

On-Prem Kubernetes at Scale

with

metal-stack.io



Who We Are

- Small consulting company based in **Munich**
- Founded in 1999
- ~30 Employees
- Mainly working in **finance and insurance sector**
- Lately became hosting provider for Kubernetes as a Service (<https://metalstack.cloud>)
- CNCF Silver Member

Stefan Majer

CTO
x-cellent technologies GmbH



Gerrit Schwerthelm

Cloud Solution Architect
x-cellent technologies GmbH



A Boost for On-Prem and European Clouds?

EETimes EUROPE News Analysis Focus Opinion Education Magazine Company Dir

Opinion Trends

Cloud Repatriation on the Rise: 83% of CIOs Plan Workload Shifts in 2024

October 2, 2024 Ignacio M. Llorente

The shift toward cloud repatriation reflects a growing concern for achieving consistent security policies, optimizing cloud spending and avoiding unnecessary waste as enterprises seek more cost-effective and flexible IT infrastructure solutions.

Share this:

WIZ Platforms Solutions Pricing Resources Customers Company Sign in Get a demo

Blog

38TB of data accidentally exposed by Microsoft AI researchers

Wiz Research found...

Join Global Switch Day on February 1st 2025

By Scott Lovink, January 30, 2025 at 1:22 pm.


Download as PDF

FEBRUARY 1ST 2025 GLOBAL SWITCH DAY

FEBRUARY 1ST 2025 GLOBAL SWITCH DAY

FEBRUARY 1ST 2025 GLOBAL SWITCH DAY

FEBRUARY 1ST 2025 GLOBAL SWITCH DAY



Broadcom's VMware Pricing Hike: What It Means for Your IT Strategy and Open Alternatives

Jerome Douinco - Follow

4 min read · Dec 12, 2024

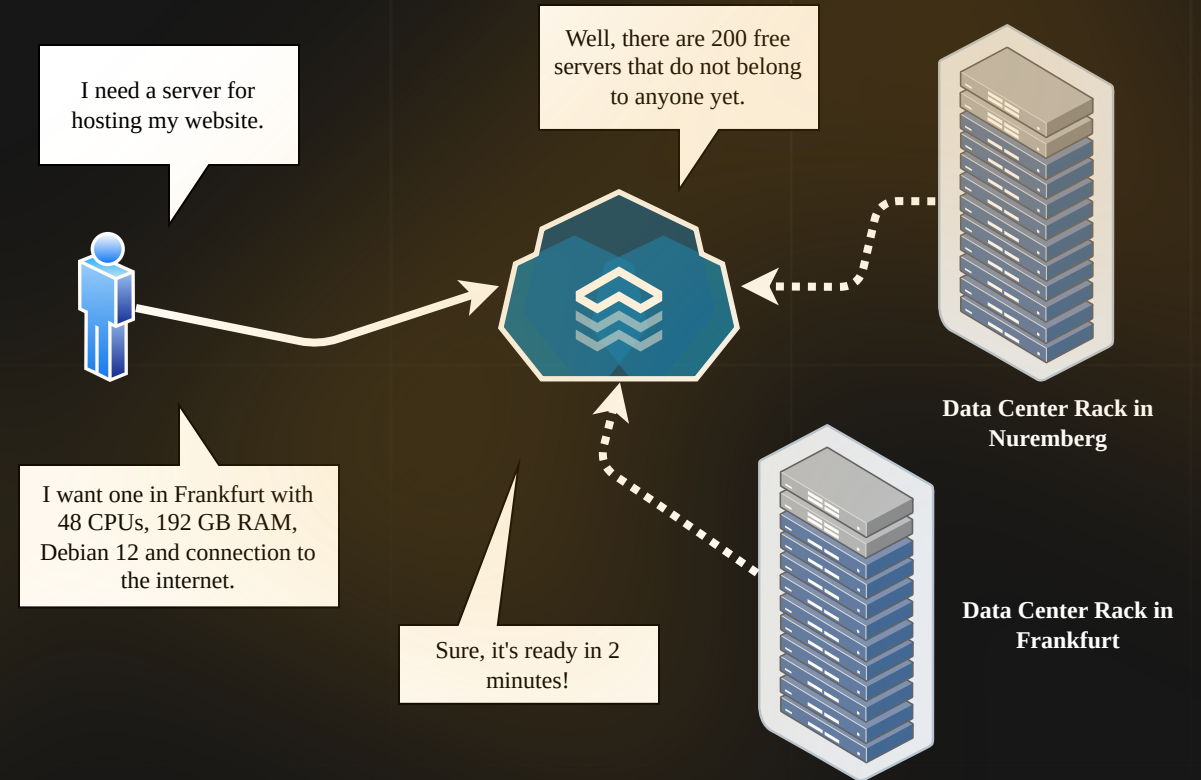
Broadcom's acquisition of VMware for \$61 billion in 2023 has marked a significant shift in VMware's licensing and pricing model. This shift, including abandoning perpetual licenses for a subscription model and



Bring the Cloud to Your Data Center

From our website:

» *metal-stack is a set of microservices that implements Metal as a Service (MaaS). It enables you to turn your hardware into elastic cloud infrastructure.* «



Why Bare Metal?

Pros

- Best **performance**
- Lowest latency
- Best possible **tenant separation**
- No noisy neighbors
- Simpler stack, no hypervisors
- Optimized for GPU Workloads
- Optimized for Virtualization

Cons

- Hardware Failures
- No overprovisioning
- No live-migration

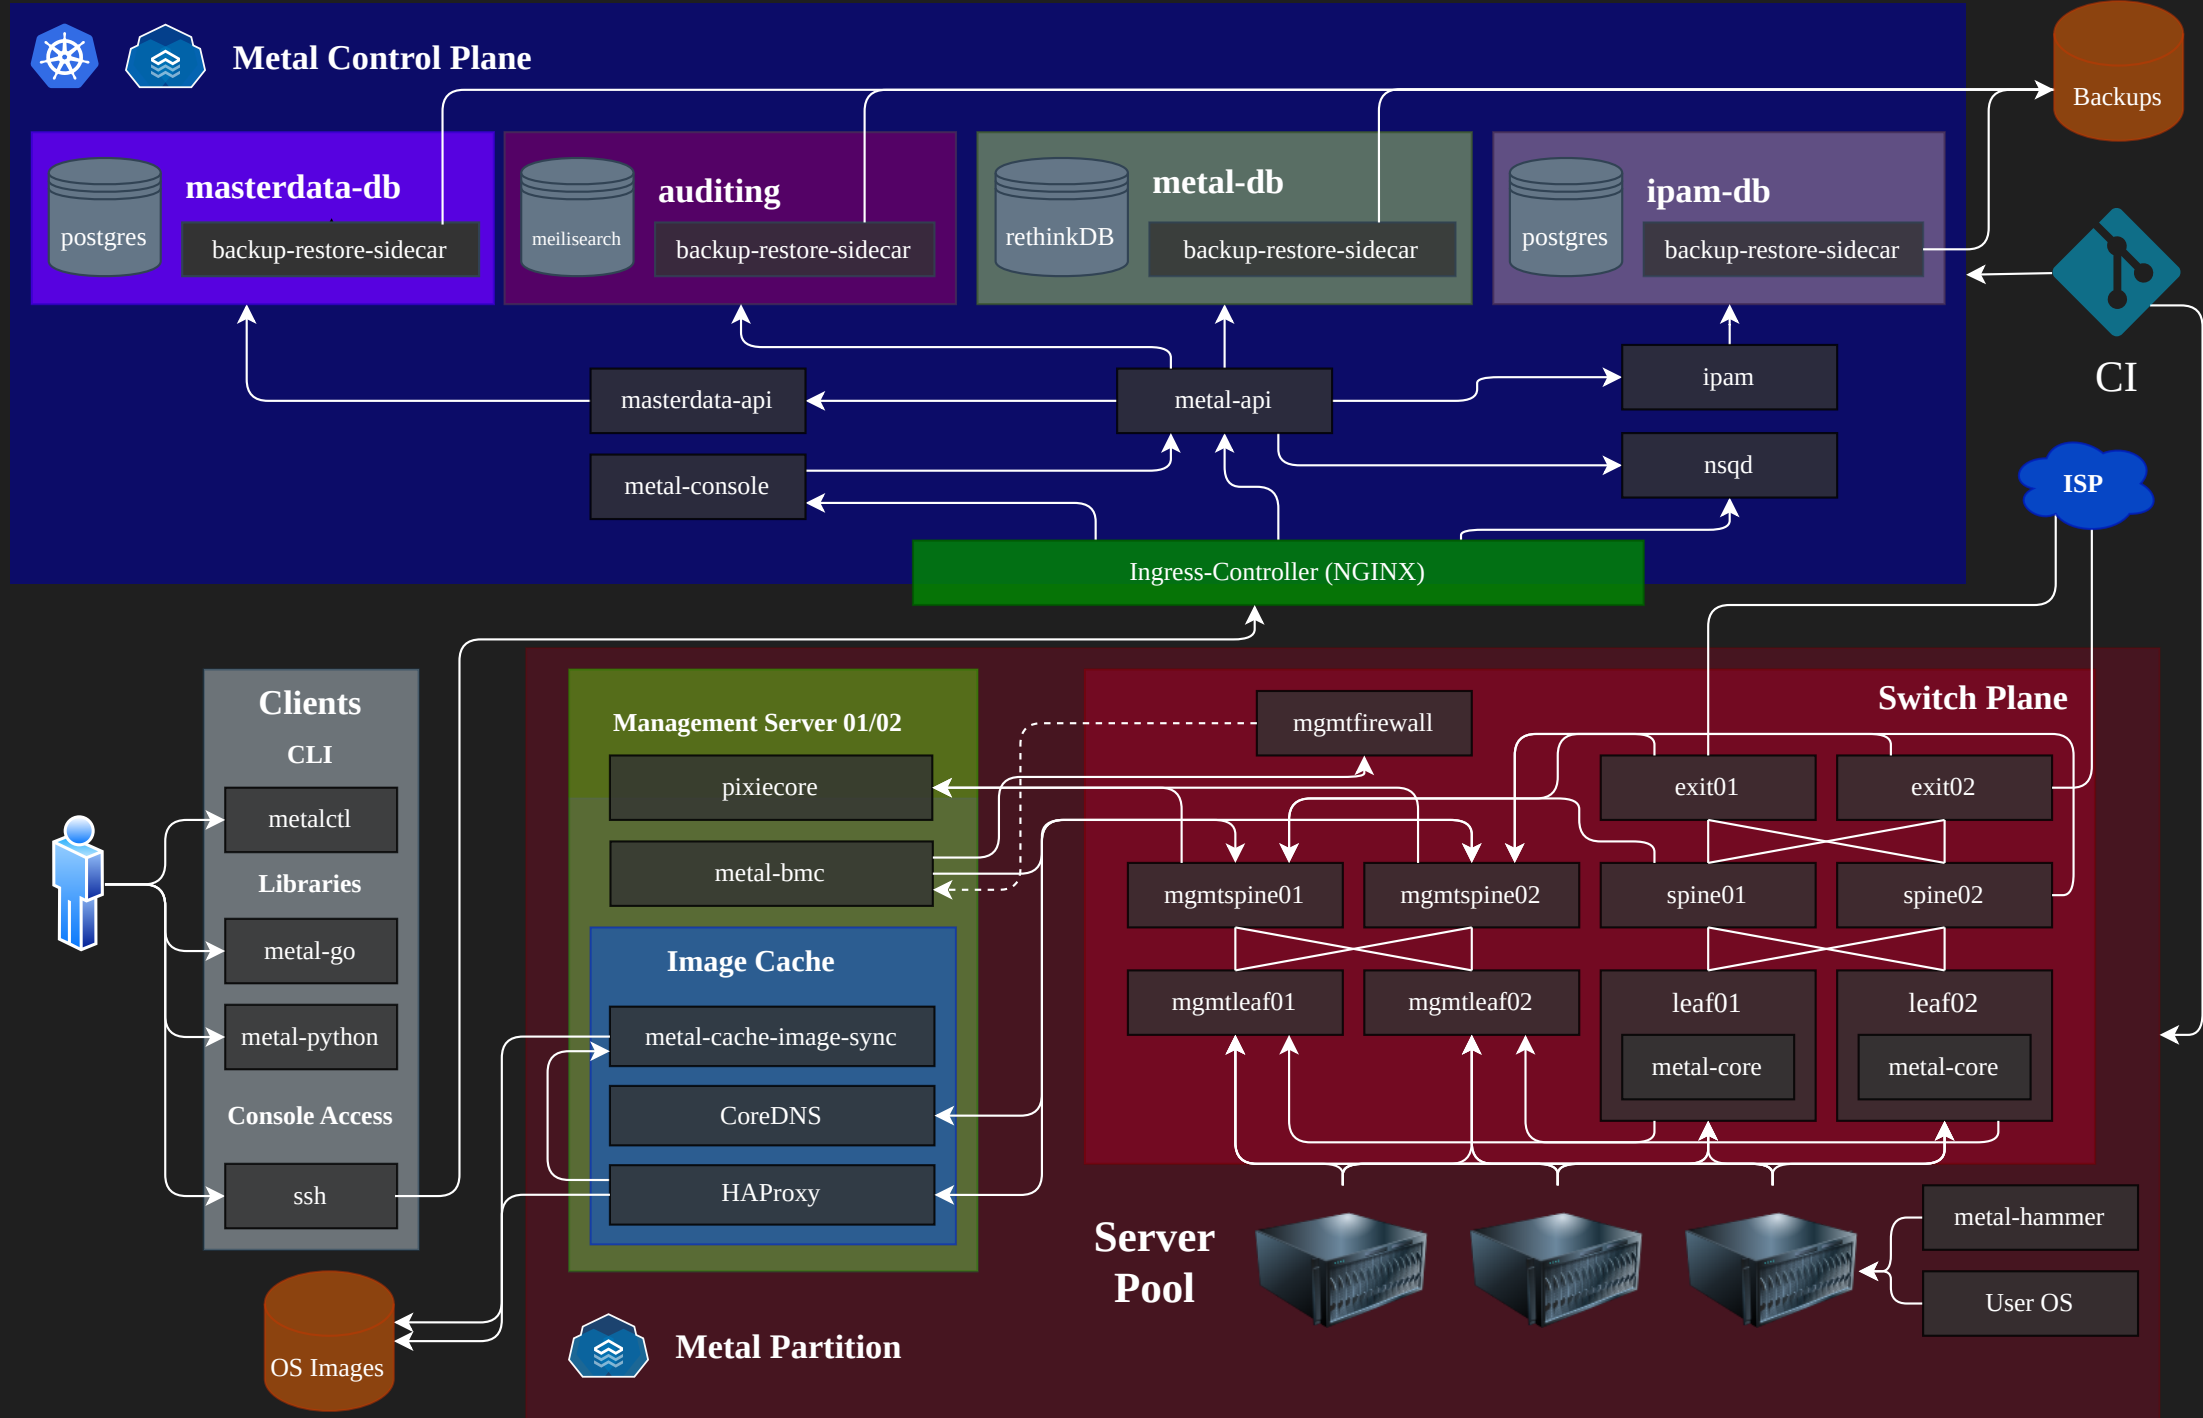
♥ **Mitigated by Kubernetes**

The Traits

- Manages **racks, machines, network, firewalls, operating system** and more
- **Fast, slim and scalable**
- No vendor lock-in
- On-Premise
- No dependencies to other cloud providers
- **Data sovereignty**
- Everyone can use it in his own data center
- Great **price** to performance ratio (the bigger the setup gets)
- Low-latency connections to existing company-internal networks
- **API-Driven** and implemented in Golang

We Use It in Production!

- With metal-stack we started this journey **in 2018** and provide a service for different critical industries
- In the current production setups we know metal-stack runs
 - **~400.000 Containers**
 - **250 Kubernetes Clusters**
 - **60 Switches**
 - **~1800 Servers**
 - **~5 PiB Persistent Storage**



Rack Design and Hardware

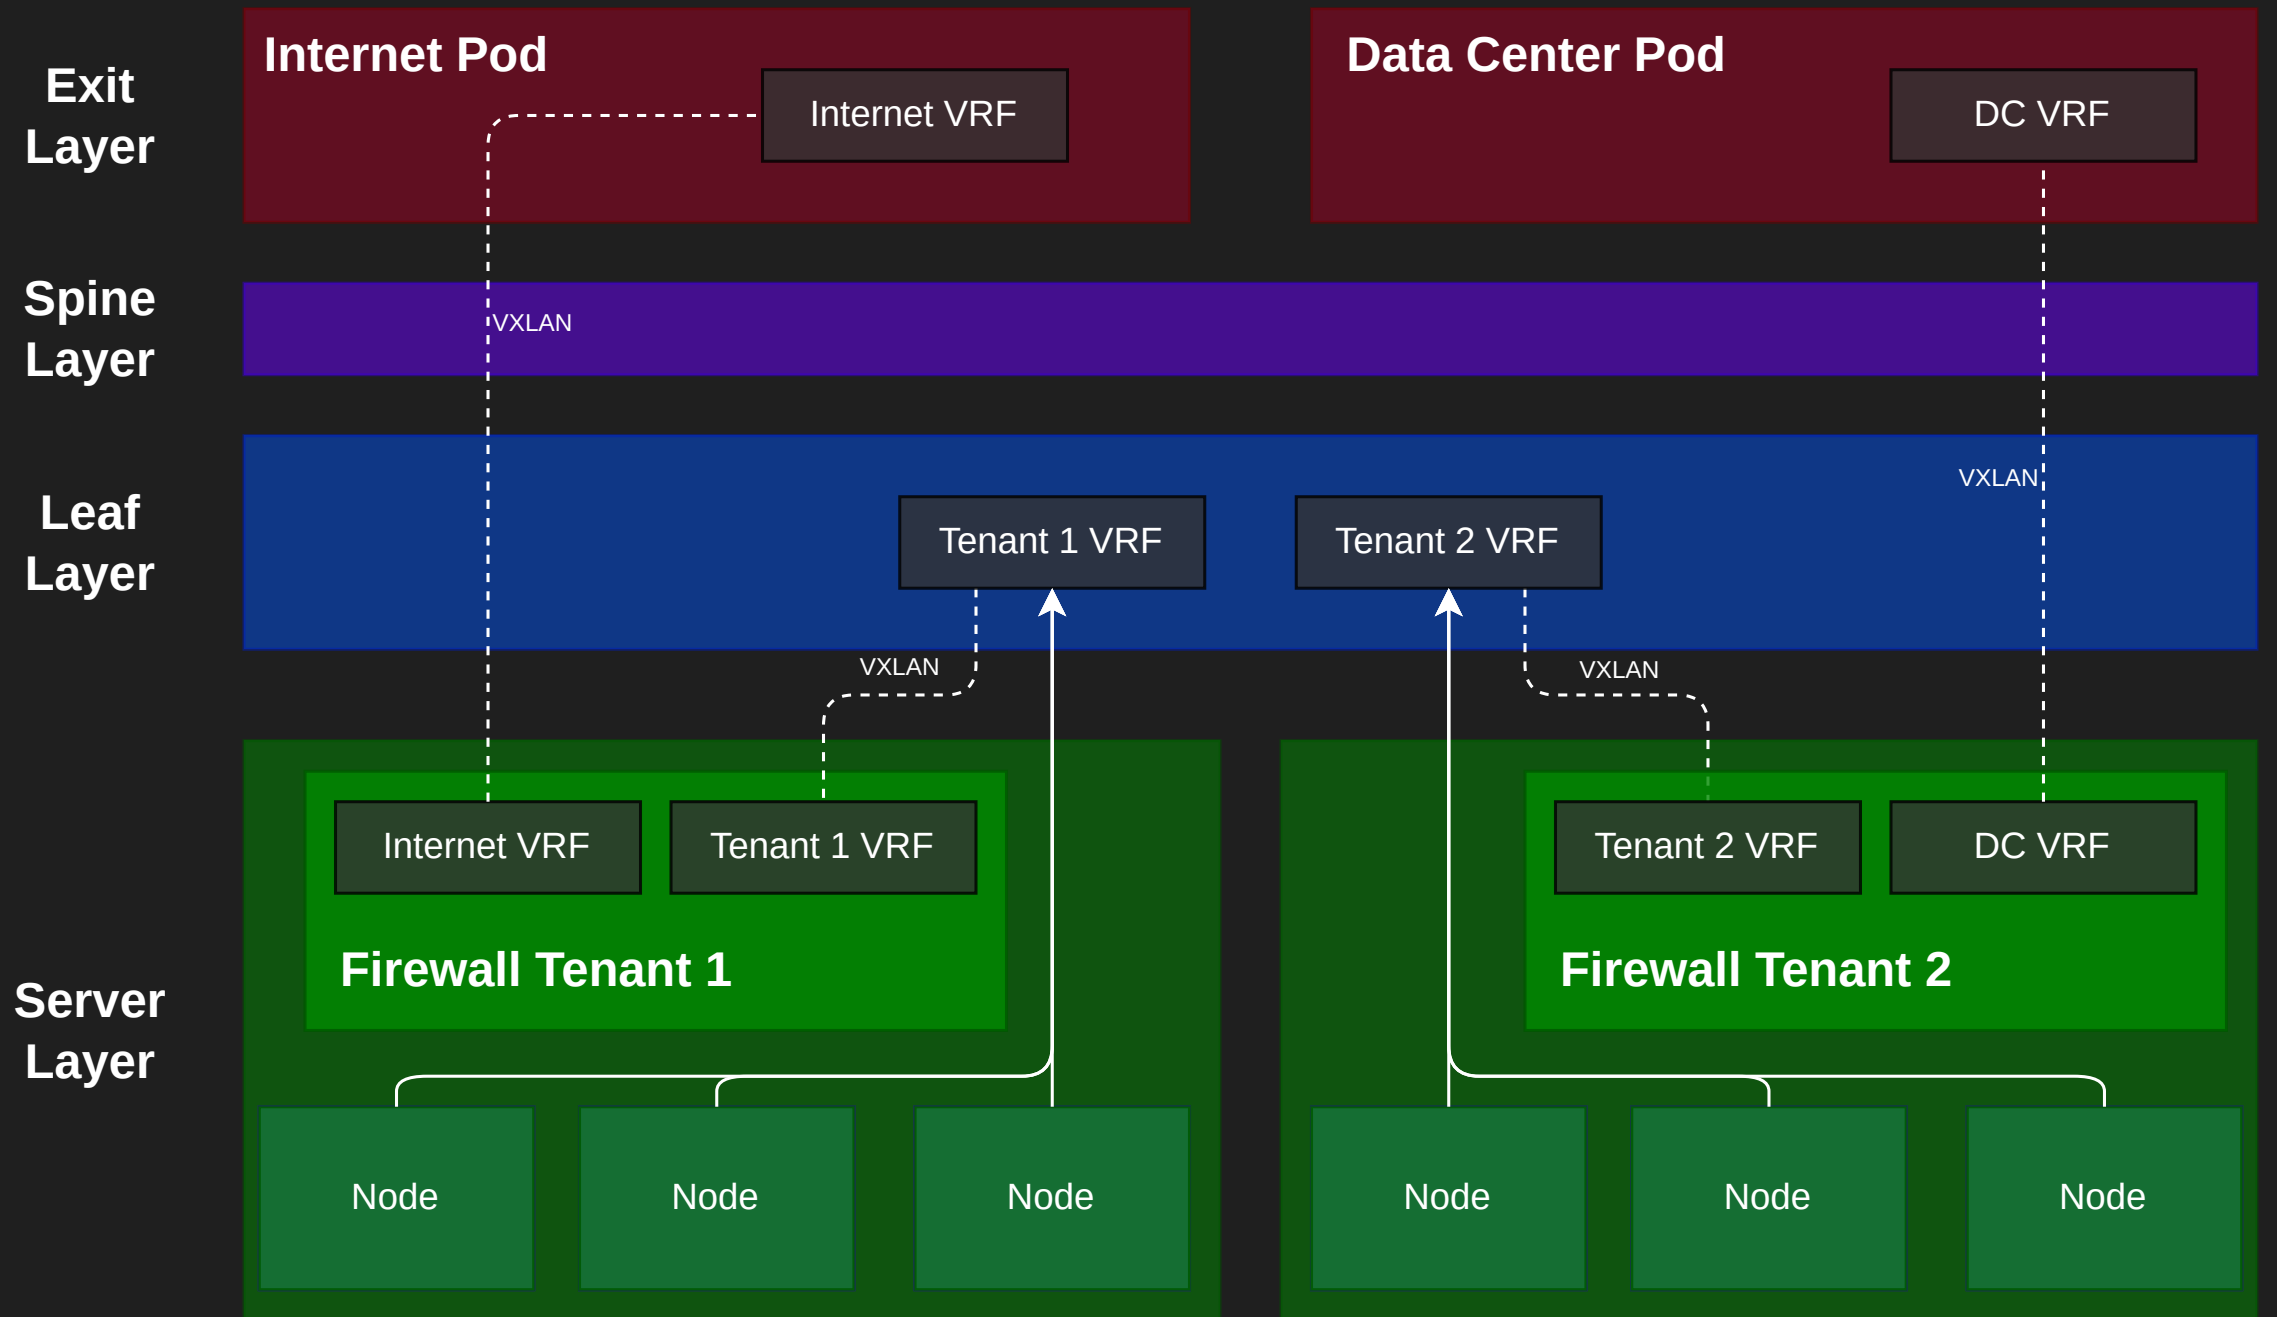
- **Scale-By-Rack-Design**
- **Failure Redundant CLOS-Topology**
- **Dual-Attached Servers**
- Overview over Supported Hardware on <https://docs.metal-stack.io>
- Abstraction Library: <https://github.com/metal-stack/go-hal>
- Automating Hardware with IPMI and Redfish is not funny 😞

Switches

- OpenCompute compatible
- SONiC OS
- 1G to 800G Port Speeds
- Managed by our Controller

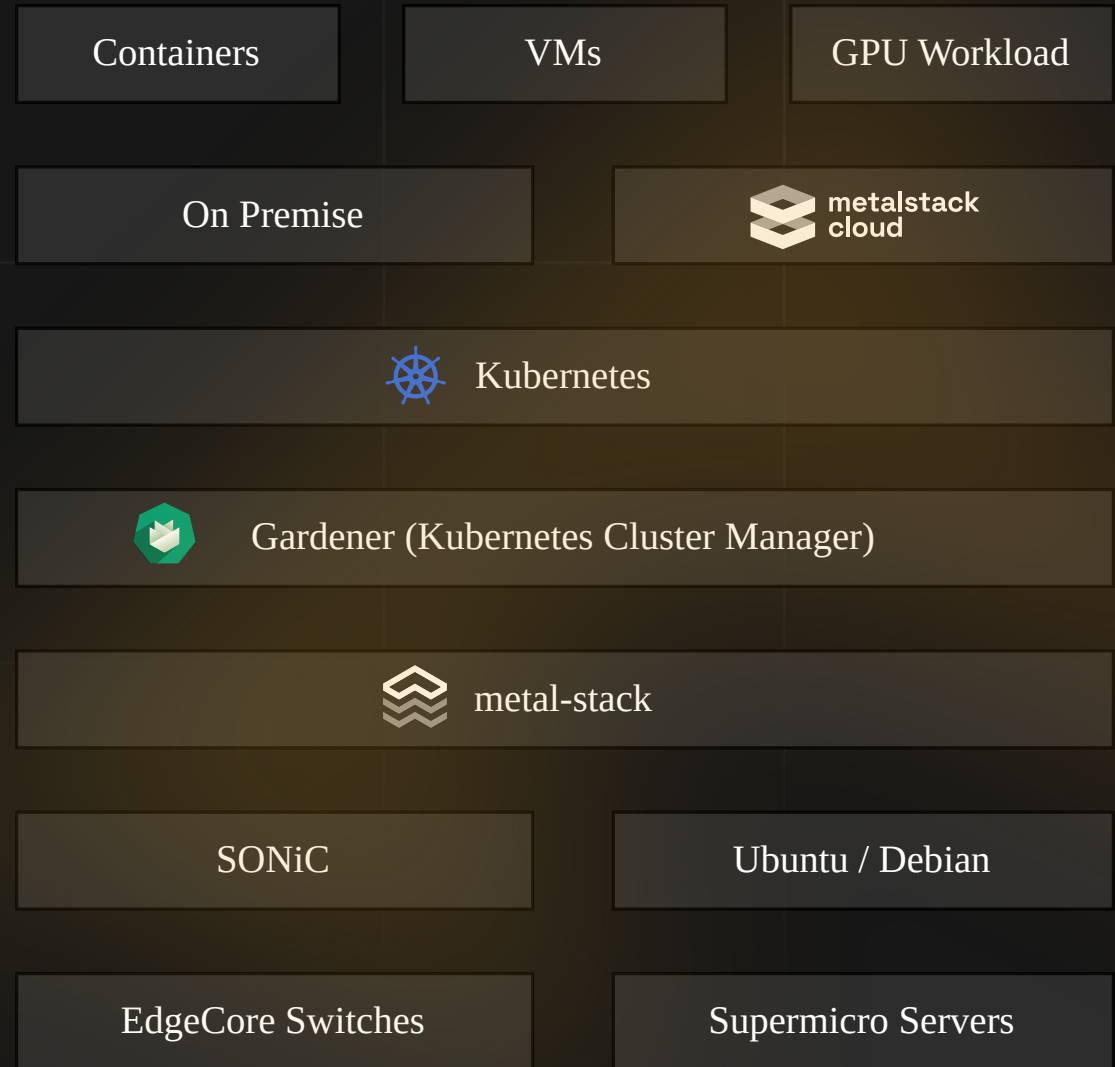
Servers

- Supermicro / OCP / Dell
- Dual Nic, initial PXE Boot
- Local NVME Disk
- Debian / Ubuntu / Your Own?
- FRRouting for BGP



Essential Building Block in the DC

- A complete Open Source Datacenter Stack
- Fast and clean Architecture, **optimized for Kubernetes**
- Developed and **maintained in Europe**, Germany
- Can run Containers, VMs via KubeVirt, and GPU Workloads



But How Can We Deploy Kubernetes?

What is Gardener?

- Open Source project by SAP (Apache-2.0)
- The website: <https://gardener.cloud/>
- Initial commit on January 2018
- Kubernetes Installer / Cluster Manager / Cluster Orchestrator
- **Multi-Cloud**: metal-stack 🎉, Open Stack, the Hyperscalers, ...
- Collaborating with metal-stack since 2019
- Regular community meetings
- Capable of managing **thousands** of Kubernetes clusters
- Enforces a clear separation between control-plane and worker nodes responsibility
- Perfect fit for building **Kubernetes as a Service**

Preconfigured Racks

It's possible to install metal-stack on your own!

But to make this process easier:

- We also build and configure racks and ship them to customer DCs
 - **Small** (8xServers, Half-Rack)
 - **Large** (200xSevers, 2 Racks)
- You only need internet connectivity, space for the racks and power

In 2022 We Made a Bold Decision...

- Setting up KaaS in the presented way is hard
- The hosted version of metal-stack = **metalstack.cloud**

We offer:

- High performance bare metal Kubernetes clusters with **ultra fast NVMe/TCP storage**
- Public user interface and **API** (gRPC and HTTP)
- Data center location in Munich
- Made and hosted in Germany (**DSGVO conformant**)
- Basically: Everything presented in this talk and more

Development Update: Recent and Upcoming Topics

- Basic **IPv6 Support**
- Installations in **offline environments**
- MEP-4: gRPC-based **V2 API with Resource Scoping**
- **Broader Hardware Support**

Try It Out

 <https://github.com/metal-stack/mini-lab>

Thanks!

📖 Docs on: <https://docs.metal-stack.io>

👤 Development on GitHub: <https://github.com/metal-stack>

📣 Community Slack Channel: <https://slack.metal-stack.io>

📦 Continuous Releases on: <https://github.com/metal-stack/releases/releases>