What happens when something goes wrong?

KubeCon Austin
2017-12-08

Marek Grabowski & Tina Zhang
gmarek@google.com ttz@google.com
Site Reliability Engineers in Google London
GH: @gmarek & @ttz21

Who are we?

"Fundamentally, it's what happens when you ask a software engineer to design an operations function."

Ben Treynor Sloss, Vice President, Google Engineering, founder of Google SRE

What will we be speaking about?

Automation

Pod Replication

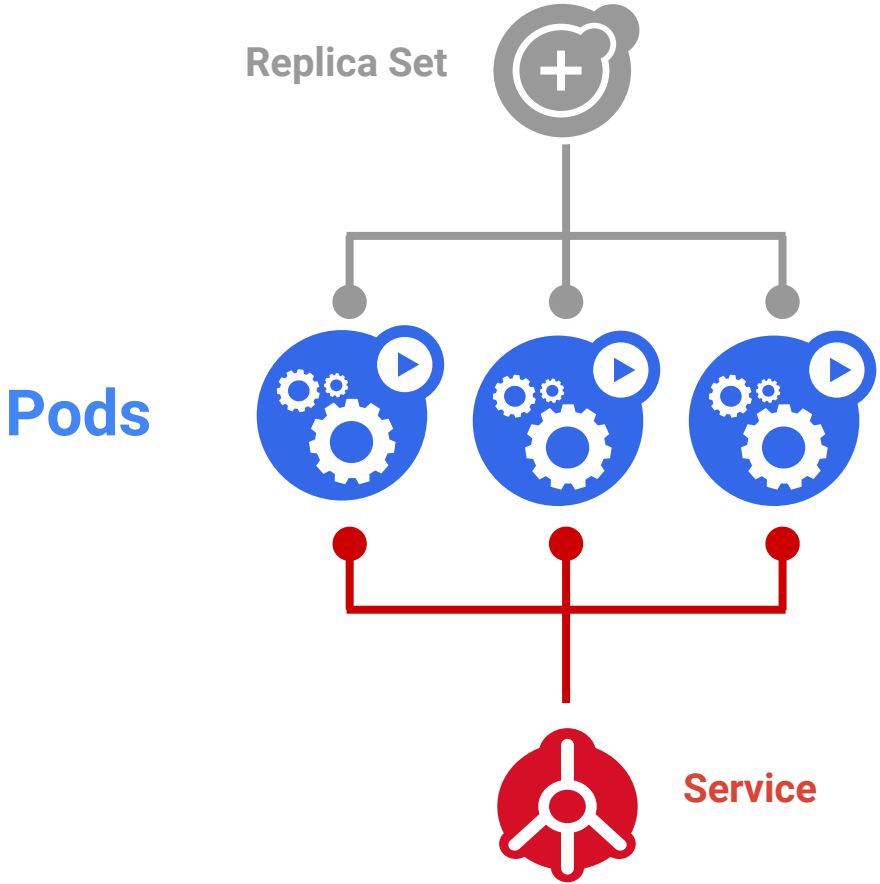

Container Lifecycle management

Kubelet

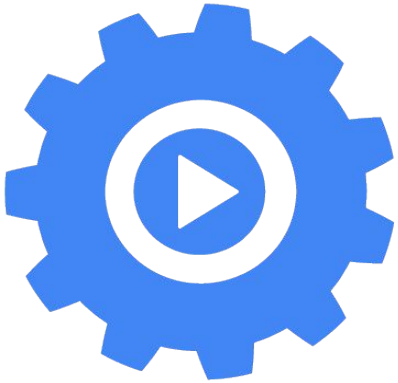

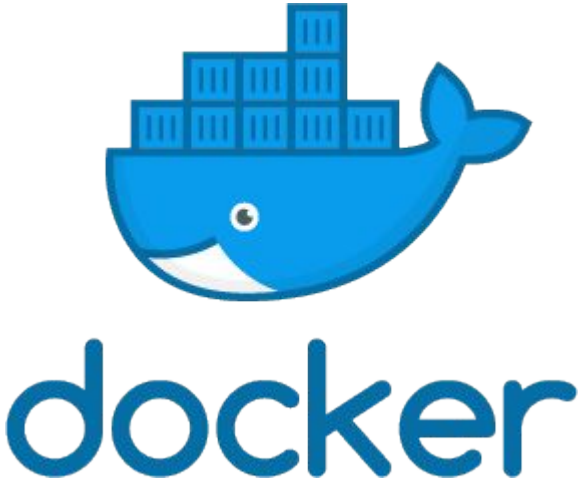

Liveness Probes

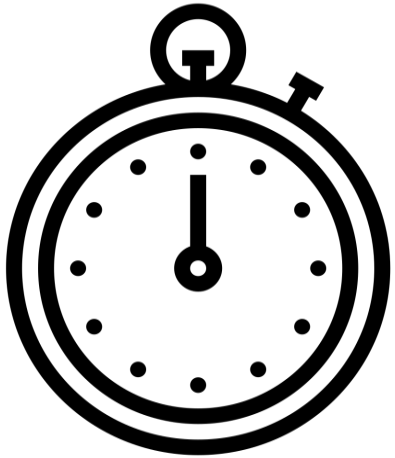

Probe Types

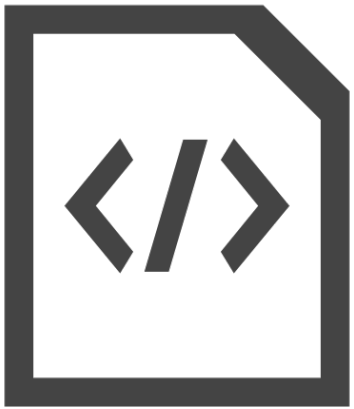

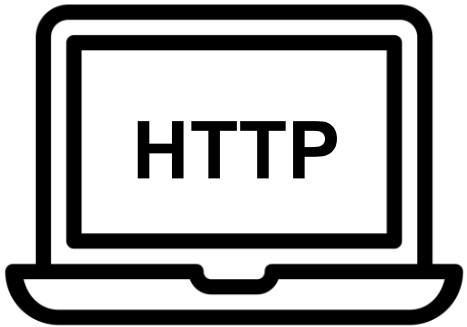

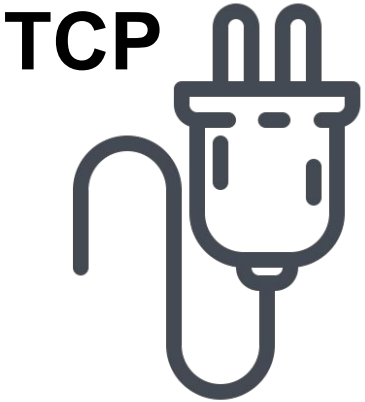

Configuring Liveness Probes (GA)

```
livenessProbe:
  httpGet:
    path: /healthz
    port: 8080
    httpHeaders:
    - name: X-Custom-Header
      value: Awesome
  initialDelaySeconds: 3
  periodSeconds: 3
  timeoutSeconds: 3
  successThreshold: 1
  failureThreshold: 3
```

```
tcpSocket:
   port: 8080
```

```
exec:
    command:
    - cat
    - /tmp/healthy
```

If Pod itself is not broken then what?

Environment?

Environmental problems

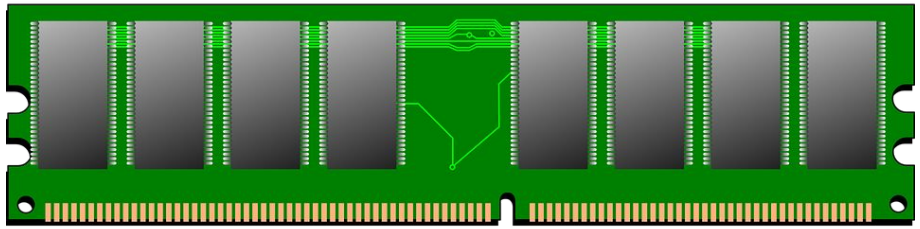

How is the environment getting fixed?

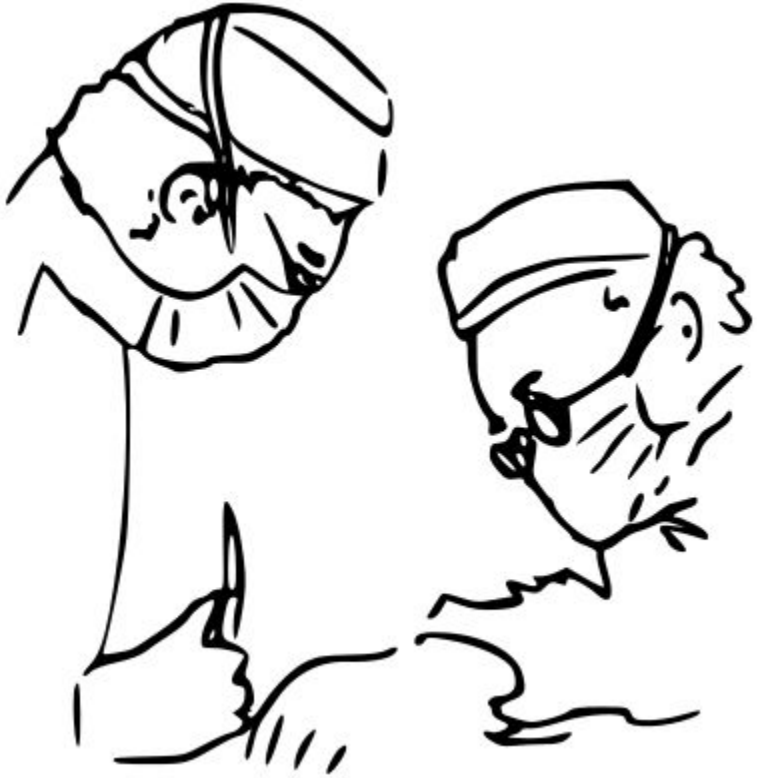

How is the environment getting fixed?

Environment is fine. What now?

Machine?

Diagnosis Daemons

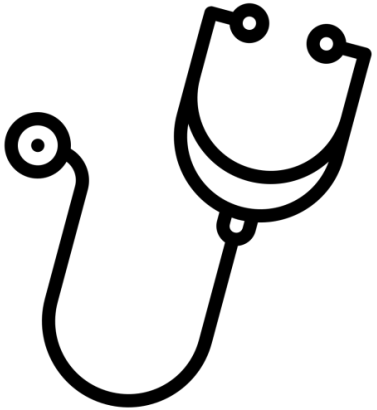

Node Problem Detector

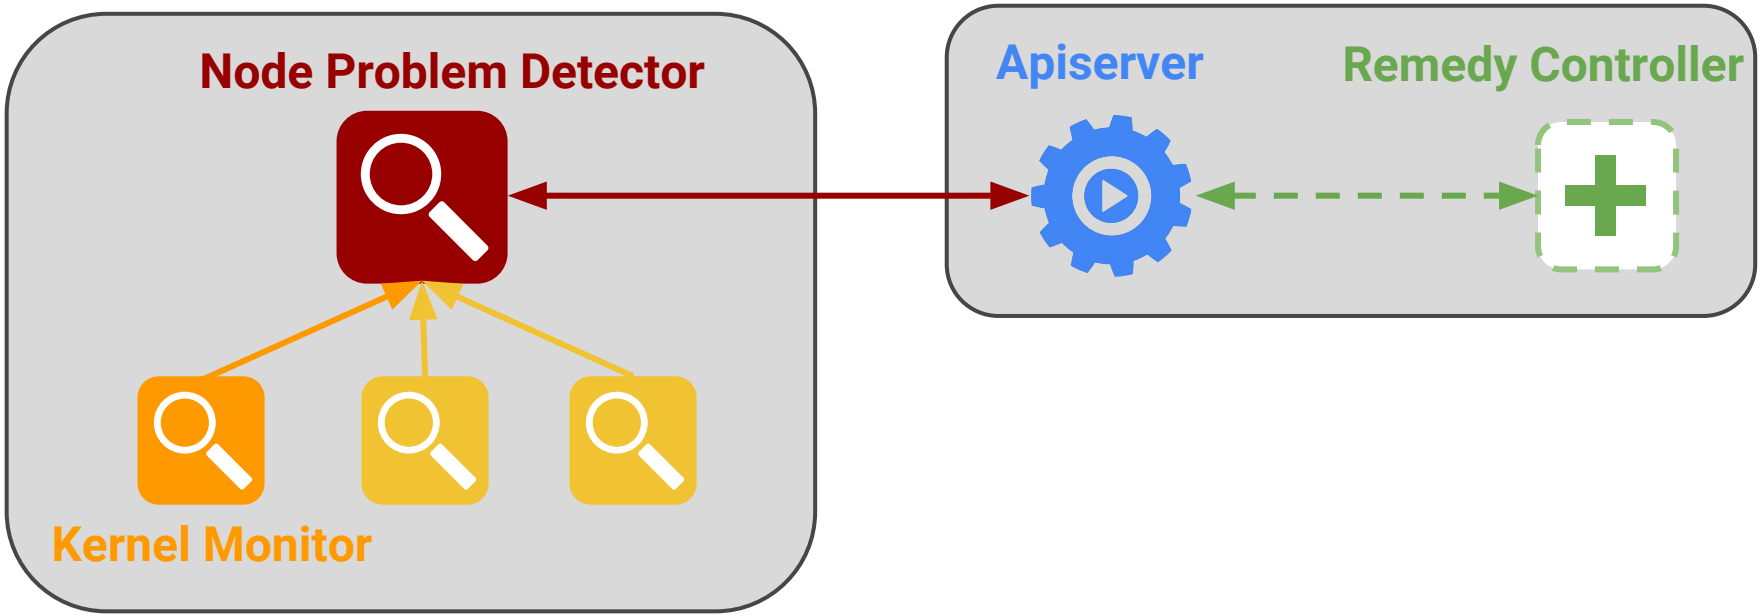

Hardware or software?

How does Kubernetes check this?

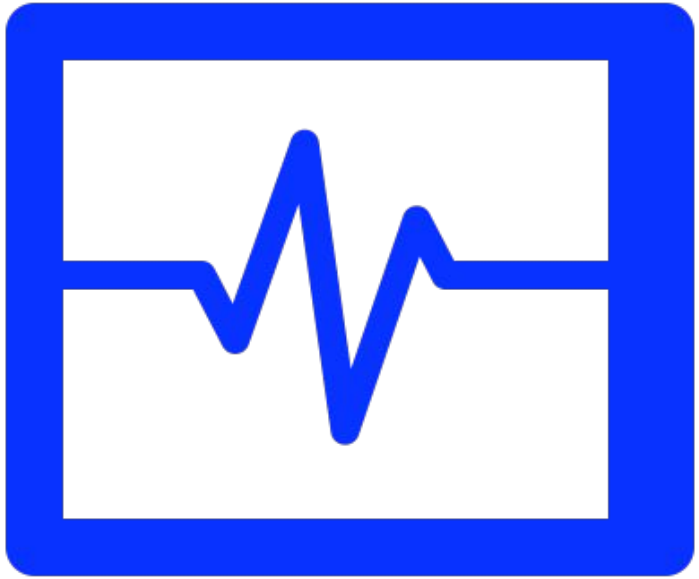

How does the system handle unresponsive Nodes?

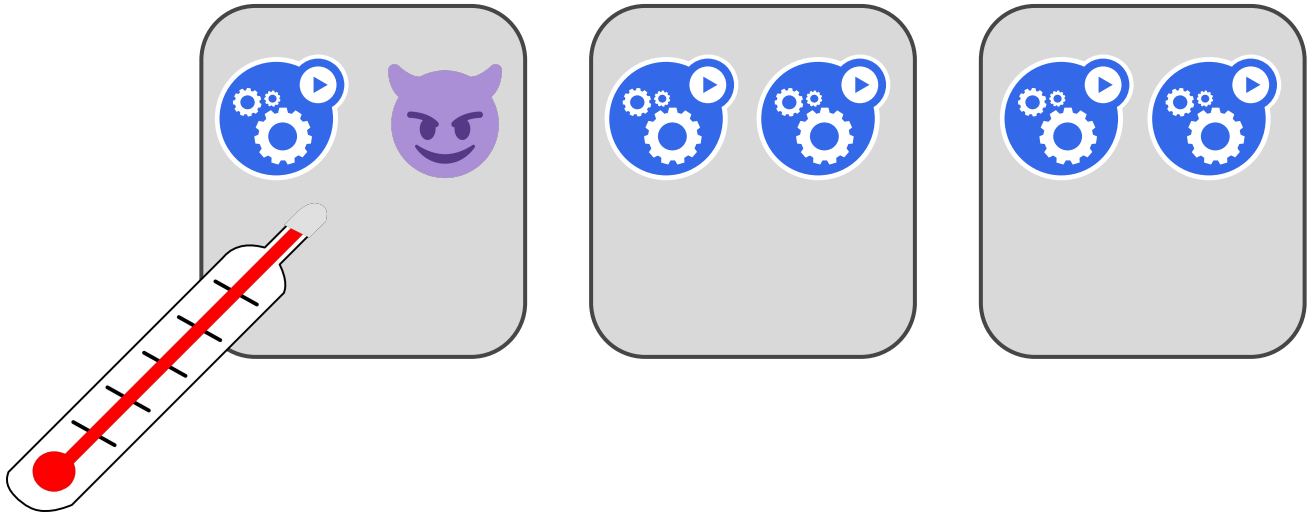

How does the system handle unresponsive Nodes?

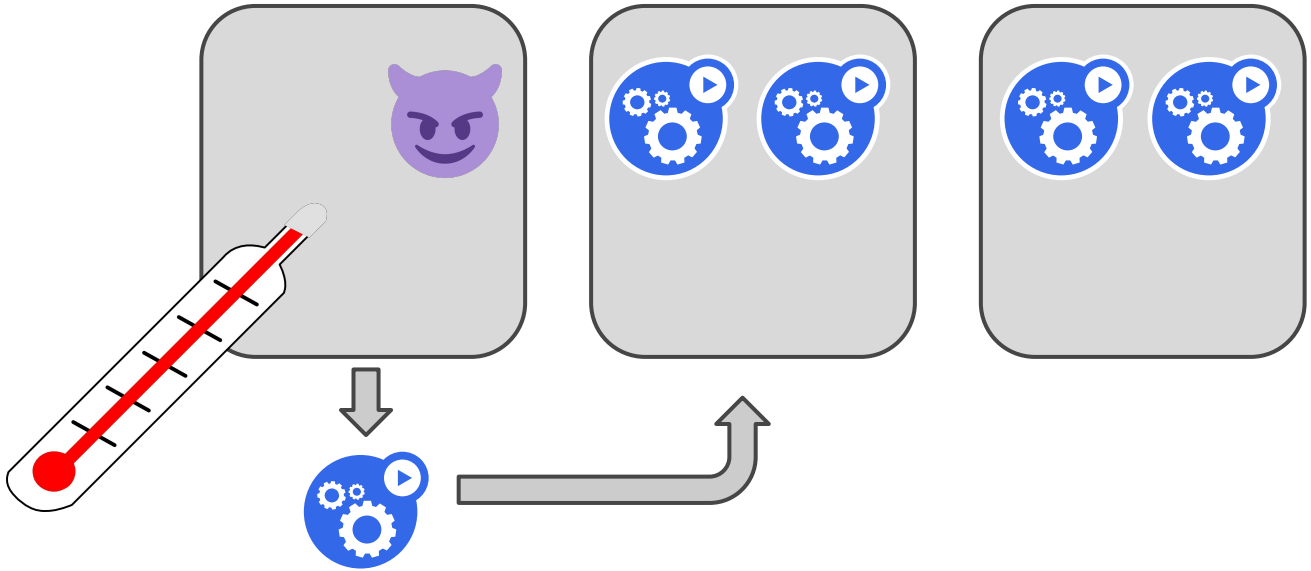

What happens when everything is very wrong?

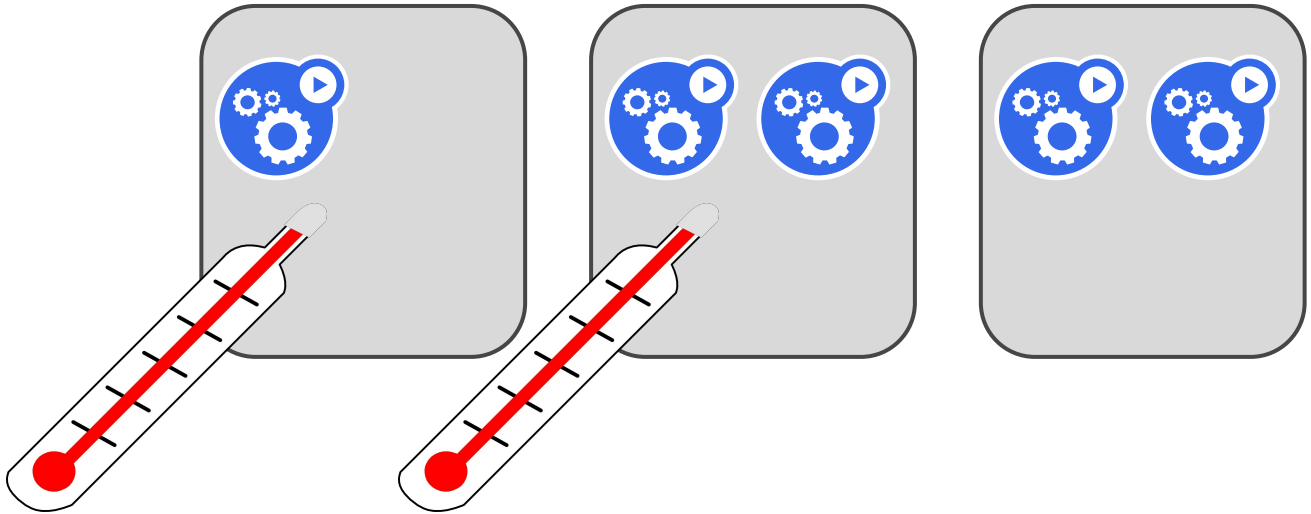

What happens when everything is very wrong?

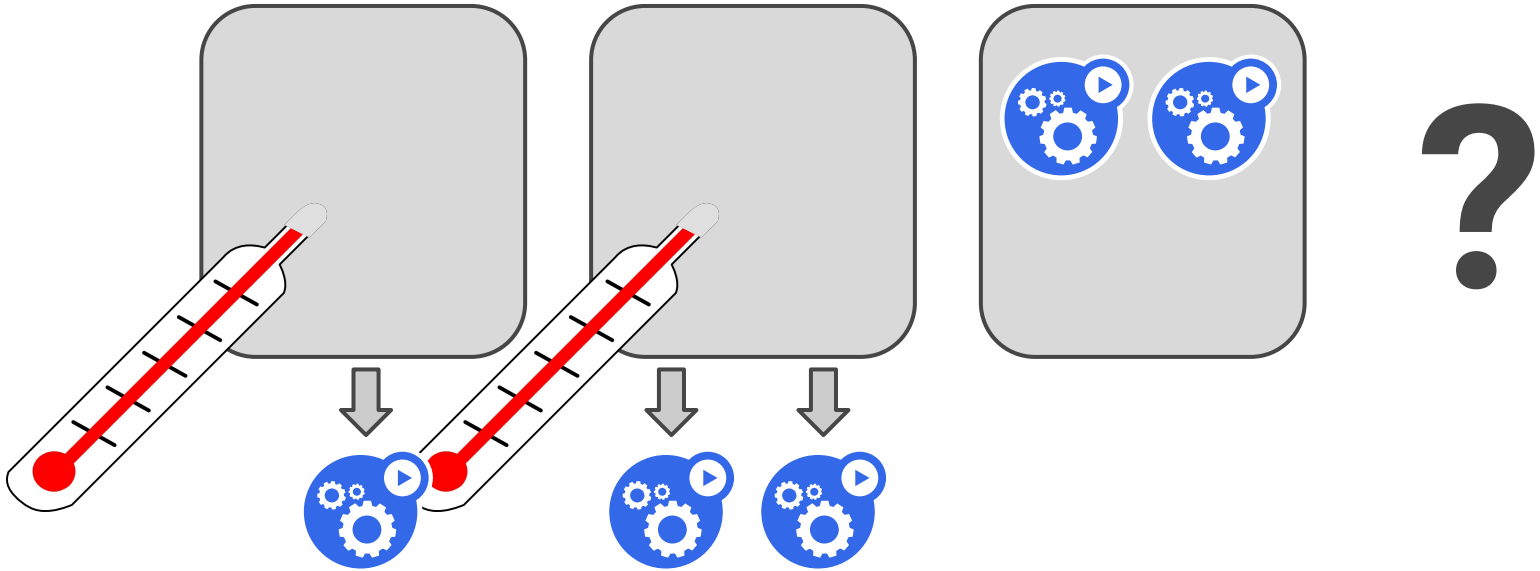

What happens when everything is very wrong?

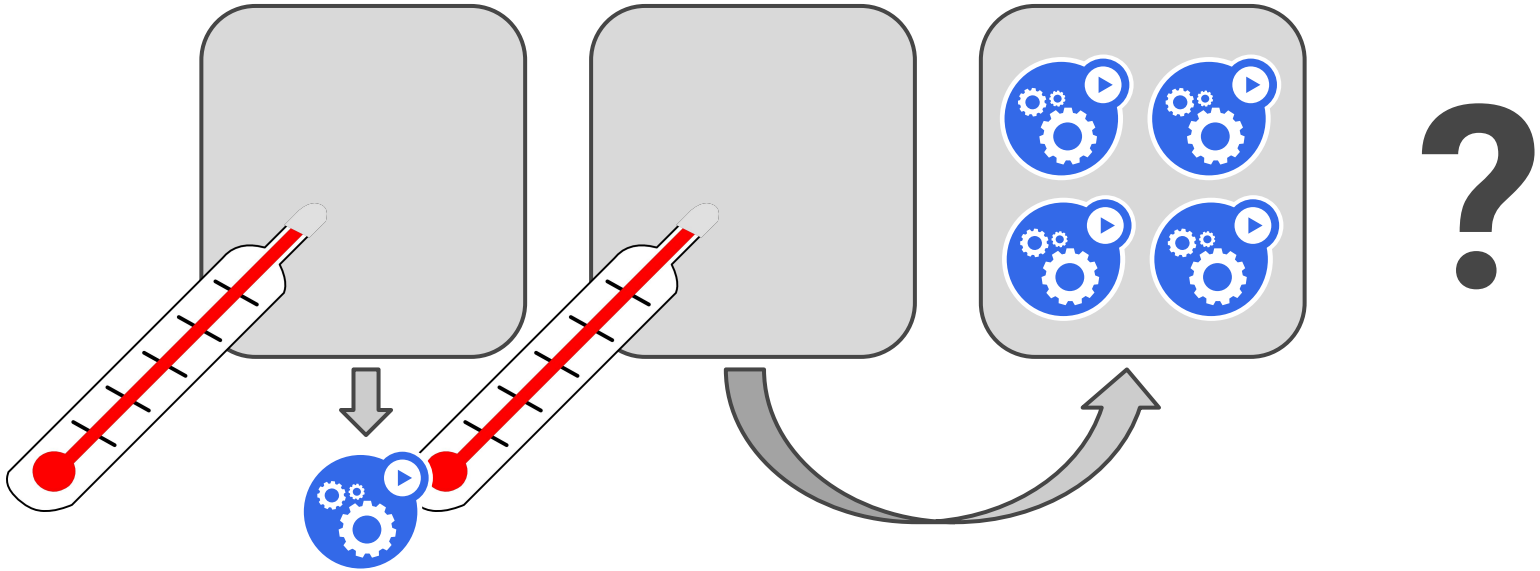

What happens when everything is very wrong?

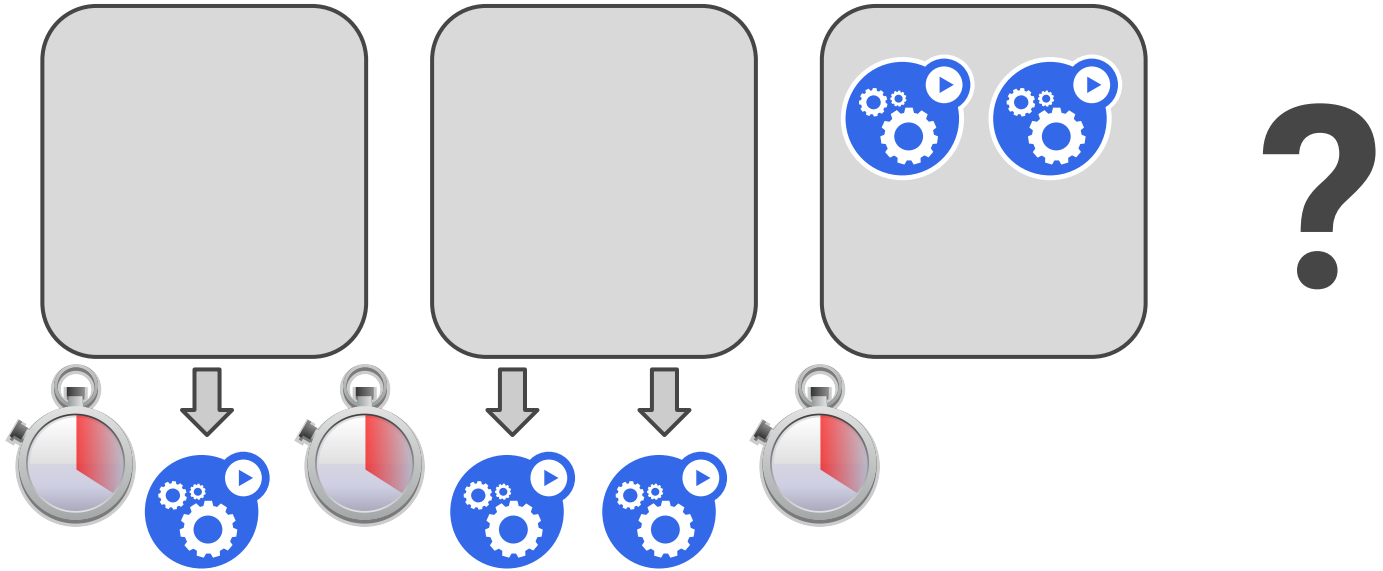

What happens when everything is very very wrong?

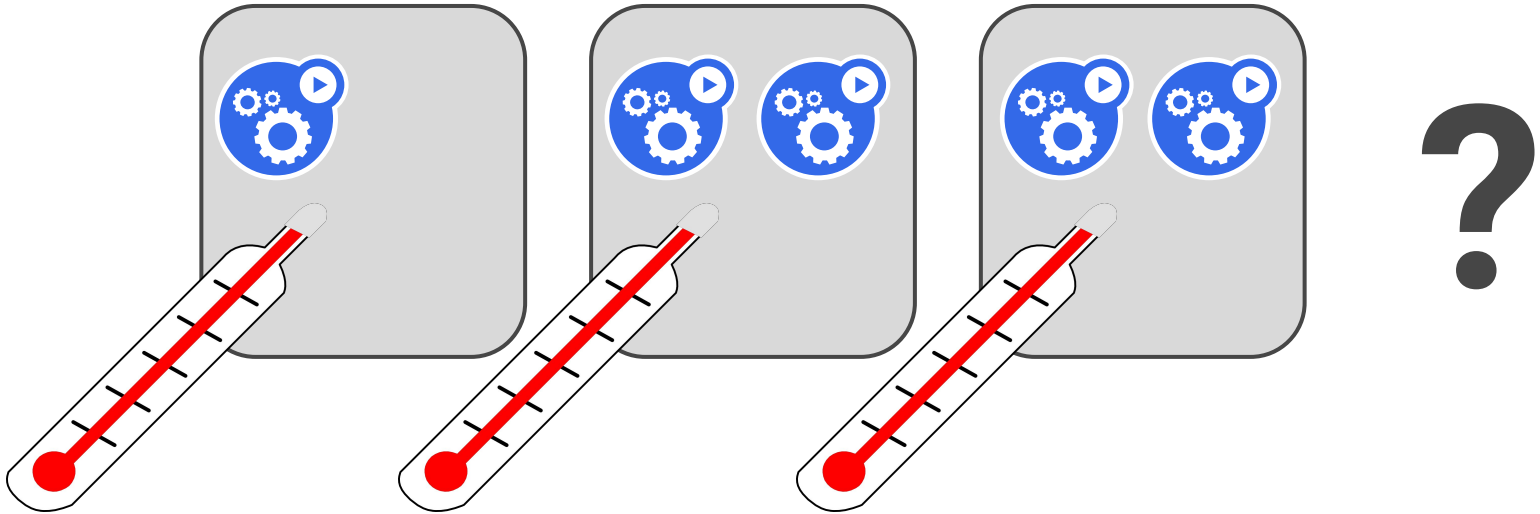

Configurability

Taints and Toleration example (alpha)

```
Node:
  ...
  spec:
    ...
    taints:
    - key:
"node.kubernetes.io/unreachable"
      effect: "NoExecute"
```

```
Pod:
  ...
  spec:
    ...
    tolerations:
    - key:
"node.kubernetes.io/unreachable"
      operator: "Exists"
      effect: "NoExecute"
      tolerationSeconds: 42
```

Control plane issues

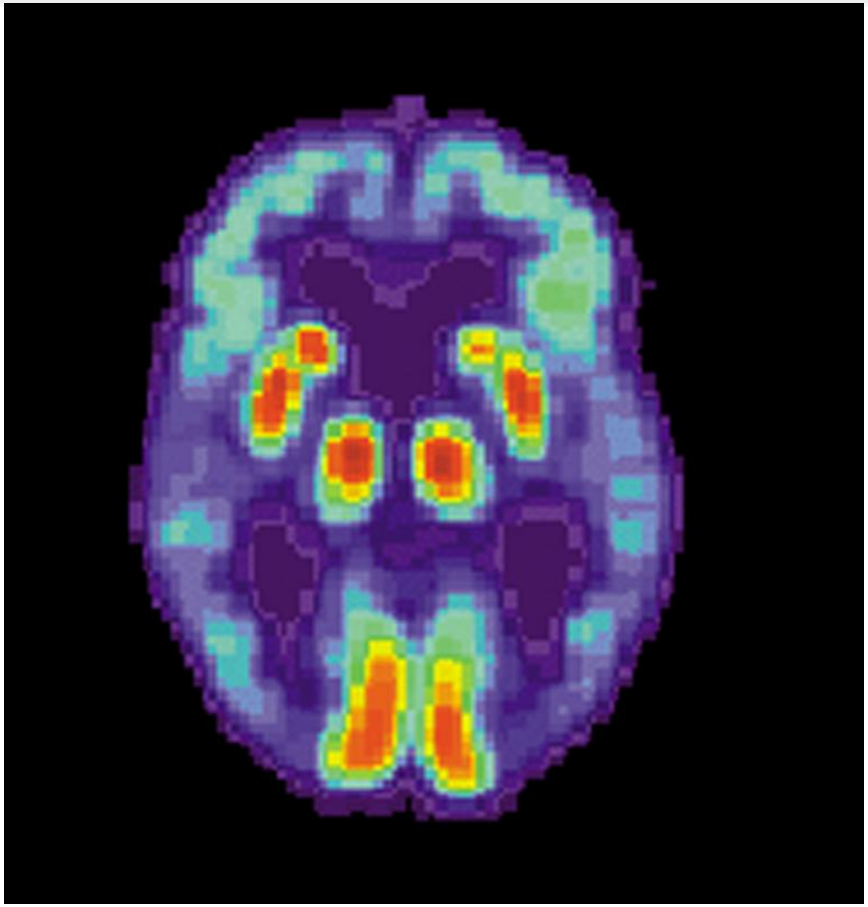

Kubernetes = self-healing clusters

Kubelet

Node Problem Detector

Liveness Probes

NodeController evictions

Nothing is perfect

What does GKE do for you when things go wrong?

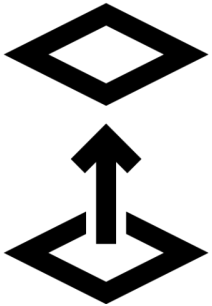

Thank you for listening