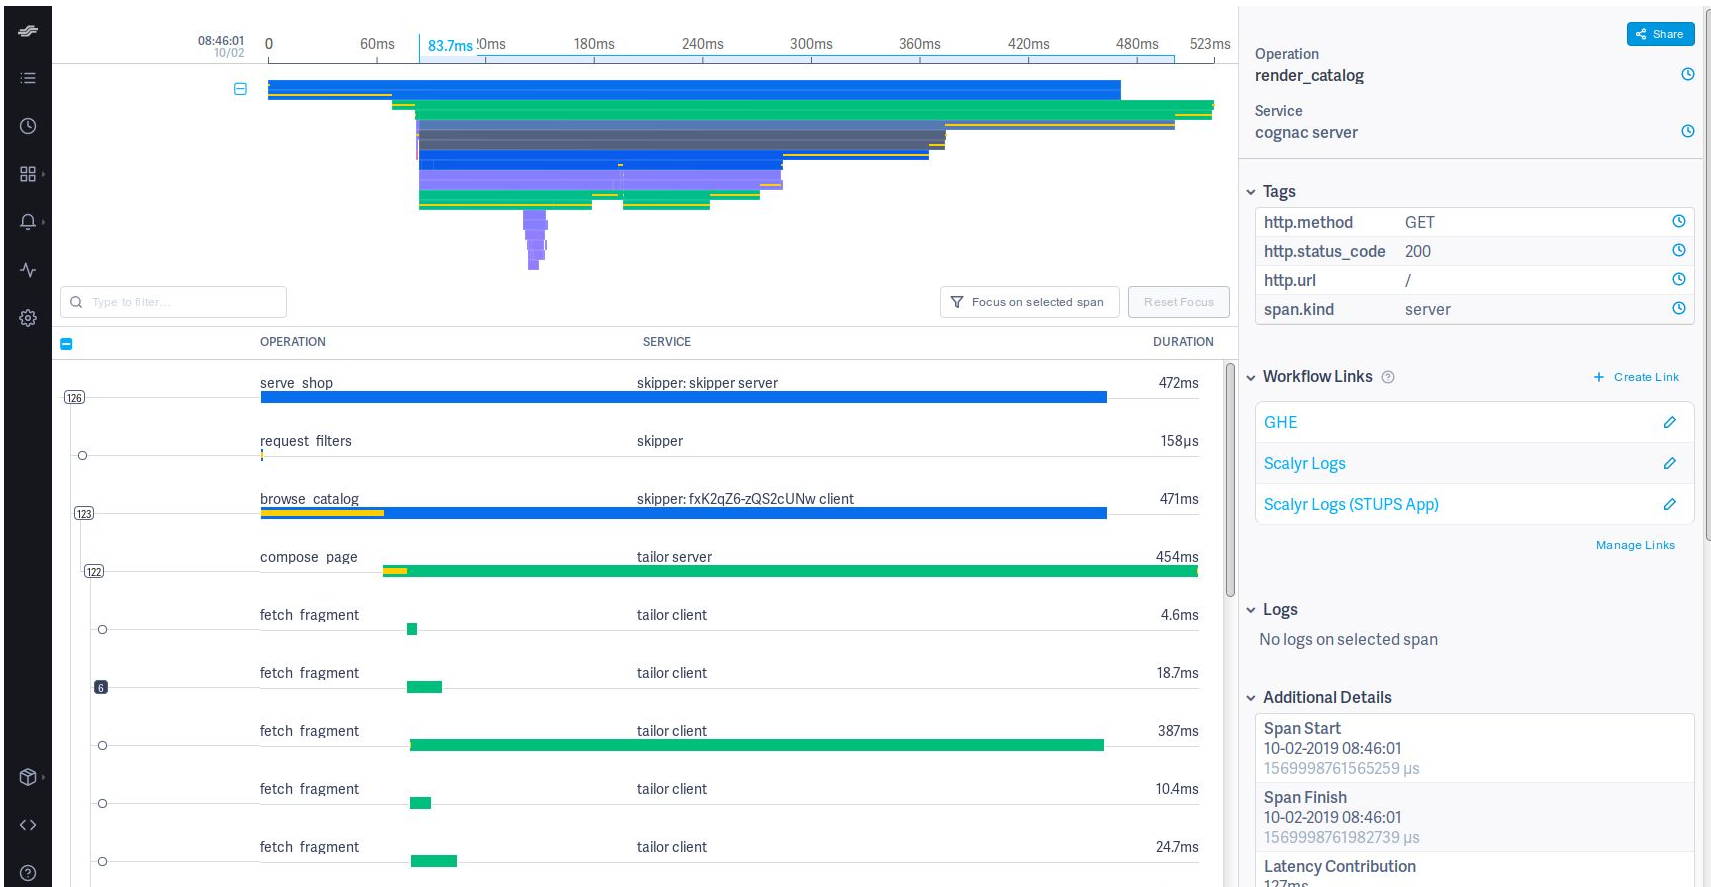
[github.com/zalando-incubator/kubernetes-on-aws/tree/dev/test/e2e](https://github.com/zalando-incubator/kubernetes-on-aws/tree/dev/test/e2e)

# MONITORING

- Alert on symptoms, not potential causes
- What do you need to monitor to ensure cluster availability?
- SLOs

Ready	Status	Restarts	Age
0/1	CrashLoopBackOff	1404	4d23h
0/1	CrashLoopBackOff	1352	4d18h
0/1	CrashLoopBackOff	1346	4d19h

# OPENTRACING



# KUBERNETES WEB VIEW

Clusters

default

Search Kubernetes objects...

CLUSTER RESOURCES

Namespaces

Nodes

PersistentVolumes

CONTROLLERS

StackSets

Stacks

Deployments

CronJobs

Jobs

StatefulSets

POD MANAGEMENT

Ingresses

Services

Pods

ConfigMaps

CRDS

PlatformCredentialsSets

postgresqls

META

Resource Types

Events

/ default / pods,stacks,deployments,services

Pods

Name	Application	Component	Ready	Status	Restarts	Age	IP	Node	Nominated Node	Readiness Gates	CPU Usage	Memory Usage	Created
<a href="#">even-master-33-5db9d68c8d-5srrh</a>	even		1/1	Running	0	5d18h			<none>	<none>	3m	313 MiB	2019-08-29 16:35:03
<a href="#">even-master-33-5db9d68c8d-72px8</a>	even		1/1	Running	0	5d18h			<none>	<none>	3m	297 MiB	2019-08-29 17:05:16
<a href="#">even-master-33-5db9d68c8d-rrh9m</a>	even		1/1	Running	0	5d18h			<none>	<none>	2m	312 MiB	2019-08-29 17:20:13

Stacks

Name	Desired	Current	Up-to-date	Available	Traffic	No-Traffic-Since	Age	Created
<a href="#">even-master-27</a>	3	0	0	0	0	131d	210d	2019-02-05 20:11:56
<a href="#">even-master-29</a>	3	0	0	0	0	131d	140d	2019-04-17 08:30:22
<a href="#">even-master-30</a>	3	0	0	0	0	131d	131d	2019-04-25 12:19:11
<a href="#">even-master-31</a>	3	0	0	0	0	50d	131d	2019-04-25 12:30:11
<a href="#">even-master-32</a>	3	0	0	0	0	64d	64d	2019-07-01 15:19:45
<a href="#">even-master-33</a>	3	3	3	3	100		50d	2019-07-15 12:18:58

Deployments

Name	Ready	Up-to-date	Available	Age	Containers	Images	Selector	Created
------	-------	------------	-----------	-----	------------	--------	----------	---------

kubectl get  
pods,stacks,deploys,..

# UPGRADE TO KUBERNETES 1.14

"Found 1223 rows for 1 resource type in 148 clusters in 3.301 seconds."

[all](#) / [nodes](#)

## Nodes



Label Columns

Labels to show as columns (comma separated) or '\*' to show all labels

Label Selector

Label selector (label=value)

Filter

Roles=worker, Version=v1.14.6

Submit

[Show CPU/Memory Usage](#)

Cluster	Name	Status	Roles	Age	Version	Internal-IP	External-IP	OS-Image	Kernel-Version	Container- Runtime	Created
		Ready	worker	4h33m	v1.14.6			Ubuntu 18.04.3 LTS	4.15.0-1045-aws	docker://18.9.7	2019-08-27 12:27:12
		Ready	worker	17m	v1.14.6			Ubuntu 18.04.3 LTS	4.15.0-1045-aws	docker://18.9.7	2019-08-27 16:44:04
		Ready	worker	152m	v1.14.6			Ubuntu 18.04.3 LTS	4.15.0-1045-aws	docker://18.9.7	2019-08-27 14:29:10

# MULTIPLE CONDITIONS

all / all / pods

**Pods**   

Label Columns

 Labels to show as columns (comma separated) or %

Label Selector

 application=coredns

Filter

 Status!=Running,Status!=Completed

Submit

```
kubectl -l application=coredns
```

Cluster	Namespace	Name	Application	Component	Ready	Status	Restarts	Age	IP	Node	Nominated Node	Readiness Gates	CPU Usage	Memory Usage	Created
	kube-system	coredns-hm545	coredns	cluster-dns	2/3	CrashLoopBackOff	136	4h10m			<none>	<none>	6m	35 MiB	2019-08-20 03:51:58

Found 1 row for 1 resource type in 146 clusters in 2.358 seconds.

# SEARCHING ACROSS 140+ CLUSTERS

## Search

Search Text  Search!

Resource Types ☒ CronJob ☐ DaemonSet ☒ Deployment ☒ Ingress ☒ Namespace ☐ Node ☐ PlatformCredentialsSet ☐ Pod ☐ ReplicaSet ☒ Service ☒ StackSet ☒ StatefulSet

✕ unselect all

**etcd-operator (Deployment)**  
/cluster: namespaces/default/deployments/etcd-operator  
Created: 2018-10-18 13:23:39 source.zalan.c /etcd-operator:v0.9.2-master-2  
name: etcd-operator

**etcd-operator (Deployment)**  
/cluster: namespaces/wpi/deployments/etcd-operator  
Created: 2019-08-12 12:30:07 e.stups.zalan.dc /etcd-operator:v0.9.3  
application: deployment-id: d-e8yt17ub9hxyt513sr27w66ea environment: staging pipeline-id: l-7bic5kvi6khdadtqzq5hy3q version: master-7

**etcd-operator (Deployment)**  
/cluster: namespaces/default/deployments/etcd-operator  
Created: 2018-10-19 14:13:50 source.zalan.dc /etcd-operator:v0.9.2-master-3  
name: etcd-operator

**etcd-operator (Deployment)**  
/clusters: namespaces/default/deployments/etcd-operator  
Created: 2018-05-04 11:01:36 tups.zalan.dc /etcd-operator:v0.6.1-2  
app: etcd component: operator

**etcd-operator (Deployment)**  
/clusters: namespaces/incentives/deployments/etcd-operator  
Created: 2018-07-03 08:12:51 .zalan.dc /etcd-operator:v0.9.3  
application: deployment-id: d-so5ukevu2piyw5bdigzxc4gx3 environment: staging version: master-26

## CLUSTER RESOURCES

Namespaces

Nodes

PersistentVolumes

## CONTROLLERS

Deployments

CronJobs

Jobs

DaemonSets

StatefulSets

## POD MANAGEMENT

Ingresses

Services

Pods

ConfigMaps

## META

Resource Types

Events

all / all / pods

## Pods



Label Columns

Labels to show as columns (comma)

Label Selector

Label selector (label=value)

Filter

Status=Pending

Submit

All Pending Pods across all clusters

Cluster	Namespace	Name	Ready	Status	Restarts	Age	IP	Node	Nominated Node	Readiness Gates	Created
	kube-system	3d6b56-h8bml	0/1	Pending	0	77s	<none>	<none>	<none>	<none>	2019-08-07 17:30:04
	kube-system	5798-iv6c5	0/1	Pending	0	17m	<none>	<none>	<none>	<none>	2019-08-07 17:13:39
	default	fb7-8w66g	0/1	Pending	0	144m	<none>	<none>	<none>	<none>	2019-08-07 15:06:34
		676f-4x8g9	0/1	Pending	0	4h46m	<none>	<none>	<none>	<none>	2019-08-07 12:45:02
		676f-8jdvk	0/1	Pending	0	4h46m	<none>	<none>	<none>	<none>	2019-08-07 12:45:02
		676f-dmjg4	0/1	Pending	0	4h46m	<none>	<none>	<none>	<none>	2019-08-07 12:45:02
		676f-qj94v	0/1	Pending	0	4h46m	<none>	<none>	<none>	<none>	2019-08-07 12:45:02
		676f-rt4md	0/1	Pending	0	4h46m	<none>	<none>	<none>	<none>	2019-08-07 12:45:02

Clusters\_all

Search Kubernetes objects..

CLUSTER RESOURCES

Namespaces

Nodes

PersistentVolumes

CONTROLLERS

StackSets

Stacks

Deployments

CronJobs

Jobs

StatefulSets

POD MANAGEMENT

Ingresses

Services

Pods

ConfigMaps

CRDS

PlatformCredentialsSets

postgresqls

META

Resource Types

Events

all / pods

Pods

Namespace	Name	Application	Component	Ready	Status	Restarts	Age	IP	Node	Nominated Node	CPU Usage	Memory Usage
kube-system	audittrail-adapter-lt2kb	audittrail-adapter		1/1	Running	151	6d9h	172.31.0.70	ip-172-31-0-70.eu-central-1.compute.internal	<none>	1m	20 MiB
kube-system	audittrail-adapter-s6tvf	audittrail-adapter		1/1	Running	23	6d10h	172.31.18.89	ip-172-31-18-89.eu-central-1.compute.internal	<none>	6m	23 MiB
default	devcon-cluster-manager-74746ff998-gg4m9	devcon-cluster-manager		1/1	Running	13	12d	10.2.211.12	ip-172-31-12-12.eu-central-1.compute.internal	<none>	0m	83 MiB
default	devcon-cluster-manager-74746ff998-4zzwn	devcon-cluster-manager		1/1	Running	13	12d	10.2.3.13	ip-172-31-4-240.eu-central-1.compute.internal	<none>	0m	63 MiB
visibility	zmon-sentry-agent-7867b8d48b-rddbg	zmon-sentry-agent		1/1	Running	11	10d	10.2.211.124	ip-172-31-12-12.eu-central-1.compute.internal	<none>	0m	17 MiB
default			main	1/1	Running	10	67m	10.2.178.190	ip-172-31-16-111.eu-central-1.compute.internal	<none>	236m	212 MiB
default	cert-manager-ff4d7884d-s6jnm	cert-manager		1/1	Running	8	12d	10.2.220.9	ip-172-31-20-103.eu-central-1.compute.internal	<none>	5m	51 MiB

[codeberg.org/hjacobs/kube-web-view](https://codeberg.org/hjacobs/kube-web-view)

# EMERGENCY ACCESS SERVICE

## Emergency access by referencing Incident

```
zkubectl cluster-access request \  
  --emergency -i INC REASON
```



## Privileged production access via 4-eyes

```
zkubectl cluster-access request REASON  
zkubectl cluster-access approve USERNAME
```



# KUBERNETES FAILURE STORIES

Learning about production pitfalls!



zerkms commented 14 days ago



@hjacobs while you're here I wanted to thank you and tell that this is the most important repository to follow for those who run their kubernetes clusters :-)



1



1

<https://k8s.af>

A compiled list of links to public failure stories related to Webometrics, most recent publications at the top.

- [8 Kubernetes c# code - Part 1 - May 2019](#)



<https://k8s.af>

# 42 stories so far

# INTERNAL TICKETS BASED ON FAILURE STORIES

## flannel setup to fully random port allocation #2213

Open

opened this issue 5 days ago · 5 comments



commented 5 days ago • edited ▼

Member



source: <https://tech.xing.com/a-reason-for-unexplained-connection-timeouts-on-kubernetes-docker-abd041cf7e02>

relevant part:

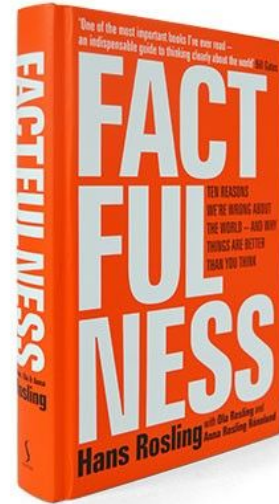
Activating full random port allocation on Kubernetes

The `NF_NAT_RANGE_PROTO_RANDOM_FULLY` flag needs to be set on masquerading rules. On our Kubernetes setup, Flannel is responsible for adding those rules. It uses iptables which it builds from the source code during the Docker image build. The iptables tool doesn't support setting this flag but we've committed a small patch that was merged (not released) and adds this feature.

# FACTFULNESS

*Things can be both better and bad!*

How would failure stories for  
your non-K8s infra look like?



<https://k8s.af>

## "I'M TORN ON THIS LIST"

It is important to learn from the mistakes from others, so I like it [...], BUT it feels like folks are **using these examples as reasons to stay away from Kubernetes.**

There could be a [...] **larger list of system failures where K8s is not involved.** Similarly there could [...] be a list of "Encryption Failure Stories" but that doesn't mean we shouldn't encrypt things. [...]

# FAILURE STORIES CAN BE FOUND EVERYWHERE



**Henning Jacobs** @try\_except\_ · 16. Okt.

So who starts the list with #AWS ECS failure stories? 🙄



**Sebastian Herzberg** @shrzbrg · 16. Okt.

Within a couple of weeks we experienced the same major incident on K8s and ECS. Yesterdays ECS optimized AMI had „releasever=nightly“ in yum.conf which pointed to a repository that ordinary users can not access. Resulting in limited ECS cluster capacity for us.



1



5





**Ivan Pedrazas**

@ipedrazas

"Daddy, tell me a horror story"

"SSL with Istio and Kubernetes"

"Is it as bad as the NFS monster one?"

"Oh no, nothing is worse than the NFS monster"

[Tweet übersetzen](#)

11:45 vorm. · 29. März 2018 · [Twitter Web Client](#)

---

**209** Retweets   **614** „Gefällt mir“-Angaben



## WHY KUBERNETES?

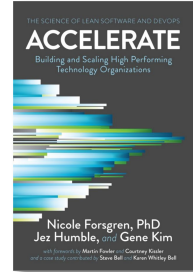
# WHY KUBERNETES?

- provides enough **abstractions** (StatefulSet, CronJob, ..)
- provides **consistency** (API spec/status)
- is **extensible** (annotations, CRDs, API aggreg.)
- certain **compatibility** guarantee (versioning)
- widely **adopted** (all cloud providers)
- works across environments and implementations

# WHY KUBERNETES?

(for Zalando)

- **Efficiency**
- **Common Operational Model**
- **Developer Experience**
- **Cloud Provider Independent**
- **Compliance and Security**
- **Talent**



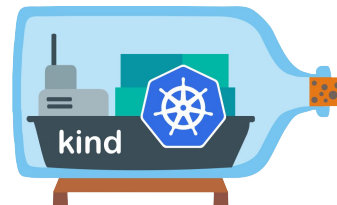
# COMPLEXITY FOR GOOGLE-SCALE INFRA?

- Managed DO cluster: 4 minutes
- K3s single node: 2 minutes



install.sh 838B

```
1 #!/bin/bash
2
3 # Install k3s
4 curl -sfL https://get.k3s.io | sh -
```





**Corey Quinn**

@QuinnyPig

Follow



Nuclear hot take: nobody will care about Kubernetes in five years.

**CZnative @ home** @pczarkowski

Replying to @tmclaughbos @iteration1 @behemphi

As I keep telling people, if you have a kubernetes strategy you've already failed. Kubernetes should be an implementation detail at the tactical level to deal with the strategic imperative of solving the problems that are halting the flow of money.

6:32 PM - 6 Feb 2019

97 Retweets 439 Likes



41



97



439



# MAYBE THAT'S GOOD?



**Kuberkus 1.16.0 is a Wrap**

@fuzzychef

Antwort an @QuinnyPig

Speaking as a kubernetes dev, that's a victory condition. It means that Kube becomes so ubiquitous, and so easy, that's it's ignorable.

[Tweet übersetzen](#)

7:06 nachm. · 7. Feb. 2019 · [Twitter Web Client](#)

1 Retweet 18 „Gefällt mir“-Angaben

# OPEN SOURCE & MORE

## Kubernetes Web View

[codeberg.org/hjacobs/kube-web-view](https://codeberg.org/hjacobs/kube-web-view)

## Skipper HTTP Router & Ingress controller

[github.com/zalando/skipper](https://github.com/zalando/skipper)

## Kubernetes Janitor

[github.com/hjacobs/kube-janitor](https://github.com/hjacobs/kube-janitor)

## Postgres Operator

[github.com/zalando-incubator/postgres-operator](https://github.com/zalando-incubator/postgres-operator)

## More Zalando Tech Talks

[github.com/zalando/public-presentations](https://github.com/zalando/public-presentations)





# QUESTIONS?



**HENNING JACOBS**  
SENIOR PRINCIPAL



[henning@zalando.de](mailto:henning@zalando.de)

[@try\\_except](#)

Illustrations by [@01k](#)



*Please*

**Remember to  
rate this session**

*Thank you!*

