PEEKING UNDER YOUR CLOUD PROVIDER'S HOOD

WHO AM I?

- Yannic Ahrens
- working since 2019 for C&H
- Mostly working on K8s stuff (K8s federations, Managed K8s)

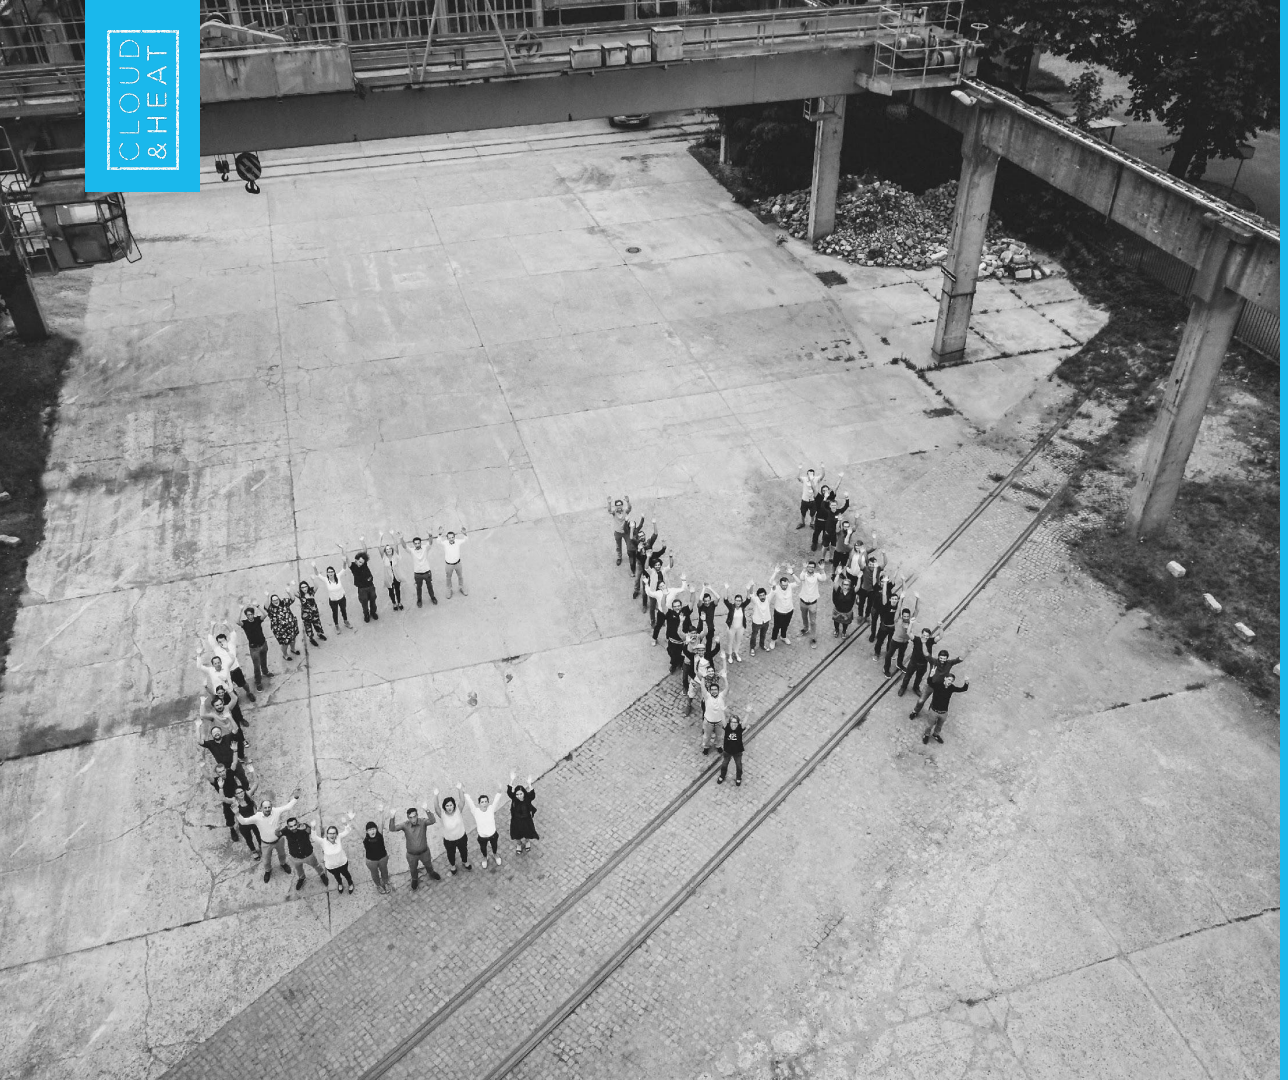

THE FUTURE OF COMPUTE

CLOUD &HEAT

Since 2011, Cloud&Heat's vision has always been to make sustainability and security drivers of digital innovation.

WHAT IS THIS TALK ABOUT?

- not an AI/machine learning talk but rather about system architecture
- sketching how the use case of a partner lead to the development of our Managed K8s platform
- giving insight into what happens on the provider side when a new cluster is deployed

BEFORE WE BEGIN

- Who of you has not heard about K8s before or does not have a rough understanding of what it does?

BEFORE WE BEGIN

- Who of you has not heard about K8s before or does not have a rough understanding of what it does?
- Who of you is using K8s or one of its derivatives?

BEFORE WE BEGIN

- Who of you has not heard about K8s before or does not have a rough understanding of what it does?
- Who of you is using K8s or one of its derivatives?
- Are you using it in development? Testing? Production?

BEFORE WE BEGIN

- Who of you has not heard about K8s before or does not have a rough understanding of what it does?
- Who of you is using K8s or one of its derivatives?
- Are you using it in development? Testing? Production?
- What are your major pains?

- developed method called Cognitive Business Robotics (CBR)
- one product is the Cognitive Secretary
- automatically scans handwritten forms and translates them to machine-readable data structures

USE CASE

PARTNERING WITH AI4BD GMBH

AI4BD's software stack (1)

- microservice architecture
- ElasticSearch, Tensorflow, ... fairly vanilla setup
- docker-based containerization
- Docker Swarm had too many constraints for use in production

USE CASE

AI4BD’s software stack (2)

- deployed their own K8s cluster on our IaaS
- AI4BD quickly realized that managing a K8s cluster is not their core business

USE CASE

MANAGED KUBERNETES

- monitoring
- life-cycle management: creation, scaling, upgrades, ...
- operations
- image registry
- ...

USE CASE

CHARACTERISTICS

- customers have high requirements on data security and privacy
- data storage and processing should happen exclusively in Germany
- dedicated servers co-located in C&H data center in Frankfurt am Main

USE CASE

MAJOR REQUIREMENTS

- protecting cluster from unauthorized, external access (obviously)
- HA (also rather obvious)
- scalable, redundant block and object storage (also shared)

USE CASE

OVERVIEW

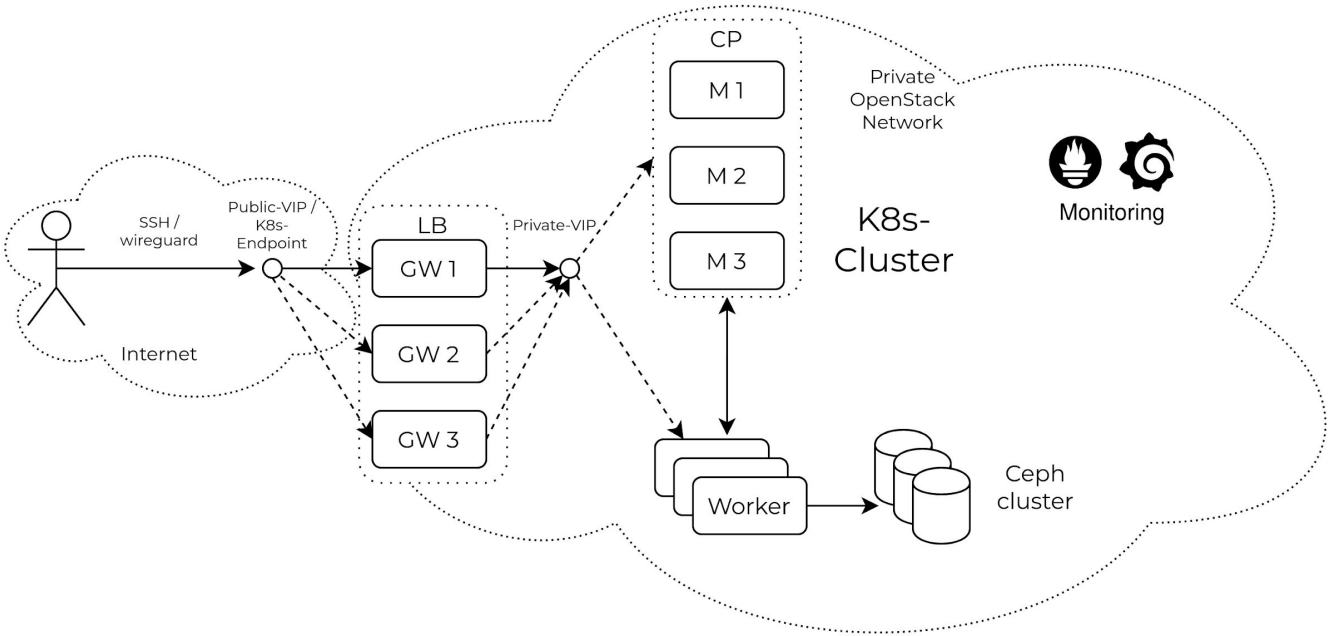

USE CASE

PROTECTION FROM UNAUTHORIZED, EXTERNAL ACCESS

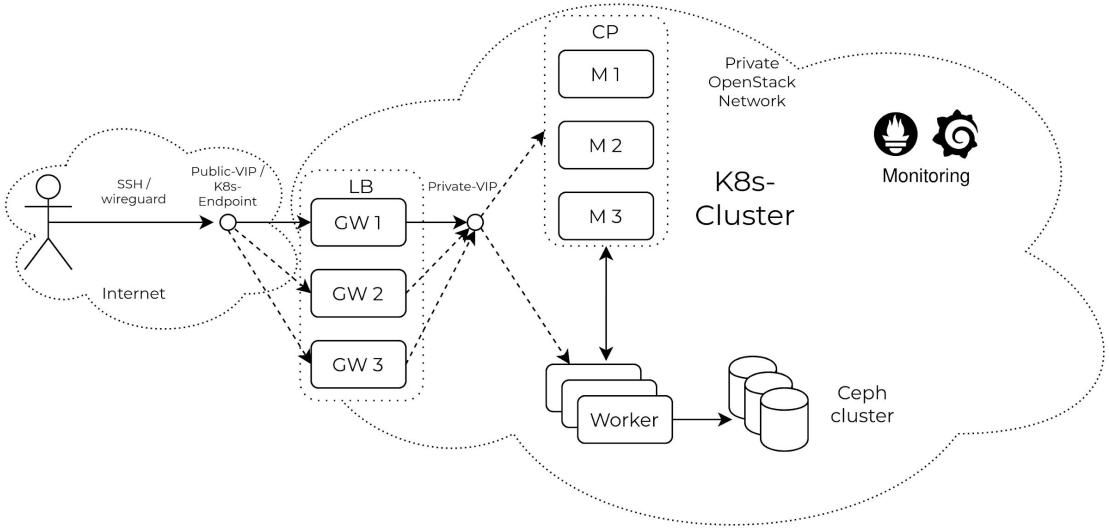

- RBAC, client certificates
- cluster resides in private network and cannot be reached directly from the outside
- wireguard VPN tunnel

USE CASE

HIGH AVAILABILITY

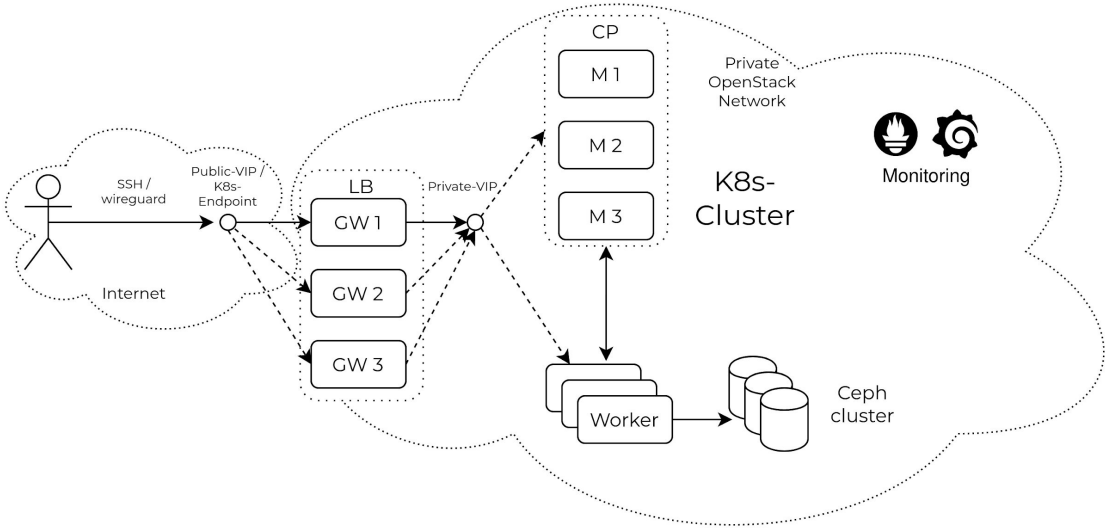

- multi-master setup
- internally: RAFT protocol (2n+1) to synchronize etcd databases
- distributed in different AZs in our DC

USE CASE

STORAGE

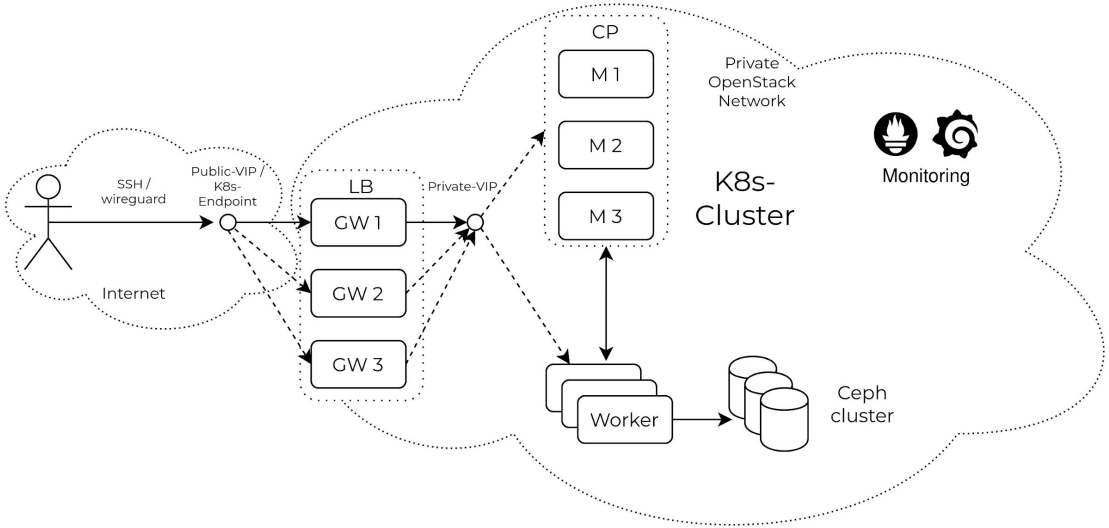

- user can request block, object storage volumes as Persistent Volume Claims with special storage class
- in the backend:
  - ceph cluster inside the K8s cluster deployed by rook
  - Cloud controller manager to interface with OpenStack environment

USE CASE

DEPLOYMENT

- Phase 1: resource provisioning with terraform (VMs, network, security groups)
- Phase 2: set up of the gateway nodes (HAProxy, keepalived, wireguard, nftables)
- Phase 3: deployment of the cluster with kubeadm
- Phase 4: smoke tests to ensure correct functionality
- orchestration of phases 2-4 with Ansible and entirely declarative

USE CASE

THE STORY SO FAR (1)

- three clusters for
  - AI4BD's internal development team (staging)
  - consultants (demo)
  - customers (production)
- successful test runs with GPUs to accelerate CBR environment

USE CASE

THE STORY SO FAR (2)

- 3,2TB NVMe storage on each host
- NVMe offers up to 64k queues with 64k entries for each queue
- allows massive parallelism, reduces latency
- ideal for microservice architectures that use distributed messaging queue

USE CASE

What’s next?

- proper Loadbalancer service type
- production-ready vGPU support

Yannic Ahrens
Cloud Architect

yannic.ahrens@cloudandheat.com

GET IN TOUCH

Contact
Zeitenströmung – Halle 15
Cloud&Heat Technologies GmbH
Königsbrücker Strasse 96
01099 Dresden
Germany

info@cloudandheat.com
+49 351 479 367 00
www.cloudandheat.com

Social Media