

Power, Run, and Manage Your AI Factory.

UPC Vektor orchestrates GPU infrastructure, fabric, workloads, tenants, and token consumption through a unified operational layer — **full-stack AI operations, from silicon to token.**

Orchestrate

Optimize

Govern

Monetize

WITHOUT A CONTROL PLANE

- GPUs idle between jobs or stranded by poor scheduling
- Training jobs starve latency-sensitive inference
- Operations spread across disconnected tools
- Costs visible in total, impossible to attribute

WITH UPC VEKTOR

- Idle and stranded capacity surfaced and reclaimed
- Placement by priority, SLA, policy, latency, and cost
- One operational view across every layer
- Cost-to-serve per tenant, workload, model, and token

94 CAPABILITIES ACROSS 15 CATEGORIES



Optimize your GPU fleet

Provision and monitor clusters, map rack and network topology, track thermal and power utilization, manage pool capacity.



Optimize token economics

Track token throughput, latency, and cost by model, tenant, and workload; compare private inference against public APIs.



Multi-tenant AI service layer

Tenant onboarding, GPU allocation, API access, RBAC, and usage metering from one control plane.



LLMs across runtimes

Abstract inference across vLLM, LiteLLM, TensorRT-LLM, and public APIs; port between private and public environments.



Govern with compliance

RBAC, SSO, MFA, per-tenant policies, centralized audit logs, and hardware-backed confidential computing.



Manage through ChatOps

Monitor fleet health, identify idle GPUs, and execute runbooks, deployments, and rollbacks with human approval.

THE CAPABILITY CATALOGUE

Platform Experience 6

Discovery & Inventory 6

Data Center Management 7

Cloud & Containers 5

Infrastructure Observability 7

AIOps 7

Automation & Configuration 3

FinOps & Sustainability 8

AI & GPU Platform 12

AI Observability 7

Tenant & Workload 7

Model Management 3

Governance & Security 6

Resilience & Privileged Security 4

Administration & Integration 6

Ready to run your AI factory with intelligence?

Book a walkthrough of UPC Vektor, or ask us to assess your current GPU estate.

UPC VEKTOR[™]
UNITEDLAYER.COM



EMAIL

contact@unitedlayer.com



PHONE

+1-888-853-7733



X

x.com/unitedlayer



LINKEDIN

linkedin.com/company/unitedlayer

Capability descriptions reflect UnitedLayer published materials for UPC Vektor. The 94-capability catalogue reflects the current consolidated feature set; framework interoperability, deployment models, and feature availability may vary by release, plan, and configuration. UnitedLayer®, United Private Cloud®, UnitedSecure™, UnitedConnect®, UnitedEdge®, and UPC Vektor are marks of UnitedLayer.

© 2026 UNITEDLAYER. ALL RIGHTS RESERVED.