Kubernetes Community Days
MUNICH 2023

SPONSORS

metalstack cloud
ELASTICBRAINS
APE FACTORY
QA|WARE
Google Cloud
Grafana Labs
AppsCode
mkdev
ISOVALENT
camp to camp
INNOVATIVE SOLUTIONS BY OPEN SOURCE EXPERTS
StormForge
okteto
lumigo
EXOSCALE
Red Hat
SysEleven
REPLY
white duck
blueshoe

Speaker: Arik Grahl
Company: SysEleven

HONEY, I SHRUNK THE DATACENTER

Operating Bare-Metal Kubernetes at Home for Fun and Data Sovereignty

echo $(whoami)

arik-grahl.de/talks/kcd-munich-2023

```yaml
---
apiVersion: v1
kind: Person
metadata:
  name: Arik Grahl
  pronouns: he/him
spec:
  work:
    company: SysEleven
    role: SRE
  background: full-stack web developer
  passions:
    - FOSS advocate
    - data hoarder and self-hoster
  contact:
    web: www.arik-grahl.de
    gitlab: gitlab.com/arik.grahl
    github: github.com/arikgrahl
    linkedin: linkedin.com/in/arikgrahl
    mastodon: chaos.social/@arikgrahl
```

Why Data Sovereignty? (1/3)

data sovereignty as an implementation of privacy

Why Data Sovereignty? (2/3)

The Guardian

The Cambridge Analytica Files
Cambridge Analytica

Revealed: 50 million Facebook profiles harvested for Cambridge Analytica in major data breach

Whistleblower describes how firm linked to former Trump adviser Steve Bannon compiled user data to target American voters

- 'I made Steve Bannon's psychological warfare tool': meet the data war whistleblower
- Mark Zuckerberg breaks silence on Cambridge Analytica

Surveillance Capitalism and the Challenge of Collective Action

New Labor Forum
2019, Vol. 28(1) 10-29

DOI: 10.1177/1095796018819461

Shoshana Zuboff[1]

Keywords
capitalism, surveillance, digital technologies, democracy, collective action, twenty-first-century society, social inequality, power, internet, surveillance capitalism

What Is Surveillance Capitalism?

In our time, surveillance capitalism repeats capitalism's "original sin" of primitive accumulation. It revives Karl Marx's old image of capitalism as a vampire that feeds on labor, but with an unexpected turn. Instead of claiming work (or land, or wealth) for the market dynamic as industrial capitalism once did, surveillance capitalism audaciously lays claim to private experience for translation into fungible commodities that are rapidly swept up into the exhilarating life of the market.[1] Invented at Google and elaborated at Facebook in the online milieu of targeted advertising, surveillance capitalism embodies a new logic of accumulation. Like an invasive species with no

right down to each individual member. The competition for surveillance revenues bears down on our bodies, our automobiles, our homes, and our cities, challenging human autonomy and democratic sovereignty in a battle for power and profit as violent as any the world has seen. Surveillance capitalism cannot be imagined as something "out there" in factories and offices. Its aims and effects are *here* ... are *us*.

Just as surveillance capitalism can no longer be conflated with an individual corporation, neither should it be conflated with "technology." Digital technologies can take many forms and have many effects, depending on the social and economic logics that bring them to life. The economic orientation is the puppet master; technology is the puppet. Thus, surveillan

Why Data Sovereignty? (3/3)

- implications: self-hosting
  - services must be self-owned
  - data must be physical present

Self-hosting Requirements (1/6)

- cost-effective
  - costs privately covered
  - no monetary equivalent for service

Self-hosting Requirements (2/6)

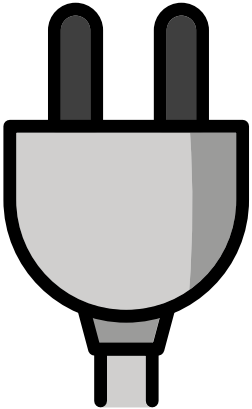

- energy-efficient
  - climate crisis is real
  - no accidental home heating

Self-hosting Requirements (3/6)

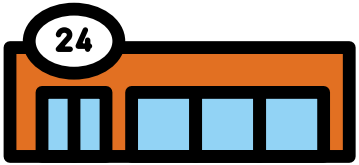

- general available
  - commodity hardware
  - spare parts should be available easily

Self-hosting Requirements (4/6)

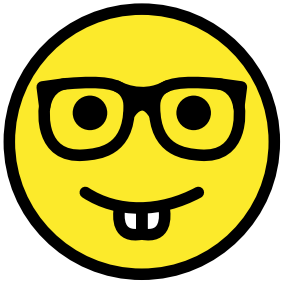

- opinionated
  - no sacrifying of cloud-native patterns
  - modern infrastructure

Self-hosting Requirements (5/6)

- performance and availability
  - sufficiently fast for medium-intensive workloads
  - best availability and partition tolerance

Self-hosting Requirements (6/6)

- non-goals
  - (infinite) scalability
  - multi-tenancy

Power Supply (1/2)

- data center usually have redundant
  - main power distribution
  - automatic transfer switch
  - diesel generator
  - uninterruptible power supplies (UPSs)
- home setup limited to UPS
  - surge protection
  - 240 W
  - integrated battery
    - lasts for 5 minutes
    - replacable

Power Supply (2/2)

- emergency power off: power outlet strip
- in cases of
  - technical malfunction (safety)
  - home raids (security)

Server Rack

- server racks
  - pricey
  - most hardware is not rack-mountable
- DIY wooden server rack
  - inexpensive
  - optimization of dimensions
  - maintenance-friendly

Networking (1/4)

- home internet involves proprietary protocols
  - VDSL2 (with Vectoring)
  - DOCSIS
- provider's router is de facto mandatory

Networking (2/4)

- unmanaged switch allows low maintenance
- most midbudget devices do not support more than 1 Gbit/s
- 16 ethernet ports allow
  - connecting of devices
  - interconnecting of different rooms

Networking (3/4)

Networking (4/4)

- WireGuard as VPN solution
  - connecting clients from outside networks
  - interconnecting private IP networks of friends

Control Plane (1/3)

- Raspberry Pi 4 Model B
  - Quadcore Cortex-A72 CPU (arm64)
  - 4 GiB LPDDR4 RAM
  - 2 USB3.0 Ports
  - Gigabit Ethernet
- 256 GB SSD via USB
- Cluster case
- extended by a 3D printed SSD mount

Control Plane (2/3)

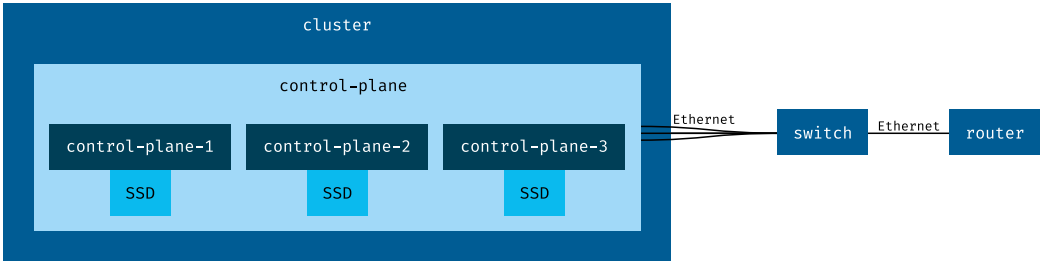

Control Plane (3/3)

- Kubernetes control plane
  - etcd
  - kube API server
  - kube controller manager
  - kube scheduler
  - CoreDNS
- managed with `kubeadm`

Storage (1/3)

- Odroid HC4
  - Quadcore Cortex-A53 CPU (arm64)
  - 4 GiB DDR4 RAM
  - 2 SATA ports
  - Gigabit Ethernet
- 512 GB SSD
- 12 TB HDD

Storage (2/3)

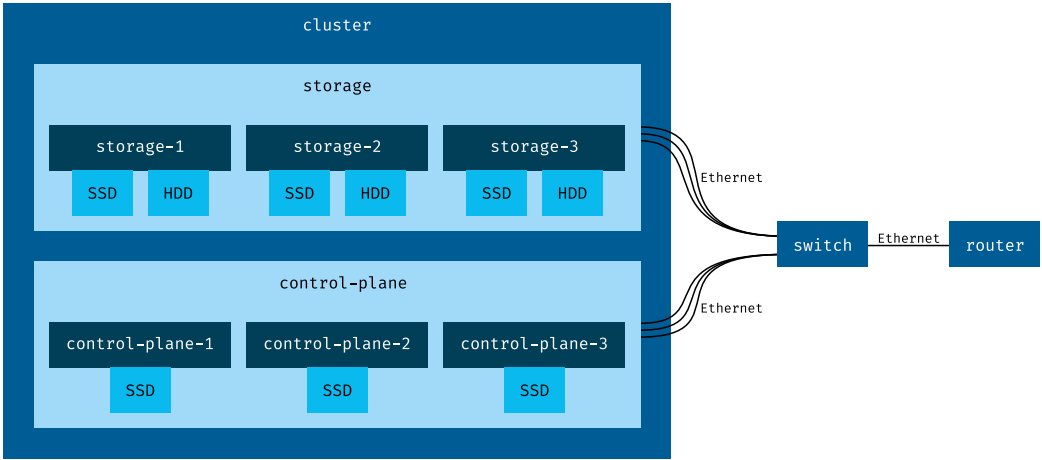
cluster
storage
storage-1
SSD
HDD
storage-2
SSD
HDD
storage-3
SSD
HDD
control-plane
control-plane-1
SSD
control-plane-2
SSD
control-plane-3
SSD
Ethernet
Ethernet
switch
Ethernet
router

Storage (3/3)

- Ceph Cluster
  - pools
    - SSD
    - HDD
  - storage types
    - block (RBD)
    - filesystem (CephFS)
    - object (RGW)
- fully managed with Rook

Compute (1/2)

- Intel NUC8I7BEH2
  - Quadcore i7-8559U CPU (amd64)
  - 16 GiB DDR4 RAM
  - M.2 port
  - Gigabit Ethernet
- 512 GB NVMe

Compute (2/2)

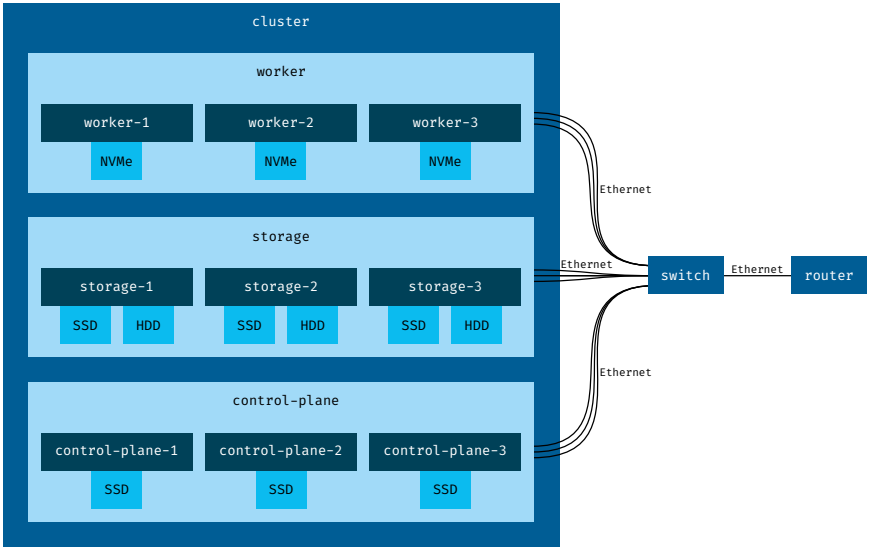
cluster
worker
worker-1
worker-2
worker-3
NVMe
NVMe
NVMe
storage
storage-1
storage-2
storage-3
SSD
HDD
SSD
HDD
SSD
HDD
control-plane
control-plane-1
control-plane-2
control-plane-3
SSD
SSD
SSD
Ethernet
Ethernet
Ethernet
Ethernet
switch
router

Maintenance Console

- network-accessable KVM (keyboard, video, mouse)
  - cheap HDMI capture card
  - USB On-The-Go (OTG) of RaspberryPI 4 Model B emulates
    - human interface devices (HIDs)
    - mass storage
  - PiKVM as linux distrubtion
    - convenient webinterface
    - low-latency interaction
    - booting of ISO images

PiKVM

Hardware Reset

- ZigBee smart plug as hardware reset
  - inexpensive
  - also enables power monitoring

Encryption at Rest

- Linux Unified Key Setup (LUKS)
  - full disk encryption
  - high security
  - good performance
- Dropbear SSH
  - boots from initramfs
  - enables remote unlock of disks

Backups (1/3)

- 3-2-1 rule
  - 3 copies
  - 2 media
  - 1 offsite
- implies encryption at rest
- proven solutions
  - borgbackup
  - restic
- Hetzner storage box as remote target
  - cost-effective
  - flexible

Borg

Backups (2/3)

```yaml
---
apiVersion: batch/v1
kind: CronJob
# ...
spec:
  schedule: "0 4 * * *"
  concurrencyPolicy: Forbid
  # ...
  containers:
    - image: docker.arik-grahl.de/containerization/restic
      workingDir: '/pvc/home-assistant-stable-webapp'
      envFrom:
        - secretRef:
            name: backup-home-assistant-stable-webapp
      env:
        - name: RESTIC_REPOSITORY
          value: sftp://u123456@u123456.your-storagebox.de:23/backups/home-assistant-stable-webapp
      volumeMounts:
        - name: data
          mountPath: /pvc/home-assistant-stable-webapp
        - name: secret
          mountPath: /root/.ssh
      # ...
```

Backups (3/3)

```yaml
---
apiVersion: batch/v1
kind: CronJob
# ...
command:
  - sh
  - -c
  - |
    #!/usr/bin/env sh -e
    restic backup \
      --host=home-assistant-stable-webapp \
      --compression=max \
      .
    restic forget \
      --host=home-assistant-stable-webapp \
      --prune \
      --keep-daily=7 \
      --keep-weekly=4 \
      --keep-monthly=12 \
      --keep-yearly=-1
```

Operation

- Ansible
  - infrastructure as code
- Helm and `helmfile`
  - most application stacks
- `kubectl`
  - unpackaged applications
  - ad-hoc workloads

Loadbalancer (1/2)

- data centers provide hardware solutions
- Keepalived as a software solution
  - IP Virtual Server (IPVS) Layer 4 loadbalancing
  - Virtual Router Redundancy Protocol (VRRP) for high-availability (HA)

Loadbalancer (2/2)

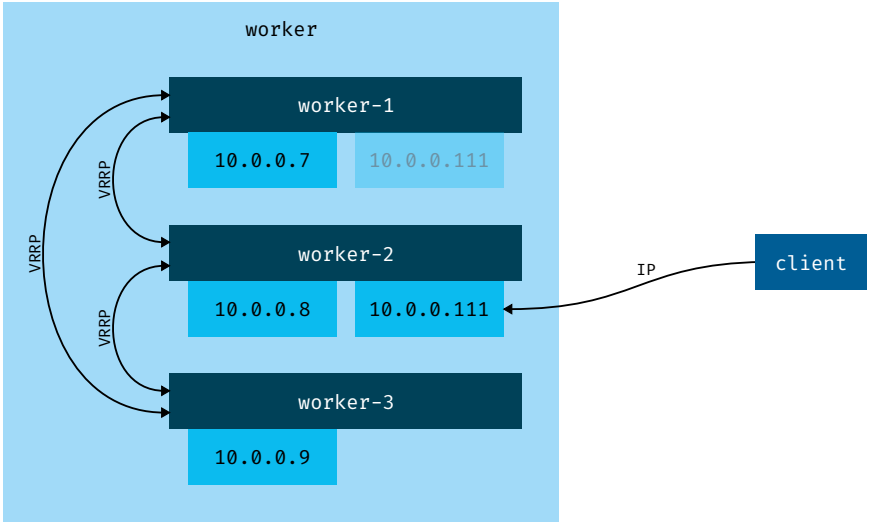

Ingress

- Ingress Nginx Controller
  - as `DaemonSet` on every worker node
  - `externalIPs` pointing to virtual IP address

Database

- central PostgreSQL cluster for all services
- CloudNative PG as a PostgreSQL operator
  - high availability (HA)
  - full lifecycle management
  - PgBouncer as connection pooler with modes
    - session
    - transaction
    - statement

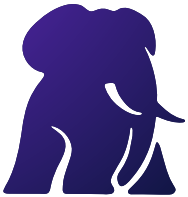

CloudNativePG

Monitoring

- Kube Prometheus Stack
  - Prometheus
  - Prometheus Node Exporter
  - Alertmanager
  - Grafana

Authentication and Authorization (1/2)

- Keycloak as Identity and Access Management (IAM)
  - OpenID Connect (OIDC) integration
  - rich integration in most services
- OAuth2 Proxy for HTTP endpoint authentication
  - generic solution for any unauthenticated services
  - authorization through defined groups

Authentication and Authorization (2/2)

```yaml
---
apiVersion: networking.k8s.io/v1
kind: Ingress
metadata:
  name: test.example.com
  annotations:
    nginx.ingress.kubernetes.io/auth-url: "https://$host/oauth2/auth?allowed_groups=tester"
    nginx.ingress.kubernetes.io/auth-signin: "https://$host/oauth2/start?rd=$escaped_request_uri"
    nginx.ingress.kubernetes.io/auth-response-headers: "x-auth-request-user, x-auth-request-email"
    nginx.ingress.kubernetes.io/proxy-buffer-size: "8k"
# ...

---
apiVersion: networking.k8s.io/v1
kind: Ingress
metadata:
  name: test.example.com-auth
# ...
- path: /oauth2
  backend:
    service:
      name: oauth2-proxy-stable
# ...
```

Development Plattform

- GitLab
  - version control system (VCS)
  - continous integration (CI) infrastructure
  - container registry

Communication (1/2)

- email
  - Postfix for relaying outgoing to upstream SMTP
  - Dovecot as IMAP server for storing email
  - Fetchmail to poll upstream inbox

POSTFIX

DOVECOT

fetchmail

Communication (2/2)

- messaging
  - Mattermost
  - Matrix

Groupware and Collaboration

- Nextcloud
  - file sharing via browser, desktop client, mobile app, WebDAV
  - calendar management via browser, CalDAV
  - contact management via browser, CardDAV
  - notes via browser, editor, mobile app
  - news reader via browser, mobile app
  - GPS tracks via browser
  - recipe management via browser, mobile app

Media and Entertainment

- Jellyfin
  - media types
    - music
    - movies
    - series
  - user interfaces
    - browser
    - AndroidTV/AppleTV
    - mobile apps
    - Kodi

Home Automation and IoT (1/2)

- wireless standards
  - ZigBee
  - ISM via Software Defined Radio (SDR)
- MQTT
  - Zigbee2MQTT
  - `rtl_433` with MQTT output
- Home Assistant
  - frontend to control, visualize, automate
    - browser
    - mobile app
  - voice assistant
    - speech to text (STT): Whisper
    - text to speech (TTS): Piper

Home Automation and IoT (2/2)

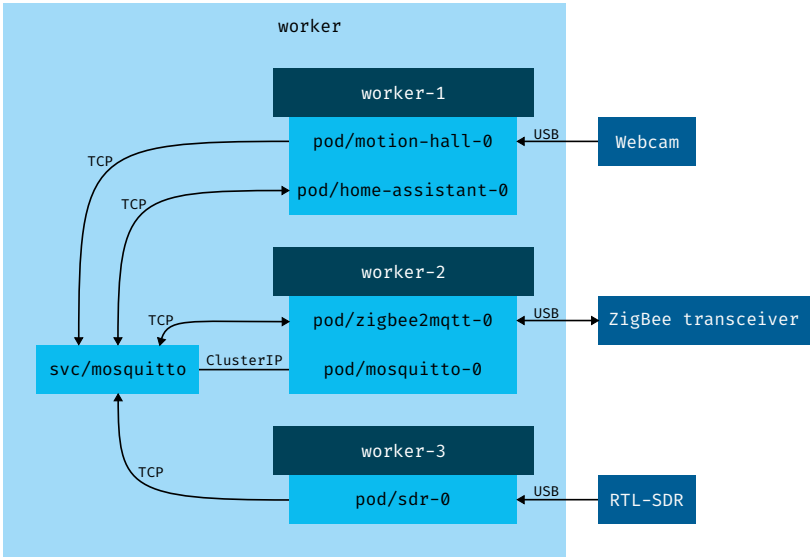

Secret Management

- Vaultwarden
  - Bitwarden API-compatible Rust-based alternative
  - user interfaces
    - browser and extension
    - mobile app
    - cli tools

Future Work

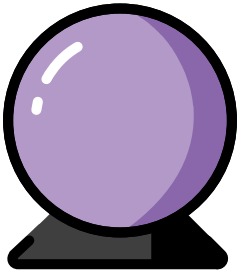

- elastic node setup
  - scale-up and scale-down of workers
  - optimization of power consumption
- modernization
  - replacement of legacy services with 12-factor apps
  - migration of shared filesystem towards object store
- automation
  - application lifecycle management

Conclusion

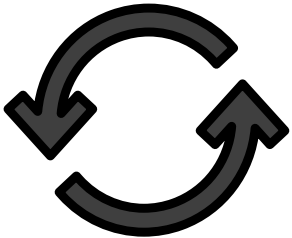

- quite some work but overall a lot of fun
- extremely maintenance-friendly setup
- great environment to learn, experiment and educate
- energy-efficient operation possible
- in home setups there is no need to sacrifice
  - storage provision
  - “HA” loadbalancing and ingress
  - rich services

HONEY, I SHRUNK THE DATACENTER

Operating Bare-Metal Kubernetes at Home for Fun and Data Sovereignty

SysEleven

Kubernetes Community Days MUNICH