

Your Gateway to Frontier AI at Scale

Train, deploy, and scale AI workloads with accelerator choice, sovereign controls, global scale, and world-class performance in one unified platform.

Top 500 HPC

#20 Globally

Core42 Compass API

50+ Models

Storage Worldwide

#3 IO500

GPUs Deployed/Planned

86,000+

Six Foundational Advantages

A full-stack AI platform that removes the barriers between AI ambition and production outcomes.

01 Accelerator Freedom

- Choose the right silicon for your workload
- NVIDIA | AMD | Qualcomm | Microsoft | Cerebras

No lock-in. Maximum flexibility.

02 AI-Optimized Storage

- High-performance storage
- Fast, reliable access to AI and HPC data

Faster training. Reliable access.

03 Instant Global Scale

- Thousands of GPUs on demand
- From prototype to enterprise-grade production

Scale when you need it, where you need it.

04 Built for Trust

- Sovereign infrastructure
- Regulatory-ready environments

AI you can deploy with confidence.

05 Core42 Compass Inference

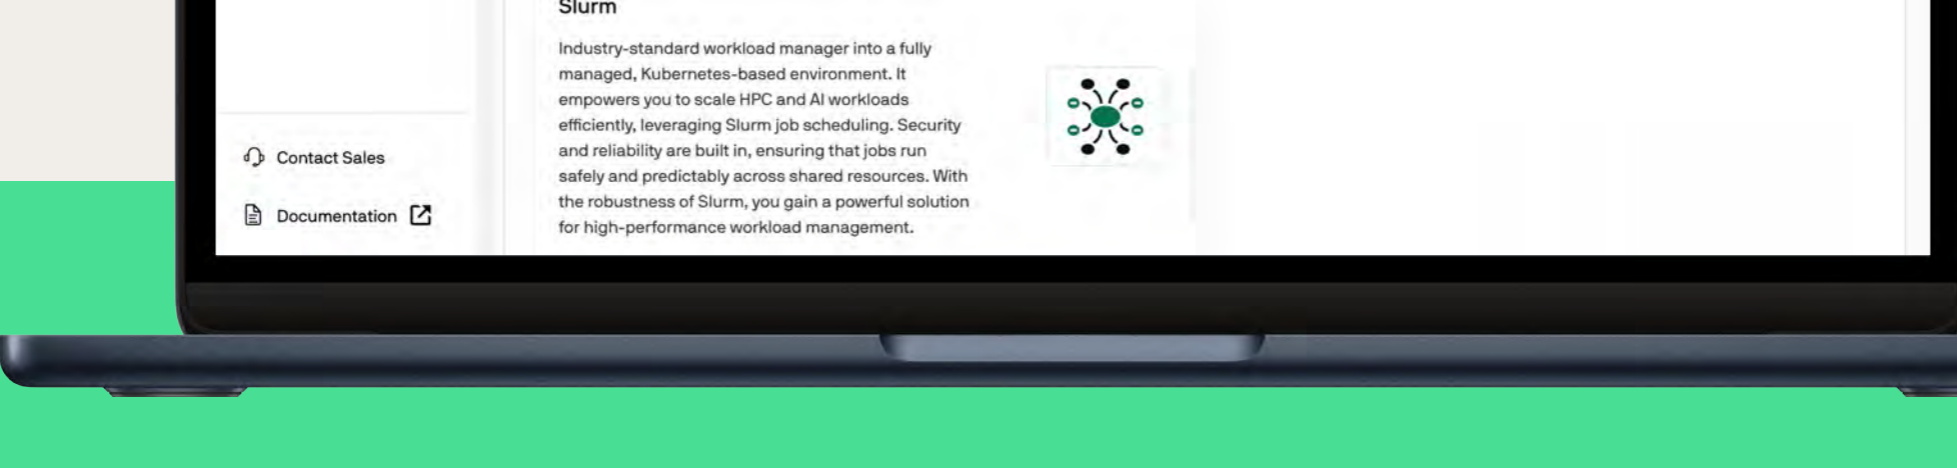
- Unified access to leading open & closed models
- Manage inference at scale

From GenAI idea to real-world deployment.

06 Fully Managed AI Platform

- Managed GPU infrastructure
- Built-in orchestration (Kubernetes, Slurm)
- Integrated observability & monitoring
- Developer-ready APIs

Remove DevOps friction across the entire ML lifecycle.



World-Class Performance Independently Verified

Top500 & IO500 Recognized Supercomputing

#20
Worldwide
TOP 500 HPC LIST
Maximis-01, AMD MI300X

#1
In the UAE
TOP 500 HPC LIST
NVIDIA DGX System (#37 global)

#2
In the UAE
TOP 500 HPC LIST
AMD MI210 System (#38 global)

#3
Worldwide
IO500 STORAGE
Sustained AI Throughput

Core42 Compass: From Pilot to Production

Fully managed AI inference platform. 50+ models. One unified API.

No rewrites. No lock-in. Sovereign controls.

Batch Processing
Cost-optimized inference for analytics, summarization, and high-volume workloads.

Fine-tuning Services
Customizing models on domain data. White-glove support from data prep to deployment.

Compass Playground
Test prompts, compare models, and validate use cases before production.

Agentic Frameworks
Multi-step workflows, tool-calling, and dynamic task execution at scale.

Private Endpoints
Access controls, guardrails, and enterprise-grade authentication built in.

Cost Visibility
Real-time metering and budget management per team, project, or workload.

Uncompromising Sovereignty
Run GenAI with in-country data residency, secure integration, and enterprise-grade governance.

Core42 AI Cloud

Powering the Next Wave of AI

Deploy AI workloads globally across sovereign data centers with built-in scale, security, and accelerator choice. From large scale training to inference in one platform.

Get Started