




NEXT Frontier: NEXTDC's AI in Education Series

The AI Innovation Race in Universities

Why Falling Behind in Data-Centre
Infrastructure Means Falling
Behind in Research



N E X T D C
where AI lives™



The New Battleground for AI Research: Infrastructure as a Strategic Advantage

The pursuit of groundbreaking AI research in **universities** is no longer a gentle scholarly pursuit; it has evolved into a global “AI innovation race.” **Universities** worldwide are vying for top talent, breakthrough discoveries, and critical research funding. In this intense competition, one factor consistently separates the leaders from those left behind: the presence of cutting-edge, purpose-built data centre infrastructure. Without the right foundation to power next-generation AI, even the brightest minds and most ambitious research agendas will stall.

For deans, chancellors, leading AI professors, and chief digital officers, the message is clear: inaction or underinvestment in robust data centre infrastructure for AI is not merely a delay; it’s a strategic retreat. The stakes are too high, and the pace of innovation too swift, to be anything but proactive.



The Hidden Cost of Falling Behind: What Universities Risk in the AI Infrastructure Race

While the benefits of robust AI infrastructure are clear, the risks of inaction are often underestimated - and they're far more damaging. In today's global AI race, failing to invest doesn't just slow progress; it accelerates decline. Here's what that looks like in practice:

- **Losing the Talent War:** Top-tier AI researchers and STEM students increasingly evaluate institutions based on their access to cutting-edge compute infrastructure. Today's students, especially in data science and engineering, expect to learn and build with state-of-the-art tools. Access to the latest high-performance GPUs from providers like NVIDIA, AMD, and others is no longer a luxury; it's a baseline requirement for serious AI research. Without infrastructure capable of supporting these high-density, high-power chips, institutions risk more than just losing elite faculty - they risk turning away the next generation of innovators. The brightest minds will go where their ideas can be trained, tested, and scaled without limitation, accelerating a talent drain and eroding long-term competitiveness in the global AI ecosystem.
- **Missed Breakthroughs and Slower Publications:** AI research thrives on speed - of iteration, of training cycles, of insight. Training large models or running complex simulations is only feasible with infrastructure designed for sustained, high-performance workloads. Where top-tier institutions process datasets in hours, under-resourced environments can leave researchers waiting days or weeks. That lag leads to missed conference deadlines, delayed publication cycles, and lost opportunities for discovery, ultimately reducing global research visibility.
- **Shrinking Grant Opportunities & Lost Industry Partnerships:** Research funding agencies and industry collaborators are increasingly looking for universities with the proven infrastructure to deliver results at scale. If your institution's compute capabilities are a known bottleneck, securing competitive grants or lucrative partnerships for large-scale AI projects becomes significantly harder. Funders want to see that their investment will translate into tangible progress, not get bogged down by insufficient hardware.
- **Diminished Institutional Prestige and Rankings:** A university's reputation is built on its research output and impact. A consistent lag in AI research, fewer high-impact publications, and a struggle to attract leading scholars will inevitably affect global university rankings and your institution's standing among peers. In the long term, this can impact everything from student enrolment to alumni engagement.
- **The Compounding Cost of Delay:** The pace of AI advancement is not linear. Each new generation of GPUs (like the jump from NVIDIA H100 to H200 with its double inference speed on specific models) creates an even larger performance disparity. The longer a university waits to invest in infrastructure capable of supporting these innovations, the more insurmountable the technical and financial gap becomes. What was a sprint turns into a marathon you can't win. This isn't just a university problem, it's a national one. As other countries invest in sovereign GPU infrastructure and AI-ready ecosystems, universities that lag risk sidelining their country's role in global innovation leadership.

Falling behind in AI infrastructure limits universities' ability to contribute to some of the most urgent challenges of our time, climate resilience, pandemic response, Indigenous language preservation, energy innovation, and more.

Inaction doesn't just affect rankings and research output. It affects society's ability to innovate, respond, and lead.



From Data Centres to AI Factories: What Modern Research Really Needs

The core of your university's AI research powerhouse isn't just the chips; it's the infrastructure that enables them. Without purpose-built environments designed to support the extreme power, cooling, and networking demands of today's AI accelerators, even the most advanced hardware will underperform or remain idle. Supporting large-scale AI workloads - like model training, inference, and simulation, requires infrastructure that's as sophisticated as the research itself. Consider these critical requirements:



Extreme Power Density:

Modern AI servers – stacked with cutting-edge accelerators from vendors like NVIDIA, AMD, and others – push power requirements well beyond legacy thresholds. Individual chips can draw over 700 watts, and when integrated into dense multi-GPU architectures, a single rack can consume 50kW or more. That's over three times the load most university server rooms were originally built to handle, rendering many existing facilities functionally obsolete for AI-scale research.

Advanced Thermal Management:

The heat generated by these chips is immense. Traditional air conditioning is insufficient. Effective AI infrastructure requires sophisticated cooling solutions like direct liquid cooling (cold plate technology) or high-capacity rear-door heat exchangers to efficiently dissipate heat and prevent thermal throttling, ensuring your GPUs run at peak performance. And as research needs evolve, the ability to scale quickly - without reengineering infrastructure or disrupting operations - becomes a strategic imperative.

High-Bandwidth Interconnection:

AI workloads are data-intensive. Fast, reliable networking both within the cluster (e.g., InfiniBand, NVLink) and to external data sources and cloud environments is crucial.

Resilience and Security:

24/7 uptime, redundant power, robust physical security, and advanced cybersecurity measures are non-negotiable for protecting valuable research, intellectual property, and ensuring continuous operation. Additionally, modern colocation facilities are optimised for energy efficiency and low carbon operations, critical for universities striving to meet their ESG and sustainability targets.

Attempting to engineer and sustain this level of infrastructure internally is often a monumental, budget-draining, and delay-prone undertaking for a university, particularly when considering the scale required for AI Factory or DGX SuperPOD™ deployments. Beyond the capital and complexity, the true cost of attempting to build AI-scale infrastructure in-house is time, delayed research milestones, lost grant windows, and slower innovation cycles.



Why Partner with a Specialised Provider like NEXTDC?

To avoid the pitfalls of falling behind and to genuinely compete in the AI innovation race, universities need a strategic partner for their data centre infrastructure, one that already has the expertise, scale, and future-ready capabilities.

"Leveraging a specialised provider like NEXTDC for your GPU-accelerated infrastructure allows your university to maximise its research ambitions without the complexities of building and managing bespoke infrastructure."

- **NEXTDC is NVIDIA DGX Data Centre Ready certified**

Our facilities are purpose-built and independently certified to support the most demanding AI workloads, delivering optimal performance and reliability for NVIDIA DGX systems, including SuperPOD™ deployments and AI Factory architectures. This certification provides assurance that our data centres meet NVIDIA's exacting standards for power, cooling, and operating conditions, backed by proven experience in aligning infrastructure to the latest and future NVIDIA architectures.

- **Shift from CapEx to OpEx & Lower TCO**

This reduces massive upfront capital investment by leveraging purpose-built, AI-ready infrastructure, delivering a lower Total Cost of Ownership (TCO) over the long term, avoiding the escalating costs and procurement delays of on-campus expansion.

- **Immediate Access to Certified AI-Optimised Facilities**

Our facilities are built to support high-density NVIDIA H100/H200 racks with advanced cooling and power efficiency. This offers immediate access to scalable power and cooling without delays or upfront capital expenditure, ensuring your GPUs run at peak performance with precise thermal management.

- **Scalable On-Demand Compute**

Grow capacity as your research expands - without waiting on construction or hardware availability. NEXTDC offers the agility to scale your AI infrastructure seamlessly, with modular, rapidly scalable infrastructure designed for AI growth, unlike limited on-campus options.

- **Expert Uptime & Operational Management**

Let specialists handle power, cooling, security, and compliance. NEXTDC's 24/7 operational excellence, highly certified security protocols, and dedicated support teams ensure maximum uptime, so your university's IT staff can focus entirely on research enablement. This effectively transfers the risk of planning, construction, and delivery of the data centre to NEXTDC, not the university.

- **Global AI Routing & Interconnection**

NEXTDC's extensive network ecosystem, including proximity to international subsea cable landings, enables ultra-low-latency, high-bandwidth cross-border AI workloads. This is crucial for collaborative projects, federated learning across institutions, and multi-region training across Asia-Pacific and beyond, providing a strategic edge through next-generation connectivity, including direct, high-speed links to research networks like AARNet.

- **Sustainability Alignment**

NEXTDC operates with sustainability at the core, from renewable energy commitments to advanced cooling systems that minimise environmental impact. By leveraging efficient, NABERS-certified data centres, universities can meet ambitious ESG goals while also supporting evolving global ranking criteria, such as the QS Sustainability Rankings. This ensures institutions can scale AI workloads responsibly while demonstrating visible leadership in sustainability.

- **Free Up Campus Space**

By shifting backend IT off-site, universities can reallocate valuable campus real estate to core academic priorities such as teaching, research labs, or student services.

The AI innovation race in universities is not waiting. Universities that invest strategically in the right data centre infrastructure today will be the ones attracting the best talent, making the biggest breakthroughs, and shaping the future of AI. Those that hesitate risk being left in the wake.



Rethinking AI Infrastructure: Smarter Partnerships, Greater Impact

NEXTDC's NVIDIA DGX-certified data centres are purpose-built to support the most demanding AI workloads, from training foundational models to deploying real-time inference at scale. With GPU-optimised design, high-density power delivery, and advanced liquid cooling, we deliver the performance and reliability needed to accelerate every stage of the AI lifecycle.

Our infrastructure is ready now, to help turn your AI ambitions into scientific breakthroughs, medical advances, and global research leadership.

Global Connectivity. Research Advantage.

Located near Australia's major subsea cable landing stations, NEXTDC enables ultra-low-latency international connectivity, essential for federated learning, cross-border model training, and real-time edge inference across Asia and the Pacific.

For universities building AI-as-a-Service platforms or leading innovation precincts, this level of connectivity isn't just infrastructure, it's the invisible engine behind world-class research.

Turn Vision into Impact

Whether you're:

- Advancing medical research,
- Educating the next generation of AI talent, or
- Launching a national AI innovation district,

Your infrastructure must match your ambition.

Work with NEXTDC to access powerful, reliable, and energy-efficient infrastructure across Australia and the Asia-Pacific. Our facilities are certified for NVIDIA DGX systems and ready to support the latest AI hardware, helping researchers and institutions move faster, scale smarter, and make a greater impact.

Download our AI Ready Checklist



Speak with a NEXTDC specialist



N E X T D C
where the cloud lives™

136 398

sales@nextdc.com

nextdc.com

This document is correct at the time of printing and is for presentation purposes only. This document does not constitute an offer, inducement, representation, warranty, agreement or contract. All information contained in this document (including all measurements, photographs, pictures, artist's impressions and illustrations) is indicative only and subject to change without notice. NEXTDC Limited, its employees, representatives, consultants and agents make no representations or warranties as to the accuracy, completeness, currency or relevance of any information contained in this document and accept no responsibility or liability whatsoever for any discrepancy between the information contained in this document and the actual data centres or services provided by NEXTDC Limited or for any action taken by any person, or any loss or damage suffered by any person, in reliance upon the information contained in this document. © 2025 NEXTDC Limited ABN 35 143 582 521.