

AI - PLATFORM

Sovereign AI Platform Engineering

Five days after which your team has not just described an inference platform but built, measured and handed one over.

DURATION**5 days****GROUP****up to 12 people****FORMAT**

In-house · Open · In the lab · German or English

THIS EXISTS AFTERWARDS

- ▶ A reference architecture document across all layers: compute, network, storage, identity, registry, inference stack
- ▶ A GPU and capacity model with a traceable sizing calculation from L40S to H100 for your load assumptions, backed by your own lab measurements
- ▶ A running inference stack your team built itself: model registry, inference server, API gateway, monitoring
- ▶ An operations handover package with runbook skeleton, update path and incident protocols from the exercises
- ▶ A decision log covering the contested architecture questions of your organization with documented reasoning

AGENDA

Day 1	Your load assumptions under scrutiny · The reference architecture layer by layer · GPU capacity model from L40S to H100 · Exercise: sizing calculation for your scenario · Day result: documented capacity model
Day 2	Azure Local as the foundation · Identity thought through end to end · Network design and egress control · Exercise: implementing the foundation in the lab · Day result: approved foundation
Day 3	Registry: provenance and signatures for models · Inference servers in an honest comparison · Build: from model to API endpoint · Load test and tuning · Day result: running stack with measurement protocol
Day 4	Observability for GPU platforms · Update path: models, drivers, firmware, containers · Tenants, quotas and internal chargeback · Incident drill: the node goes down · Day result: operations dashboard and incident protocol
Day 5	The runbook becomes complete · Architecture review against your requirements · Acceptance simulation: handover to operations · Transfer plan for Monday · Final result: referenceable blueprint

**YOUR TRAINER****Dino Bordonaro**

Microsoft MVP for eight consecutive years (Microsoft Azure · Cloud and Datacenter Management), Microsoft Certified Trainer, almost 30 years in IT. Builds sovereign AI systems on own hardware (Azure Local, Foundry Local, NVIDIA GPUs) and trains on exactly that. More than 800 professionals trained, every concept proven in his own enterprise lab in a TÜV-audited datacenter.

€39,500 flat, plus VAT

The course builds competence and a blueprint, not a production environment in your data center.
Implementation and managed services are separate BORDONARO offerings and are contracted separately.

AI-PLATFORM