

TECHNICAL SPECIFICATIONS

HARCH INTELLIGENCE

Infrastructure Datasheet

Africa's largest AI hyperscale data center network. 1,798 GPUs across 5 hubs, PUE <1.15, 8ms latency to Europe, 89% below industry carbon intensity. Powered by dedicated renewable energy at \$0.018/kWh.

1,798 GPUs | 5 Hubs | PUE <1.15 | Tier IV+

harchcorp.com/intelligence

Harch Intelligence Platform Overview

Harch Intelligence operates Africa's largest AI hyperscale data center network, designed from the ground up to deliver sovereign AI compute with industry-leading efficiency and sustainability. The platform comprises five purpose-built GPU compute hubs strategically located across Morocco, each co-located with dedicated renewable energy installations operated by Harch Energy. This vertical integration enables us to deliver compute at an energy cost of \$0.018/kWh, which is 72% below European averages, while maintaining a Power Usage Effectiveness (PUE) below 1.15 and an average carbon intensity of approximately 47 gCO₂/kWh, which is 89% below the global data center industry average.

1,798 GPUs Operational	<1.15 PUE	8ms Latency to EU	81.5% Renewable Energy
----------------------------------	------------------------	-----------------------------	----------------------------------

GPU Hub Specifications

Each hub is purpose-built for AI training and inference workloads, featuring H100 and H200 GPU clusters, hybrid liquid-air cooling systems, Tier IV+ security, and direct connections to submarine cable systems. Hubs are located at sites selected for optimal renewable energy availability, climate conditions favoring free cooling, and proximity to international network infrastructure. The five-hub architecture provides geographic redundancy and enables our proprietary carbon-aware scheduling system to dynamically route workloads to the cleanest available compute at any given moment.

Hub	GPUs	Renewable	Carbon (gCO ₂ /kWh)	Cooling	Status
Ouarzazate	800	97.2%	18	Hybrid Liquid+Air	Phase 1
Dakhla	400	94.8%	32	Hybrid Liquid+Air	Phase 1
Benguerir	350	88.5%	55	Hybrid Liquid+Air	Phase 1
Tanger	200	82.1%	95	Hybrid Liquid+Air	Phase 1
Casablanca	48	45%	210	Air (Free Cooling)	Operational

Table 1: Harch Intelligence GPU hub specifications and operational status

Power and Cooling Infrastructure

The power infrastructure supporting Harch Intelligence is a key differentiator. Unlike conventional data centers that purchase energy from the grid, Harch Intelligence is powered exclusively by dedicated renewable energy installations operated by Harch Energy. This eliminates exposure to grid carbon intensity, volatile energy prices, and supply interruptions. The total power pipeline extends to 500MW at full build-out, with 200MWh of battery backup providing ride-through capability for any renewable intermittency. The hybrid liquid-air cooling system leverages Morocco's favorable climate, where average temperatures of 22 degrees Celsius enable approximately 8,500 hours per year of free cooling, reducing cooling energy consumption by up to 40% compared to traditional mechanical refrigeration systems.

Parameter	Specification	Industry Comparison
Total Power Pipeline	500MW (Phase 3)	Equinix: 280MW total
Phase 1 Commissioning	100MW, Q3 2027	-
Energy Cost	\$0.018/kWh	EU avg: \$0.065/kWh (72% lower)
PUE	<1.15	Industry avg: 1.56 (26% better)
Free Cooling Hours	~8,500 hrs/yr	N. Europe: ~5,000 hrs/yr
Avg Temperature	22 degrees C	Frankfurt: 10 degrees C (heating needed)
Solar Irradiance	2,800 kWh/m2/yr	S. Europe: 2,000 kWh/m2/yr
Backup Power	200MWh battery	Industry standard: diesel generators

Table 2: Power and cooling specifications with industry comparisons

Network and Connectivity

Harch Intelligence's network infrastructure is designed to serve both African and European markets with ultra-low latency. Morocco's position at the northwestern tip of Africa, connected to Europe by 14 submarine cable systems with a combined capacity exceeding 80 Tbps, provides an unmatched connectivity advantage. Our Dakhla hub is just 8 milliseconds from European internet exchanges via the South Atlantic Express cable, making it a natural extension point for European financial services, gaming, and real-time AI inference workloads. Dual diverse path routing ensures that no single cable failure can interrupt service, and direct peering with major European IXPs eliminates unnecessary hops and latency.

Connectivity Feature	Specification
Submarine Cable Systems	14 systems, >80 Tbps capacity

Latency to Europe	<8ms (Dakhla to EU IXPs)
Latency to Financial Centers	<30ms (London, Frankfurt, Paris)
Redundancy	Dual diverse paths on all routes
Peering	Direct at major European IXPs
Backbone	100G+ coherent DWDM backbone

Table 3: Network connectivity specifications

Security and Compliance

Harch Intelligence facilities are built to Tier IV+ standards, providing the highest level of redundancy and fault tolerance available in the data center industry. Physical security measures include multi-factor access control, biometric authentication, 24/7 CCTV surveillance, and mantrap entry systems. Logical security is ensured through Harch Technology's Cybersecurity Suite, which includes DDoS protection, intrusion detection and prevention, encrypted data at rest and in transit, and compliance monitoring for GDPR, ISO 27001, SOC 2 Type II, and Moroccan data protection regulations. Data sovereignty is guaranteed by keeping all data processing within Moroccan jurisdiction, with no cross-border data transfers unless explicitly authorized by the customer.

Certification/Standard	Status	Scope
ISO 27001	In Progress	All hubs
SOC 2 Type II	In Progress	All hubs
GDPR Compliance	Compliant	Data processing operations
Tier IV+ Design	Certified	All new construction
PCI DSS	Planned	Payment processing workloads
Moroccan DPA	Compliant	All operations

Table 4: Security certifications and compliance status

GPU Hardware Specifications

Harch Intelligence's GPU fleet is built on NVIDIA's H100 and H200 accelerator platforms, providing the computational density required for large-scale AI training and inference workloads. Each GPU node is configured with NVLink and NVSwitch interconnect for maximum intra-node communication

bandwidth, and interconnected via a 400Gbps InfiniBand fabric for distributed training across multiple nodes. The platform supports all major AI frameworks including PyTorch, TensorFlow, and JAX, with pre-configured environments for common workloads such as large language model training, computer vision, and scientific computing.

Component	H100 Specification	H200 Specification
GPU Memory	80 GB HBM3	141 GB HBM3e
Memory Bandwidth	3.35 TB/s	4.8 TB/s
FP16 Compute	1,979 TFLOPS	1,979 TFLOPS
FP8 Compute	3,958 TFLOPS	3,958 TFLOPS
Interconnect	NVLink 4th Gen	NVLink 4th Gen
TDP	700W	700W
NVLink Bandwidth	900 GB/s	900 GB/s

Table 5: NVIDIA H100 and H200 GPU specifications

Carbon-Aware Scheduling

Harch Intelligence's proprietary carbon-aware scheduling system represents a paradigm shift in sustainable AI compute. Rather than treating all compute as equivalent, our scheduler continuously monitors the real-time carbon intensity of energy generation at each hub and routes workloads to the cleanest available compute. This is made possible by our vertical integration with Harch Energy, which provides real-time telemetry from solar and wind installations at each hub location. The result is an average carbon intensity of approximately 47 gCO₂/kWh across the platform, with our Ouarzazate hub achieving just 18 gCO₂/kWh, making it one of the cleanest AI compute facilities on the planet. For comparison, the global average for data center operations is approximately 440 gCO₂/kWh, meaning Harch Intelligence operates at 89% below the industry standard.

Key Differentiator: Our carbon-aware scheduling is not based on renewable energy credits or offset purchases. It is genuine operational carbon reduction achieved by routing workloads to hubs where renewable energy is being generated in real-time. When the sun is shining in Ouarzazate, training jobs are routed there. When the wind is blowing in Dakhla, inference workloads shift to that hub. This dynamic optimization is only possible because we own and operate both the energy and compute infrastructure.

Request a Technical Briefing

For detailed technical specifications, pricing, or to schedule a technical briefing with our infrastructure engineering team, please contact us at harchcorp.com or reach out to our solutions architects directly. We offer custom configurations for enterprise AI workloads, government sovereign cloud deployments, and research institution compute partnerships.