Hydra
Confidential Metal

Zero trust model for AI workloads

INTRO

AI infrastructure is maturing fast, and security is moving from afterthought to first principle. For regulated industries, sensitive data handlers, and government operators, the stakes are straightforward: data must remain confidential and intact, full stop.

What if you didn't need to trust the underlying provider at all? What if your workloads and sensitive data were cryptographically secured by design, with negligible performance impact?

That's Confidential Metal.

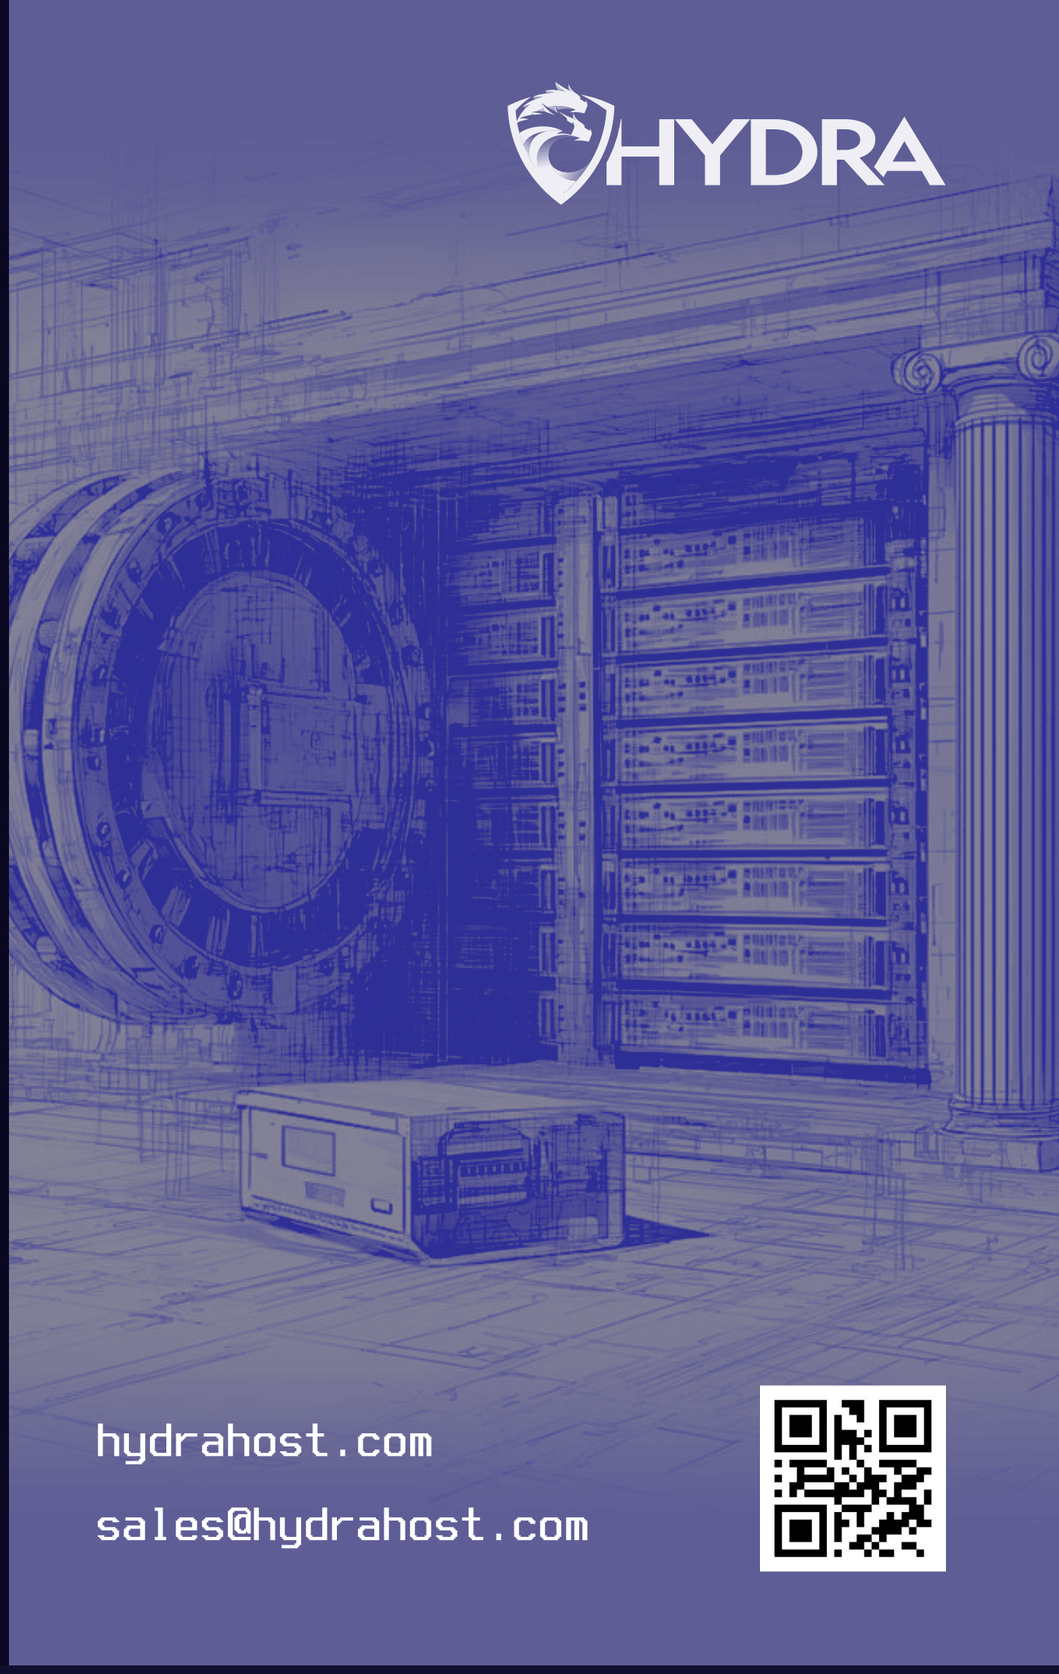

hydrahost.com

sales@hydrahost.com

DEFENSE IN DEPTH

Confidential Metal layers four independent security boundaries. Compromising any single layer leaves the others intact.

WHY HYDRA

Dedicated bare metal servers	Single-tenant hardware with full network-level isolation. Your workload shares nothing with anyone. No hypervisor overhead, no noisy neighbors, no side-channel exposure from co-located tenants. The machine is yours.
Agentless deployment, your keys only	You own the operating system. Hydra provisions the server, hands you the keys, and steps away. We maintain zero ongoing access to your environment. No management agents, no backdoors, no remote shells. Once provisioned, the server answers to you alone.
Encryption at rest	Full disk encryption is configured out of the box. Your proprietary data, model weights, training artifacts, and inference logs remain encrypted on disk at all times. Even in the event of physical seizure or forced access to the underlying storage, the data is unreadable without your keys.
Confidential Compute	This is where Confidential Metal goes beyond conventional bare metal security. Hydra preconfigures Intel Trust Domain Extensions (TDX) and Confidential Compute on compatible servers. Together, these technologies create hardware-enforced encrypted execution environments where sensitive data, including model weights, training data, and inference inputs, remains encrypted in memory and while actively being processed. Not even Hydra, the datacenter operator, or a compromised OS kernel can access what's inside the trust domain. The CPU and GPU enforce confidentiality at the silicon level. This is not a software sandbox. It is a cryptographic boundary enforced by the hardware itself.

REMOTE ATTESTATION

Confidential Metal doesn't ask you to take our word for it. Before sending sensitive data or model weights to a server, you can cryptographically verify that the execution environment is genuine and unmodified.

Hydra's attestation stack chains together three independent roots of trust. Intel TDX generates a hardware-signed quote proving the CPU trust domain was initialized with the expected firmware and configuration. Confidential Compute provides a separate GPU attestation report confirming the GPU is running in confidential mode with authentic, unmodified driver and firmware. Canonical's Ubuntu Confidential VM image supplies a measured boot chain, giving you a verifiable record of every software component loaded from UEFI through kernel to userspace.

Together, these three attestation sources give you a complete, independently verifiable proof that the hardware is real, the firmware is unmodified, the software stack is what you expected, and confidential protections are active. You verify all of this before your data ever touches the machine. If any component has been tampered with, the attestation fails and you never deploy.

This is what makes the zero trust model complete. You don't trust Hydra. You don't trust the datacenter. You verify the hardware, cryptographically, every time.

HOW CONFIDENTIAL METAL COMPARES

	Confidential Metal	Cloud Confidential VMs	Standard Bare Metal
Tenant isolation	Dedicated server	Shared host, VM-level isolation	Dedicated server
Hypervisor in TCB	No	Yes	N/A
Provider access	Zero ongoing access	Provider retains host/ hypervisor access	Varies by provider
CPU confidential compute	Intel TDX preconfigured	Available (TDX/SEV)	Typically not configured
GPU confidential compute	NVIDIA CC preconfigured	Limited availability	Typically not configured
Memory encryption	Full CPU + GPU memory	VM memory only	Not standard
Remote attestation	Direct hardware attestation (Intel, NVIDIA, Canonical)	Hypervisor-mediated	N/A
Performance overhead	Encryption only (negligible)	Virtualization + encryption	None
Bare metal performance	Yes	No	Yes
Physical seizure protection	Full disk encryption, your keys	Provider-dependent	Varies

USE CASES

Confidential Metal is purpose-built for developer and inference platforms serving:

Customers with sensitive assets

Proprietary model weights and private datasets are high-value targets. Confidential Metal ensures they remain encrypted in use, not just at rest or in transit. Platforms can offer their customers verifiable confidentiality guarantees backed by hardware attestation, not just policy documents.

Regulated industries

Healthcare, financial services, legal, and other sectors with strict data privacy and compliance requirements (HIPAA, SOC 2, FedRAMP, GDPR) need infrastructure that enforces confidentiality by design. Confidential Metal provides the technical controls that auditors and regulators increasingly expect.

Government and defense workloads

Classified and sensitive government AI workloads require hardware-level confidentiality guarantees that go well beyond standard cloud security postures. Confidential Metal delivers infrastructure compatible with the most demanding threat models in the public sector.

The rightfully paranoid

If your threat model assumes the infrastructure provider is compromised, you're not paranoid. You're realistic. Confidential Metal is built for exactly that assumption.