

EUROPE'S FOUNDATION FOR THE NEXT WAVE OF AI FACTORIES

At Penta Infra, we believe in being successful together by growing together. Our AI-Ready Colocation Suite provides the high-density power, advanced cooling, and scalable infrastructure needed for artificial intelligence workloads.

As an experienced partner in state-of-the-art datacenters, we deliver four progressive tiers that take you from traditional computing to full AI-factory readiness, all with NVIDIA-certified infrastructure and sovereign deployment in Europe.





YOUR AI FUTURE SHOULDN'T BE A CLOUD RENTAL. OWN IT.

Escape escalating cloud costs and regulatory uncertainty. Build your sovereign AI factory on a foundation you control.

The public cloud was a starting point for AI, but scaling has revealed its true cost. As your AI ambitions grow, so do the bills and the risks. Scaling AI on shared infrastructure means accepting someone else's pricing, someone else's roadmap, and someone else's jurisdiction.

Our AI-Ready Colocation Suite puts you back in control with predictable costs, dedicated capacity, and the legal clarity that comes with sovereign European deployment. We provide the certified, high-density power and direct-to-chip cooling your NVIDIA clusters demand, all within a secure, sovereign European environment. It's time to build your AI future on your own terms.

AI-READY COLOCATION

Four levels built for your AI journey

Whether you're exploring AI for the first time or deploying sovereign AI infrastructure, our tiered approach provides a clear path forward with proven technology and scalable capacity.



Tier 1

Enhanced HPC Rack/The AI Sandbox

ENTRY-LEVEL FOR AI EXPLORATION

For organisations validating their AI strategy. Test models, run analytics, and explore AI's potential on powerful hardware without the cost and complexity of a facility retrofit. This is your risk-free first step.

- Up to 80 kilowatts (kW) per rack
- Traditional air cooling with optional liquid upgrades
- Dual power feeds for reliability



Tier 2

AI Starter

YOUR FIRST PRODUCTION AI CLUSTER

For teams moving from pilot to production. This is the turn-key solution to deploy your first true AI factory, engineered to get your initial 4-8 rack cluster online and delivering business value, fast.

- 80–100 kW per rack
- Liquid-to-air cooling systems
- High-capacity power distribution (800 ampere busways)
- Deploy 4–8 racks for pilot AI projects
- NVIDIA DGX platform ready



Tier 3

AI Production POD

ENTERPRISE-SCALE AI TRAINING AND INFERENCE

For businesses competing at scale. When AI is core to your strategy, this NVIDIA-certified, production-grade POD provides the performance and redundancy to rival a custom-built facility, at a fraction of the time and capital cost.

- 100–150 kW per rack
- NVIDIA-certified liquid-to-liquid cooling
- Production-grade redundancy (4-to-make-3)
- Deploy 8–18 rack clusters
- Full compliance with NVIDIA reference designs



Tier 4

Extreme

FUTURE-PROOF LEADERSHIP TIER

For national programs and sovereign AI initiatives. This tier offers rare global capacity for next-generation platforms, ensuring you stay ahead of the hardware curve and can support the most demanding AI workloads of the future.

- 150–300 kW per rack
- Advanced direct-to-chip cooling
- Modular power substations
- Support for next-generation NVIDIA platforms
- Sovereign AI and national program ready

DEPLOYMENT OPTIONS

Scale at your pace

FLEXIBLE FOOTPRINTS

- Start with a single rack
- Expand to full rows with aisle containment (8-12 racks)
- Deploy complete POD's (modular units of 16-32 racks)

01

AI-Ready Racks

Start your AI validation immediately. Our high-density racks are pre-engineered for rapid deployment, supporting the 80kW+ power and advanced cooling (air or liquid) required for your first NVIDIA GPU clusters. This is your risk-free first step to test models and accelerate your AI strategy without re-engineering your own facility.

02

AI-Ready Cages & PODs

Move from pilot to production with a turn-key, private AI factory. Our secure Cages and PODs are purpose-built to host multi-rack, NVIDIA-certified clusters, delivering 100kW+ per-rack densities with redundant, direct-to-chip liquid cooling. This is the foundation for your enterprise-scale AI, balancing raw performance with European data sovereignty.

03

Dedicated AI Suites

For national-scale programs and sovereign AI initiatives, our Dedicated Suites offer a fully bespoke environment. We partner with you to design and build a multi-megawatt, private AI hall, engineered to your exact specifications. This solution provides the ultimate control over your security, privacy, and cooling architecture, creating a long-term, custom-built foundation for your most strategic AI assets.

POWER INFRASTRUCTURE

- Traditional 400-volt systems
- High-efficiency DC power options
- Dedicated power blocks for POD deployments

REAL-TIME MONITORING

- API access for infrastructure telemetry
- Compatible with NVIDIA Base Command Manager(tbd)
- Integration with standard datacenter management platforms
- Track power, cooling, and environmental metrics

COOLING INTEGRATION

- Bring your own cooling equipment or use facility systems
- Clear connection points defined(power/cooling/telemetry)
- Compatible with NVIDIA-certified equipment

SUSTAINABILITY

Green by design



100% RENEWABLE ENERGY

- Wind and solar power purchase agreements at all locations
- Net-zero operational emissions



INDUSTRY-LEADING EFFICIENCY

- Power Usage Effectiveness (PUE) as low as 1.2
- Traditional data centers average 1.6+ PUE
- Lower PUE = more efficient operations



HEAT REUSE PROGRAMS

- Waste heat warms Brussels and Hamburg neighborhoods
- 60 - 80% of thermal output redirected to district heating



CIRCULAR OPERATIONS

- Rainwater harvesting systems
- Solar panels integrated into building facades

OUR LOCATIONS



Penta Infra BRU01

Europe's triple-certified AI node in Belgium

For EU research clusters, startups requiring green credentials, sovereign AI initiatives

Tier Availability:

- Tier 1: Available now
- Tier 2: Q4 2025
- Tier 3: Q1 2026
- Tier 4: Planned 2026

Penta Infra HAM01 High-density AI hub of Northern Germany

For Nordic connectivity, large footprint requirements, low latency to Northern Europe

Tier Availability:

- Tiers 1 & 2: Available now
- Tiers 3 & 4: Live 2026



Penta Infra AMS01

Strategic AI gateway for Europe in The Netherlands

For cloud connectivity, fast data exchange across Europe, premium network access

Tier Availability:

- Tiers 1 & 2: Q1 2026
- Tier 3: Q4 2026
- Tier 4: Feasibility 2026



Capacity	7,2 MW	4,4 MW	8 MW
Efficiency	1,2 PUE (among Europe's best)	1,25 PUE	1,25 PUE
Certifications	• BREEAM Excellent • EDGE Advanced • Uptime Institute Tier III • ISO 27001 • ISO 9001	• EN 50600 • ISO 22237 • Uptime Institute Tier III • ISO 27001 • ISO 9001 • SOC 2 Type II	• Uptime Institute Tier III • ISO 27001 • ISO 9001 • SOC 2 Type II
	Location: Near Brussels Airport and EU institutions	Innovation: District heat reuse project with Hamburg Energy Campus. Use of natural refrigerants	Connectivity: Direct access to Amsterdam Internet Exchange (AMS-IX)



Penta Infra AI-ready data center

Penta Infra data center

SERVICES & SUPPORT

Always available when you need us



24/7 OPERATIONS

- Remote hands support
- AI-ready commissioning assistance
- Guaranteed response times



NVIDIA INTEGRATION

- DGX, GB200/GB300 deployment guidance
- Cooling equipment qualification support
- Base Command Manager integration (tbd)



NETWORK CONNECTIVITY

- Cross-connects within facility
- Metro fiber to major hubs
- Cloud on-ramps (AWS, Azure, Google)



CUSTOMER PORTAL

- Easy ticketing system
- Real-time infrastructure monitoring (tbd)
- Power and cooling budget tools (tbd)

KEY METRICS

WHAT WE DELIVER

OUR COMMITMENT

Uptime SLA

99.982%

Power redundancy

N+1 to 4-to-make-3

Cooling redundancy

N+1 pumps and systems

Monitoring

Real-time API access

Compliance

ISO 27001, GDPR-ready



TERMINOLOGY

Understanding the key principles

POWER DENSITY (KW/RACK)

How much electricity each equipment rack can use. AI requires 2-6× more power than traditional IT.

POD

A modular data center unit containing multiple racks with shared power and cooling.

TIER III

Industry reliability standard ensuring 99.982% uptime with no single point of failure.

LIQUID COOLING

Advanced cooling using liquid instead of air. Essential for AI because chips generate extreme heat.

PUE (POWER USAGE EFFECTIVENESS)

Efficiency metric. A PUE of 1.2 means for every 1 watt used by IT equipment, only 0.2 watts are used for cooling and infrastructure. Lower is better.

N+1 REDUNDANCY

One extra backup system, so operations continue even if one component fails.





GET IN TOUCH

AI INFRASTRUCTURE TEAM

 penta-infra.com

 sales@penta-infra.com