Your HPC AI Strike Team On Demand ...

The engineers behind the compute delivery.

CloudSyntrix is the AI engineering and deployment arm for the people building AI at scale: neoclouds, data center operators and the HPC ecosystem behind them. Your hardware arrives. We turn it into production capacity that is racked, cabled, liquid-cooled and validated.



WHERE THE TIME GOES

The hard part starts after the GPUs are bought

Power, cooling, fabric, validation and handoff are where deployments stall. CloudSyntrix puts engineering Strike Teams on site to take a facility or a cluster from empty white space, or from a retrofit, to production.

WHAT WE DO

Three layers, one accountable team

Physical infrastructure

- Rack-and-stack and high-density integration, including NVL72-class racks
- Direct-to-chip liquid cooling and CDU coordination
- Structured fiber, AOC and DAC cabling with labelling, documentation and validation testing

GPU and fabric

- InfiniBand and Ethernet (Spectrum-X, RoCE) fabric builds
- ConnectX adapters and BlueField DPUs
- Storage integration across DDN, VAST, Pure, NetApp and WEKA
- Configuration and acceptance testing

Software layer

- Kubernetes with the GPU Operator, Slurm, NVIDIA Base Command Manager and Run:ai
- NVIDIA AI Enterprise and NIM microservices for serving inference

Offered capabilities, drawn from CloudSyntrix service materials. They are not claims about completed deployments, and no performance figures are made here.

HOW WE WORK WITH YOU

Two ways to take part

Led by CloudSyntrix

A Strike Team assembled by CloudSyntrix, project management included, running the delivery from site readiness to handoff.

Inside your program

CloudSyntrix sits behind you as your delivery team, white-label where agreed, so your customer relationship stays yours.

We scale up for a build, then hand over a validated environment with its documentation. Terms are agreed per engagement.

COMMON QUESTIONS

Where we fit

We already have an in-house team.

Good. We are built to plug in. We add surge capacity for a build so your people stay on operations.

NVIDIA or our OEM handles the deployment.

They handle the system. We handle what is around it: the facility, fabric, cooling, validation, and the path to production.

WHO WE HELP

Start from where you are

Neoclouds	*Your GPUs are committed. What does the path from dock to production look like?*
Data center operators and colos	*Are your halls ready for liquid cooling and high-density racks, or do tenants ask for more than the building was designed for?*
OEMs, ODMs, NVIDIA partners and storage vendors	*When your deal closes, who handles the deployment?*
Existing clients	*What is on your 2027 roadmap: a refresh, more inference capacity, or additional sites?*

Where are you most constrained right now: getting capacity live, retrofitting existing halls, or onboarding customers onto it?

Tell us, and we will map the work with you. Start at **cloudsyntrix.com**.