ungleich

Kubernetes & IPv6 in 2023

Nico Schottelius

IPv6 only HOSTING

ungleich≠

Background

- **ungleich.ch unmanaged hosting**
  - Data Center Light - unmanaged cloud / virtual machine hosting
  - IPv6OnlyHosting - same without legacy IP
- **ungleich.ch managed hosting**
  - Mainly Open Source Products
- **100% IPv6**
  - Every service is IPv6 reachable
  - Same requirement for the kubernetes workload

ungleich≠

ungleich Managed Hosting - apps

- Various applications such as
  - Matrix
  - Nextcloud (+ Collabora)
  - Django-Hosting
  - Matrix-Chat
  - Monitoring (Prometheus)
  - Netbox
  - Mastodon
  - Redmine
- And other things like…
  - routing
  - VPNs

ungleich≠

ungleich Managed Hosting - stack

- **Hosting on our own cloud infrastructure**
- **Using cdist for automation**
- **Why not stay on it? Various reasons …**
  - Developers (we hire) like containers
  - Developers (who build applications) like containers
    - Notable: ceph “requires” containers
  - Containers like containers
  - GitOps is an efficient workflow
- **In a nutshell:**
  - **Using containers + k8s seems more reasonable**

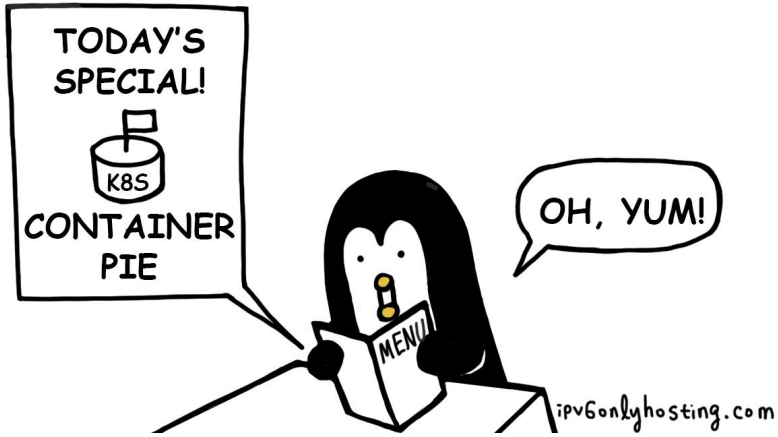

#

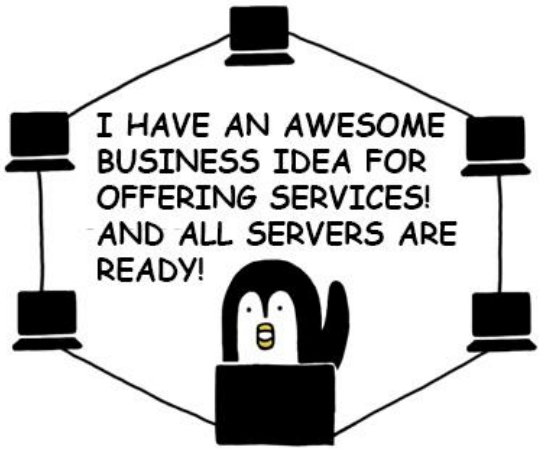

##

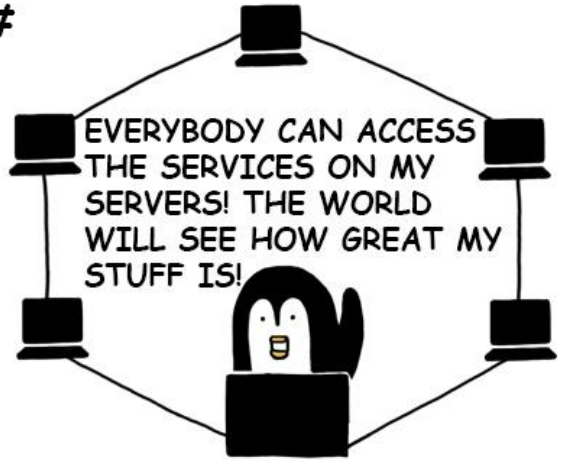

###

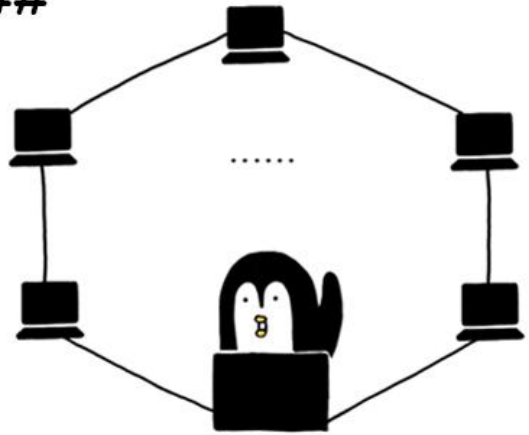

####

SHIT NO THEY CAN'T.
I'M BEHIND NAT...

THIS PENGUIN NEEDS IPV6.

ipv6onlyhosting.com

ungleich≠

Kubernetes and not IPv6

- **Initially there was no ~~light~~ IPv6 - 2015**
  - Kubernetes 1.0 was released 2015
- **Bootstrapping with IPv6 did not work for many years**
  - Some daemons would accept options, some won’t
  - The bootstrap process did not pass the options to the daemons
  - Documentation was not existing

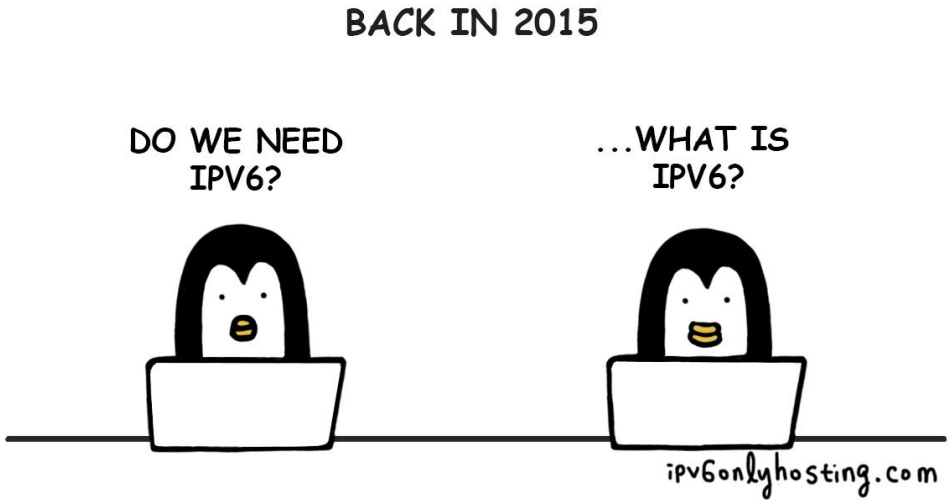

ungleich≠

Kubernetes and IPv6

- IPv6 only clusters officially supported since v1.18 - 2020
- IPv6 only clusters officially supported since v1.23 – 2021
  - https://kubernetes.io/blog/2021/12/08/dual-stack-networking-ga/
  - https://kubernetes.io/docs/concepts/services-networking/dual-stack/

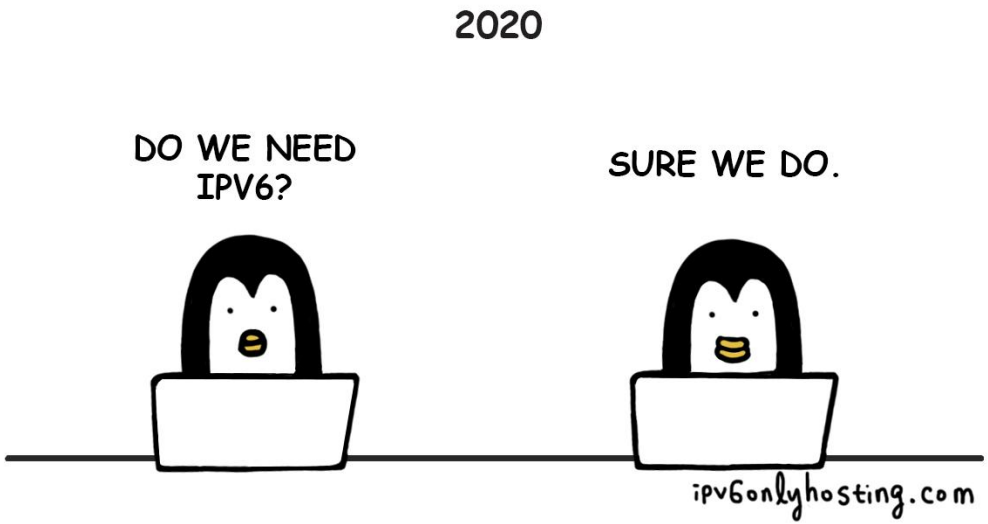

ungleich≠

IPv6 support in Kubernetes

- IPv6 as provided by the k8s stack
  - The good news: Kubernetes fully supports IPv6 nowadays
- IPv6 as supported by the k8s applications
  - Most applications work IPv6 only out of the box
  - Some applications need configurations
    - nginx - listen directive
    - Ceph - bind options

ungleich≠

Kubernetes Networking

- Kubernetes uses **CNI**
  - Container network interface
  - https://github.com/containernetworking/cni
- Anything can provide IPAM and connectivity
- If your CNI supports IPv6,

  kubernetes supports IPv6

- Some examples?

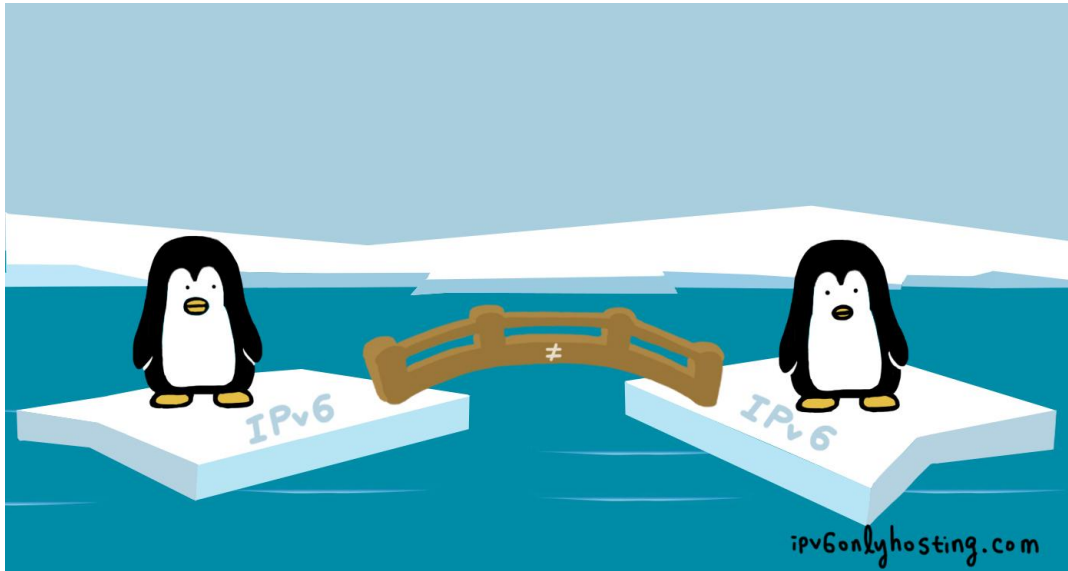

ungleich≠

Example CNI: calico

- Works in all three modes:
  - IPv6 only
  - Dual stack
  - Legacy IP only
- Has a lot of features
  - VLAN
  - IP-IP
  - VXLAN
  - Wireguard
  - BGP
- Can feel a bit overwhelming

Example CNI: cilium

- In theory supports IPv6
- Somewhat more lightweight
- Intense use of eBPF
  - “fast dataplane”
- **Does not work on IPv6 only nodes**
- But has a lot of potential in the future

ungleich

Example CNI: reference plugins

- bridge, hostdev, …
  - Very simple configurations
  - Often not designed to be used “cluster wide”
- https://www.cni.dev/plugins/current/
- static, host-local
  - Work with IPv6
- Worth checking out when starting with k8s

ungleich≠

Kubernetes non-IPv6 Concepts

- Kubernetes is designed with running **pods in a private network**
  - However if you use GUA, pods are fully public
- Services often **communication unencrypted** using http or similar
  - Services are also reachable unencrypted from the world
- Ingress is deeply embedded into k8s
  - Strictly speaking not necessary with IPv6
  - Conceptually an abstracted webserver / proxy

ungleich≠

Kubernetes pods

- Running pods with IPv6 (only) works fine
- To isolate pods from the Internet, one can use the NetworkPolicy
  - A glorified in-k8s firewall
- You can use a /64 for the pod network
  - Wait for it …

ungleich

Kubernetes services

- Kubernetes services are very practical
  - Abstraction of pods that come and go
- Services are **intended** to be **internal**
  - Nothings stops us from using them in the Internet though
- Service IPs are taken from a separate pool
  - Limited to a maximum of /108
  - /64 is not supported due to resource management bugs
  - Basically reserving space for every entry…
- More details on https://ungleich.ch/u/blog/kubernetes-without-ingress/

ungleich

Kubernetes services hack with IPv6

- Cluster domain is used by service
- Allows automatic name usage in the Internet:
  - At ungleich we use the domain “k8s.ooo”
  - Every cluster has a name, such as “c2.k8s.ooo”
  - Services are added below “svc.c2.k8s.ooo”
  - Sample service is named: “sample-service.default.svc.c2.k8s.ooo”
- **Fully automatic global DNS with Coredns**

ungleich

Kubernetes DNS

- Kubernetes has an internal DNS server
  - As a resolver
  - But also for internal service names
- Using coredns
- Supports **DNS64**
  - Needs manual change to the ConfigMap
  - Upgrade warn of features not set by default
  - But work flawlessly
- More k8s DNS info on https://ungleich.ch/u/blog/kubernetes-dns-entries-nat64/

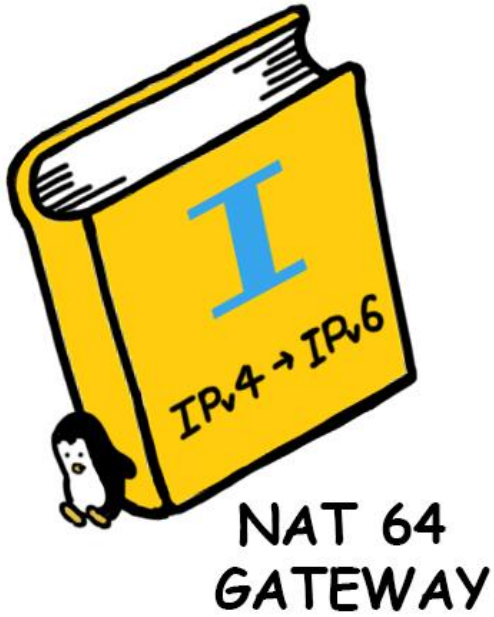

ungleich≠

Kubernetes Load Balancer with IPv6

- Load Balancer is a construct for cloud providers
- Allowing to steer inbound traffic
- In a nutshell, a load balancer …
  - has a public IP address
  - Dispatches traffic to inside the cluster
- The Service pattern with IPv6 can do the same thing

ungleich≠

Kubernetes kube-proxy

- All nodes (“hosts”) can receive traffic for any service
- kube-proxy
  - Instead of using routing, k8s uses proxying by default
  - Leads to loss of the original requestor IP address
- One can build kube-proxy less clusters with routing, but…
  - this is only supported for legacy IP clusters in calico

ungleich≠

Kubernetes works fine with IPv6.
But the concepts are from a legacy IP time.

ungleich≠

Future IPv6 support

- In theory cilium supports NAT64
  - Not tested due to the IPv6 only host bug
  - Tools such as Jool (kernel space) might not be feasible
  - eBPF certainly an interesting way to go
- Services, Load Balancer and Ingress
  - Concepts need to be questioned and revised

ungleich≠

More of this?

- Get in touch …
  - via email: ipv6@ungleich.ch
  - via matrix chat: https://IPv6.chat (**#ipv6:ungleich.ch)**
- Exchange ideas on mastodon: https://IPv6.social
- Share your IPv6 experiences on https://IPv6.blog