webinale WEB WEEK by devmio

IPC by devmio

Digital Sovereignty
where to put my Kubernetes

Robert Lemke · Flownative

STATUSPAGE now
incident opened automatically

GRAFANA now
error rate 100% on euw1

SLACK now
»anyone awake?«

PAGERDUTY now
escalation — nobody acknowledged

UPTIME now
still down (12m)

CILIUM now
egress gateway flapping

We did everything **right**

Flownative


Robert Lemke
Co-founder Flownative
Founder Neos CMS Project
Creating software since 1985
Happy PHP user since 1998

Where to put my **Kubernetes?**

...suddenly a political question

Flownative

S sitegeist

Projects
Database Servers
Redis Servers
Domains
Backups

Overview | Builds | Webhooks | Add-Ons | Plan

PRODUCTION

Beach Nürnberg

Neos Frontend

main

Group **Frontend** · 3 dependents

Manage | Clone

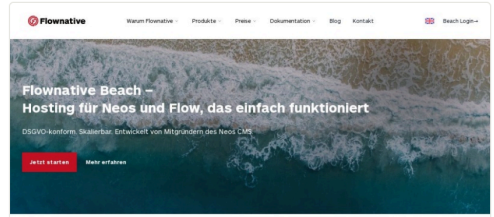

Einblick in Kundenprojekte

Bevor wir uns für ein neues Produkt oder einen Dienstleister entscheiden, schauen wir uns an, wie der Anbieter so "tickt".

Recent Activity

Show all ›

Deploy v3.2.1	vor 3d
PHP 8.4 enabled	vor 5d
Scale 1 → 2 Replicas	vor 8d

Last Hour

Online

CPU	Mem	Net
23%	68%	12 MB/s

2/2	Large	Auto
REPLICAS	TYP	SCALING

Monitoring Grafana

Scale

Releases

Check for new commits | Auto-Deploy ON

LIVE

v3.3.0

Feat: Add dark mode support

Committed by John Doe, 2 days ago

main
abc1234
Flow 8.2.1 · Neos 8.18

Deployed by Jane Doe, 2 days ago

Re-Deploy

PREVIOUS

v3.2.1

Fix: Resolve OOM in image processing

Committed by Jane Doe, 5 days ago

main
def5678
Flow 8.2.1 · Neos 8.18

Digital sovereignty is not an **add-on**

It's the result of concrete engineering decisions.

Flownative

Sovereignty is a spectrum, not a switch

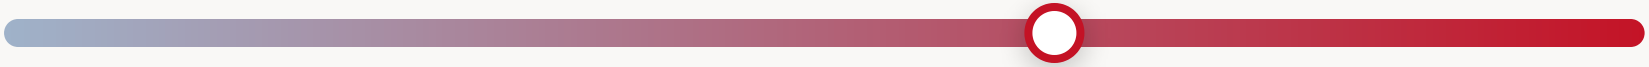

US hyperscaler
maximum convenience

Self-hosted in the EU
maximum control

Where your data sits is not the same as **sovereignty**

The US CLOUD Act reaches US companies – wherever the servers stand.

Flownative

»Trump can shut down any organization that depends on US technology – by decree.«

Open Source Business Alliance, after Microsoft blocked the ICC prosecutor's mailbox (2025)

We mix up three different fears

data protection

independence

lock-in

Our journey off the hyperscalers

2015

Google Cloud

2021

Open Telekom Cloud

2025

Hetzner Cloud

2026

Hetzner Bare Metal

Google Cloud was a great start

...but expensive and we feared the lock-in

Flownative

GDPR made it real

Our most important customer processes
personal data of millions of persons.

Flownative

Is there a **European** hyperscaler?

Not really hyper

Nobody has the breadth of AWS, Azure or Google.

~~Open Telekom Cloud~~

T-Cloud Public

... no capacity, ancient Kubernetes.

Flownative

So we learned to do it **ourselves**

...with three people

Flownative

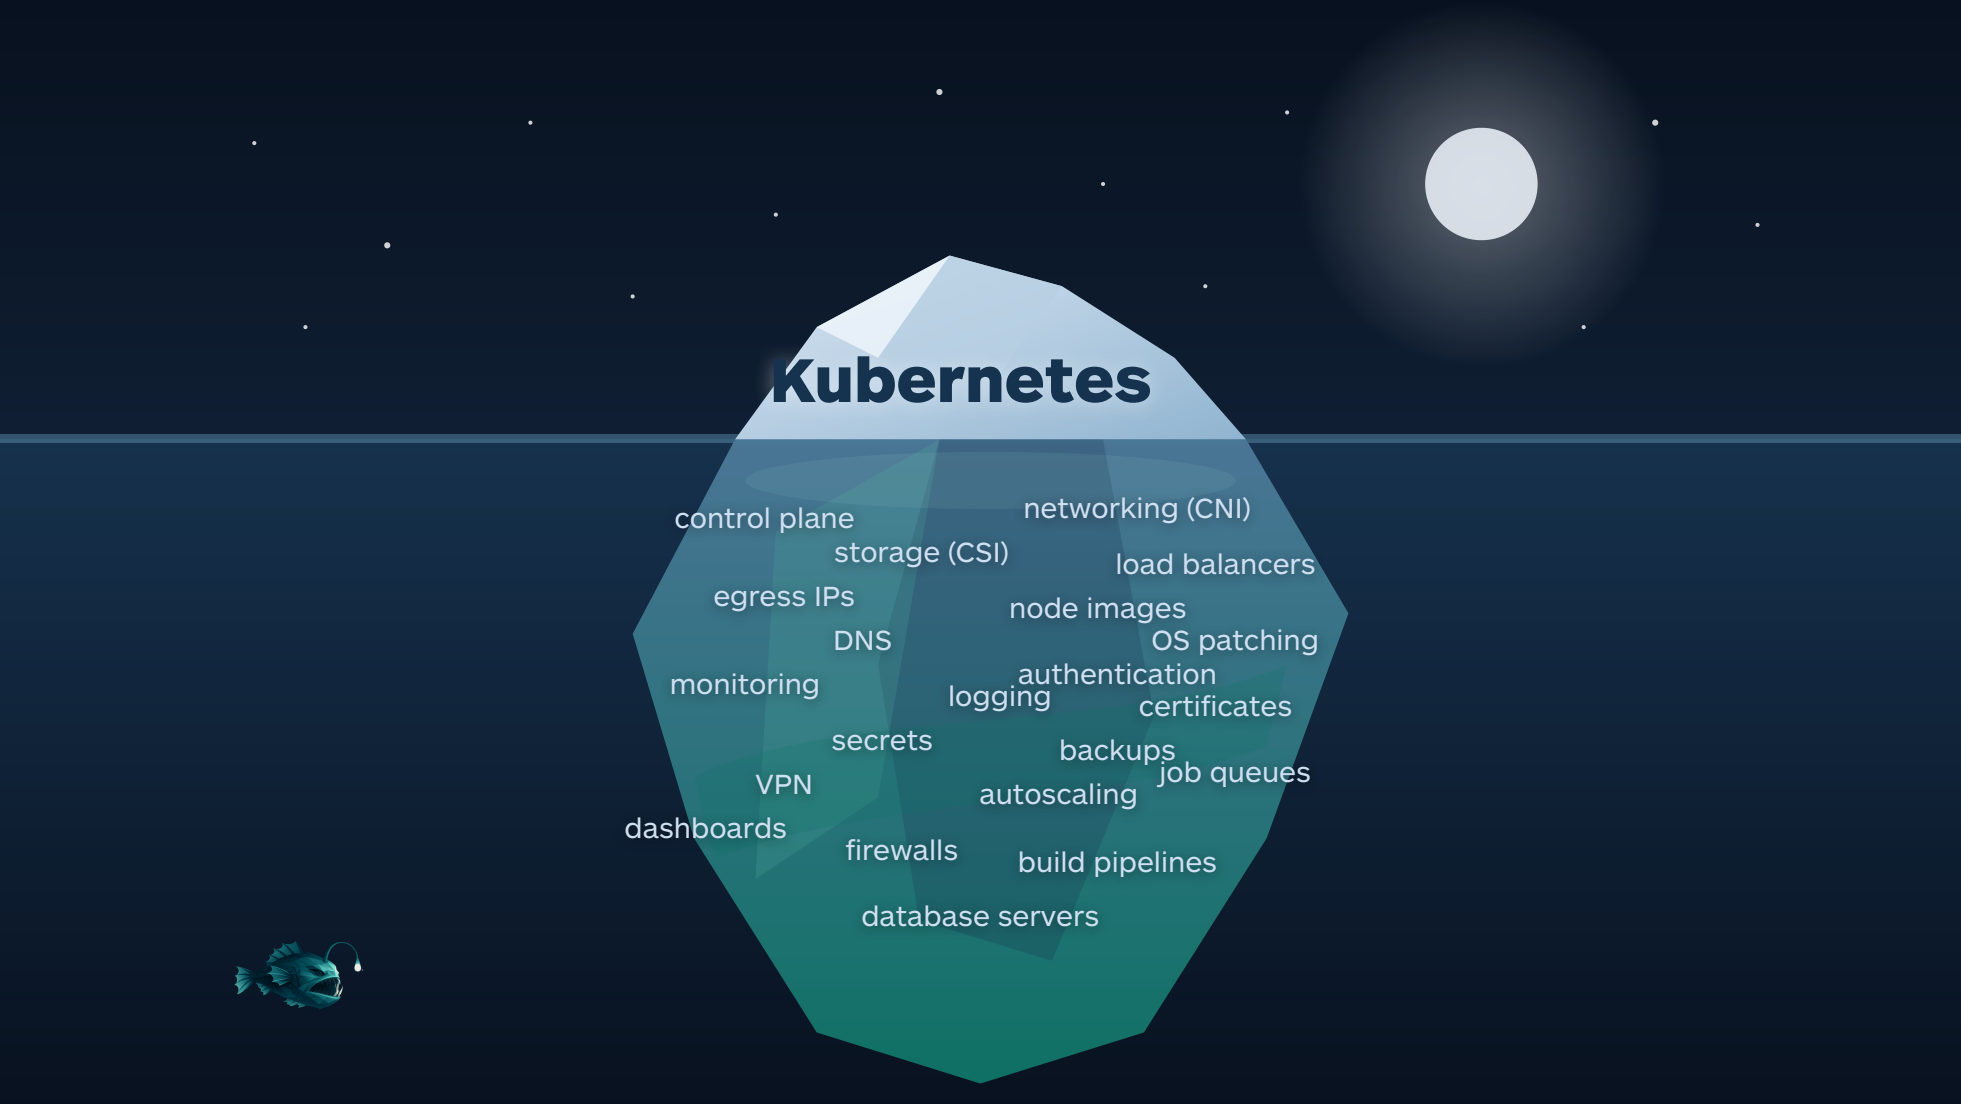
Kubernetes
control plane
networking (CNI)
storage (CSI)
load balancers
egress IPs
node images
DNS
OS patching
monitoring
authentication
logging
certificates
secrets
backups
job queues
VPN
autoscaling
dashboards
firewalls
build pipelines
database servers

You stop using a platform.
You **become** one.

Cluster API

Kubernetes managed by Kubernetes.

A management cluster runs controllers that create and heal workload clusters — declaratively, as YAML.

Routing egress through a clean IP

```yaml
apiVersion: cluster.x-k8s.io/v1beta1
kind: Cluster
metadata:
  name: flownative-beach-abc123
  namespace: flownative-beach-abc123
spec:
  topology:
    class: flownative-beach-hcloud-202601291v1
    classNamespace: beach-workload-clusters
    version: "1.34.5"
    controlPlane:
      replicas: 3
    workers:
      machineDeployments:
        - class: hcloud-worker
          …
    variables:
      - name: controlPlaneRegions
        value: ["nbg1"]
        …
```

Talos: an OS with **no shell** — only an API

Immutable, minimal, built only for Kubernetes. No SSH, no node images to maintain.

We're not on one stack — we're migrating between two

Cluster API + Debian images
Hetzner Cloud · since 2025

migrating →

Talos Linux
Bare Metal · now

CAPI provisions VMs from node images we build and patch. Talos is an immutable OS — nothing to image, nothing to SSH into.

Either way, you still assemble all this

Cilium	networking (CNI), kube-proxy replacement, WireGuard encryption
Storage	Hetzner Volumes (cloud) or Piraeus/LINSTOR on NVMe (bare metal)
hcloud-CCM*	load balancers & node lifecycle (*forked — see later)
FluxCD	GitOps delivery
Monitoring	metrics & logs — Mimir, Loki, Grafana

Hetzner Cloud vs. Bare Metal — the trade-off we lived

Hetzner Cloud

+ elastic autoscaling
+ new servers in minutes
– weak disk I/O for databases
– capacity limits, pricier per core

Hetzner Bare Metal

+ ~35% cheaper per core, fast local NVMe + ECC
+ stable and predictable
– you own replication & capacity planning
– more hardware gotchas

+30-60%

the markup for managed Kubernetes.
Worth it? That depends on your team.

Back to that night

We had inherited a **spammer's** IP address

Scale up, and you get whatever address is free – including someone else's bad reputation.

In the cloud, you don't own your **reputation**

Our control panel still ran on Google

...and our Hetzner cluster couldn't reach it

The fix

...assigned automatically by our own operator

...and that was just the **first** surprise

We had to **fork** Hetzner's cloud controller

Its Robot API is account-wide — reinstall, rescue, even cancel any server. We refused to put that token in the cluster.

In some data centers there is **no remote power-on**

...only Hetzner support can switch the server back on

At 256 GB of RAM, a silent **bit-flip** is a certainty

So every node runs ECC memory — one more thing a hyperscaler quietly handled for you.

11 July 2025, 02:45

A third of our cluster — gone.

Kubernetes tried to self-heal. It couldn't.

```
$ create server …
Error: failed to create server: "resource_unavailable"
```

Hetzner had no spare capacity in Nuremberg.

Right now, compute is scarce – **everywhere**

Even Azure hits capacity limits. The rush to sovereign clouds (and AI) makes it worse.

The same problem that pushed us off OTC

...came back at Hetzner

What we changed

Keep buffer capacity

Bare metal for stability

Spread across regions

So – should **you** do this?

Three of several options

MOST CONVENIENT

US hyperscaler

Everything managed. Least sovereignty. Not really GDRP compliant.

BALANCED

EU managed K8s

Scaleway · IONOS · STACKIT. Sovereign, at a premium.

MOST CONTROL

EU self-hosted

Hetzner. Cheapest, most control — you run the platform.

1,500 € → 33 €

17.5 TB of egress. Google Cloud vs. Hetzner. Same traffic.

Three people is the **minimum**

Not for the workload — for the on-call rotation.

...and only because we automate **everything**

self-healing systems · 3 x 25 years experience

Don't self-host if...

- small team, no platform engineer
- Kubernetes isn't your edge
- you just need it to work

Do it, if...

- you have (or want) a platform team
- scale makes cost matter
- control or sovereignty is a hard requirement

Whatever you choose:
design for exit

Kubernetes + GitOps + IaC = reversible

Digital sovereignty is not an **add-on**

Flownative

It's worth it — go in with **open eyes**

Sovereignty has a price. Pay it on purpose.

Flownative

Take this home

Sovereignty is a spectrum — **choose on purpose**

No EU hyperscaler — **you bring the platform**

Self-host only with a team that can carry it

Design for exit from day one

Thank you

Robert Lemke · robert@flownative.com ·

Questions?

Flownative

Bonus

Backup slides — real config from our cluster, in case you asked.

Routing egress through a clean IP

```yaml
apiVersion: cilium.io/v2
kind: CiliumEgressGatewayPolicy
metadata:
  name: default
spec:
  destinationCIDRs: ["0.0.0.0/0"]
  egressGateway:
    nodeSelector:
      matchExpressions:
        - key: "beach.flownative.com/cilium-egress-gateway"
          operator: Exists
  selectors:
    - namespaceSelector:        # all app namespaces, except
        matchExpressions:       #   kube-system, flux-system, monitoring
          - key: kubernetes.io/metadata.name
            operator: NotIn
            values: [kube-system, flux-system, monitoring]
```

The gotcha that takes your public IP

```yaml
machine:
  network:
    interfaces:
      - interface: enp7s0    # real name — NOT eth0
        dhcp: true           # omit this in a patch → server loses its public IP
        vlans:
          - vlanId: 4000
            addresses: ["10.0.1.10/24"]
            mtu: 1400
```

Networking math you now own

```
# Hetzner vSwitch physical MTU:        1400
# minus WireGuard encryption overhead:  -60
routingMode: native
MTU: 1340      # get this wrong → silent packet loss, debugging hell
```