

AI FACTORY OR AI-READY DATA CENTRE

Structural Reality Verses Semantics



THE TERM “AI FACTORY” IS GAINING MOMENTUM

Led by [Nvidia](#) and echoed by hyperscalers like [AWS](#), investors and AI-as-a-service providers, it describes a new class of purpose-built facilities designed to orchestrate the entire AI lifecycle under one roof, from data ingestion to model training and inference. Unlike traditional data centres running mixed workloads, these environments are framed as production engines for intelligence at scale.

In some regions, such as the Nordics, the term has taken hold. In others, operators still refer more practically to “AI-ready data centres”. But the real issue is not what we call them. What matters is the structural shift now underway.

AI workloads are forcing fundamental changes in how data centres are designed, retrofitted and operated – from cooling architectures and network design to power density, deployment timelines and workforce capability.

Beneath the headlines are immediate engineering challenges: building infrastructure that can physically sustain AI workloads, and doing so at speed. For organisations aiming to compete in the AI economy, success will depend not on terminology, but on their ability to address these realities.



1. THE CHANGING RULES OF AI INFRASTRUCTURE

AI workloads are fundamentally different from the traditional, CPU-led environments most data centres were designed to support.

GPU infrastructure generates extraordinary heat density. The airflow and cooling strategies that were sufficient even three or four years ago are no longer viable. Traditional air cooling simply cannot manage the thermal output of modern GPU clusters. Liquid cooling is no longer a forward-looking enhancement. It is becoming a baseline requirement.

The same shift is happening in networking. High-density, high-performance connectivity – most notably [InfiniBand](#) in AI environments – underpins GPU performance. This is not interchangeable with conventional copper installations. It requires different design principles, deployment approaches and specialist skillsets.

In practical terms, AI-ready infrastructure demands:

- ✓ Structured cabling that supports airflow, density and scalability
- ✓ Liquid cooling designed for sustained high-density GPU clusters
- ✓ High-performance, low-latency networking architecture
- ✓ Rack and layout strategies engineered for extreme compute intensity

Airflow, cable density, rack configuration and cooling are now interdependent. You cannot optimise one without understanding its impact on the others. AI readiness is not about inserting GPUs into an existing footprint; it is about rethinking the entire environment.

For operators of existing facilities, this often means retrofitting – and retrofitting quickly. Data centres that cannot support liquid cooling and high-density networking will struggle to attract AI tenants, particularly the fast-growing neo clouds.

CASE STUDY

InfiniBand expertise for AI infrastructure, Sweden

When a leading AI hyperscaler expanded into Falun, Onnec delivered a bespoke InfiniBand cabling solution for its new European campus. Building on experience from the client's previous site, the team deployed trained engineers to ensure consistent standards and quality under a tight delivery schedule.

This included:

12MW
ACROSS 3 DATA HALLS

17,200
MPO8 FIBRE CABLES

2,150
LC-LC FIBRE PATCH CORDS



2. SPEED NOW TRUMPS COST

The commercial dynamic behind AI infrastructure is unlike anything the sector has experienced in recent memory.

Traditional data centre projects could run on six to twelve-month RFP cycles. There was time for extended planning, phased procurement and gradual mobilisation. AI deployments do not operate on that timeline.

Once capacity is secured, mobilisation can now be required within weeks.

This urgency is driven by capital. Billions are being invested into AI-as-a-service providers across Europe and globally. These organisations are under pressure to deploy quickly and begin generating returns.

In this environment, the decision criteria have shifted:

- ✓ **Speed of deployment outweighs marginal cost savings**
- ✓ **Proven delivery capability outweighs theoretical design expertise**
- ✓ **Scalability outweighs incremental optimisation**

**THE PRIMARY QUESTION IS NO LONGER “WHO IS CHEAPEST?”
IT IS “WHO CAN GET US LIVE SAFELY AND FAST?”**

The financial stakes reinforce the urgency. Individual GPU deployments represent enormous capital value. Delays caused by hardware lead times – currently stretching into weeks and likely to increase – can materially affect commercial timelines. Logistics risks and installation errors are amplified when the hardware involved carries multimillion-pound valuations.

The result is a market behaving more like a technology land grab than a conventional infrastructure cycle. Everyone is racing to secure their position. And in this race, hesitation is expensive.

CASE STUDY

Mobilising for AI growth, Spain

As part of a leading AI hyperscaler's rapid European expansion, Onnec delivered a high-density InfiniBand deployment in Barcelona – mobilising quickly to meet an aggressive programme aligned to accelerated AI growth.

This included:

8MW
ACROSS 2 DATA HALLS

31,948
MPO8 FIBRE CABLES

4,010
LC-LC FIBRE PATCH LEADS



3. ORCHESTRATION AND CHOREOGRAPHY

If infrastructure is the visible component of AI readiness, orchestration is the hidden one.

GPU lead times must align with data centre access. Specialist cabling availability must align with installation teams. Cooling upgrades must be complete before hardware arrives. Skilled engineers must be ready to mobilise, often with very limited notice.

Planning has become the most complex element of AI deployment because so many variables are interdependent, this means:

- ✓ **Tight coordination between hardware suppliers and infrastructure teams**
- ✓ **Real-time flexibility around access and mobilisation**
- ✓ **Dedicated design alignment to specific accounts**
- ✓ **Capacity to deliver concurrent projects across regions**

Unlike traditional projects, where timelines were predictable and procurement cycles stable, AI deployments demand agility. Installers may receive confirmation only weeks before needing to be on site. Design continuity becomes critical; even temporary knowledge gaps can slow progress.

This is where the “factory” analogy holds some weight – not in terms of automation, but in choreography. Hardware, infrastructure, skills and supply chain must move in synchronisation. If one component falters, the entire deployment is at risk.

4. SKILLS CONSTRAINTS AND THE NEED FOR SCALABLE EXPERTISE

One of the least discussed constraints in AI infrastructure is people.

InfiniBand and other AI-specific networking installations are not legacy skillsets. Many infrastructure teams have had to upskill rapidly, learning on live deployments and accelerating knowledge transfer across regions. GPU installations are materially different from conventional CPU environments, both in technical complexity and risk profile.

At the same time, demand is exploding across geographies, particularly in regions offering competitive renewable energy and land availability. Delivering multiple AI deployments simultaneously requires not just competence, but depth of resource.

AI readiness is therefore about more than hardware. It requires:

- ✓ **Specialist knowledge in high-density AI networking**
- ✓ **Teams trained in liquid cooling integration**
- ✓ **Resource depth to support simultaneous regional expansion**
- ✓ **A delivery culture built around speed and precision**

**IT IS NOT ENOUGH TO HAVE ONE EXPERIENCED TEAM.
YOU NEED SCALABLE EXPERTISE.**

Organisations that anticipated this shift early, invested in upskilling and built strong regional partnerships now have a structural advantage. Those that did not are racing to catch up.

WHY EXPERT PARTNERS NOW MATTER MORE THAN EVER

In the AI era, infrastructure mistakes are amplified.

Install failures damage reputations. Delays disrupt funding cycles. Misjudged cooling strategies constrain scalability. In a market where tenants are heavily capitalised and moving at pace, there is limited tolerance for error.

Whether described as an AI factory or an AI-ready data centre, the terminology is secondary. What matters is whether the infrastructure works – technically, operationally and commercially. It must:

- ✓ **Sustain extreme density**
- ✓ **Integrate cooling and networking seamlessly**
- ✓ **Be deployed at speed**
- ✓ **Be delivered consistently, without failure**

EXPERT PARTNERS ARE NOT AN OPTIONAL LAYER IN THIS EQUATION. THEY ARE ALSO A RISK MITIGATION STRATEGY.

Partners that understand high-density networking, liquid cooling integration, rapid mobilisation and concurrent regional delivery are not simply service providers; they are enablers of AI growth. Their ability to anticipate bottlenecks, coordinate supply chains and execute under compressed timelines directly impacts time to revenue.

Call it an AI factory or call it an AI-ready data centre. Either way, the winners will be those who focused on the practical realities – and partnered with the expertise required to get it right.

ABOUT ONNEC

Onnec is a leading independent IT infrastructure solutions and services partner, delivering for major organisations worldwide, every day. We specialise in structured cabling, managed services and advanced network solutions.

Our highly skilled designers, project managers and technicians, supported by world-class vendor partnerships, design, deliver and manage business-critical IT environments across the globe. Regardless of size or location, every project receives the same level of dedication, expertise and commitment to excellence that defines Onnec's work worldwide.

Onnec's expertise spans all environments and can support customers with:

- Structured cabling design and installation
- Installation of cabling, ODFs, PDUs, GPUs and containment solutions
- Network hardware installations, changes and support
- Connectivity and equipment upgrades and changes
- Smart Hands support services

