

kcore bare-metal cloud on Cloud Hypervisor

Real metal. Declarative VMs with vCPUs. One API for the company that wants its own cloud.

Tacconi Consulting Ltd · kcorehypervisor.com

A Cloud Hypervisor cloud — on your hardware

Other companies get a real bare-metal cloud: physical hosts, Cloud Hypervisor guests, desired state.

Not a VM sitting inside someone else's hypervisor. The metal is yours.

You declare VMs with vCPUs, memory, disks, and GPUs. The control plane applies them.

Updates are atomic. A bad generation rolls back.

Modern tools for AI on that cloud

NVIDIA GPUs

Pass GPUs into Cloud Hypervisor VMs today. NVIDIA vGPU next, so one card can be sliced across guests.

Kubernetes + GPU Operator

Declarative clusters on the bare-metal cloud. NVIDIA GPU Operator for scheduled AI workloads.

Slurm for AI

Build the metal and VM layout Slurm needs for training and batch. Controllers, partitions, shared storage as desired state.

Agent sandboxes

Isolated VMs for AI agents. MCP exposes the same API so agents can provision, run, and tear down sandboxes.

One declarative surface for humans and agents

gRPC API is the source of truth. CLI and UI call it.

MCP: advise → plan → apply. Cursor and VS Code talk to the same path.

Use it to stand up the bare-metal cloud, then the AI layers on top.

Same model for VMs, host updates, and the tooling you attach.

Open core. Paid for production support.

Community

£0

Apache-2.0. Build or take GitHub Release ISOs. Labs and evaluation.

Standard

£249 / socket / year

Signed ISOs, stable channel, tested updates. 5 tickets, 1 business day.

Premium

£599 / socket / year

Unlimited tickets, 4-hour response, remote support.

Enterprise from £999 / socket / year (min 8). Per physical CPU socket, not per core.

Your metal. Cloud Hypervisor. AI-ready.

Product: `kcorehypervisor.com`

Source and MCP: GitHub `kcorehypervisor/kcore`

Partners and migrations: `team@tacconiconsulting.com`

Non-NDA customer material for eligibility review only.