

NIIN


National Industry Innovation Network

NIIN Sovereign AI Cloud Study Tour Report

Lessons from the US about Building Critical Sovereign
AI Infrastructure to power Research and Innovation in
Australia

March 2026





1. Forewords

Message from a delegate

The impact of AI on universities is becoming more visible and more consequential. Early adoption focused largely on staff use of generative AI tools to improve productivity. From there, AI use has extended into education where it is beginning to reshape learning, assessment, and the broader campus experience. At the same time, universities are exploring AI applications across core operational and service functions. Together, these developments signal that AI is no longer peripheral to the university mission – it is becoming (if not is already) foundational to how universities teach, operate, and engage.

Equally critical is the role of AI in research. As AI reshapes every sector of the economy, governments, industries, and communities are increasingly reliant on universities to generate independent evidence about its impacts, risks, and benefits. This includes not only applying AI to research problems, but also researching AI itself – its reliability, safety, governance, and societal consequences.

To do this well requires more than enthusiasm for AI – it requires access to secure, affordable, controllable infrastructure and trusted data environments. Researchers either cannot obtain the GPU capacity they require, or they cannot afford it at scale. Where access is available, it is often fragmented, externally controlled, or poorly aligned with Australian data sovereignty expectations. This creates a growing trust deficit where researchers and institutions are unsure where data is held, who ultimately controls the infrastructure, and whether regulatory and ethical obligations are being met. Alongside this is a data

gap – insufficient access to high-quality, well-governed data needed to support meaningful AI research across disciplines.

The Australian Government’s Strategic Examination of R&D highlights the importance of research infrastructure, but leading institutions such as Stanford University and University of San Francisco made it clear that infrastructure alone is not enough. Skills, tools, governance frameworks, and institutional confidence are equally important if universities are to capture value from AI whilst managing risk.

The study tour brought together ten universities with a shared interest in understanding how leading institutions and global providers are addressing these challenges, and what lessons might be applied in the Australian context. There is a strong case for collaboration – particularly where infrastructure control, security and data governance are concerned.

The AI Study Tour provided a clear view of the current frontier of AI-enabled education, research, and operations. It reinforced the responsibility of universities to lead the adoption of AI for public good, while also acting as critical researchers of AI itself. This dual role – using AI to improve institutional performance whilst generating independent evidence to inform society – will be central to how universities contribute to the national capability, economic resilience, and public trust in the years ahead.



Jonathan Pheasant

Vice President (Corporate Services) Flinders University

Message from Cisco

Immediately after the Study Tour, some of our delegation attended NVIDIA's GTC Conference in San Jose.

I have attended many technology conferences over the years, but this year's NVIDIA GTC was notable from the outset. A sold-out arena at the SAP Center, more than 2,000 speakers, and 1,000 technical sessions signalled that this was more than a product announcement – it was a statement of intent from an industry moving with considerable purpose.

Aside from announcing new AI accelerated compute hardware availability from Nvidia, Jensen Huang's keynote emphasised OpenClaw – an open-source autonomous AI agent platform that skyrocketed to popularity even faster than Linux.

Where conventional AI tools respond to queries, these agents complete tasks, make decisions, and take action with minimal (but appropriate) human oversight. Jensen was unambiguous about the significance of the moment, describing OpenClaw as having achieved in weeks what took Linux thirty years, and asserting that every technology company must now develop a considered OpenClaw (Agentic AI) strategy. The enterprise implications are substantial. Capability without governance however, is not a strategy –

it is an exposure. For those of us focused on enterprise and sovereign AI, this was the most consequential part of the conversation.

NVIDIA's response NemoClaw (a hardened package of nvidia's NemoTron AI model), along with OpenShell – a security control layer delivering a serious attempt to support locally run, sovereign, autonomous AI agents.

Beyond the keynote, what distinguished this year's GTC was its breadth and technical depth. Over 200 research posters and thousands of breakout sessions reinforced that this is not simply a commercial event, it is where the industry convenes to determine what comes next.

For Cisco, our presence in that room – and Cisco's designation as a design partner in the agentic AI security layer with Defense Claw – highlights where Cisco sits within the AI infrastructure & governance stack now being assembled at industrial scale.


The conclusion is clear: the period of AI experimentation is drawing to a close. What is being constructed now is the foundation for something considerably larger.




David Siroky

CTO AI, APJC Cisco

Universities that participated in the study tour:





2. Context

AI is reshaping the operating environment for universities across every dimension including how research is conducted, how students learn, and how institutions manage operations. The pace of change driven by AI is without precedent: model releases have increased 500% in two years and the share of Australian organisations classified as sophisticated AI adopters rose from 4% to 22% in a single year .

The Australian Government's National AI Plan set out three priorities for Australia to turn AI

adoption into competitive advantage: Building smart digital and computing foundations; Backing Australian AI capability; and Attracting major AI investment.

The Plan outlines the imperative for coordinated effort across government industry and universities. The Cisco AI Study Tour to the United States was an opportunity for Australian universities to plan their own response to AI and collectively agree on what sovereign AI capability might look like. This report captures major themes and potential next steps.

Study Tour briefings and site visits:



Cisco Executive Briefing Centre. Cisco's Silicon Valley HQ and anchor for Cisco's US\$9.3B annual spend on research and development, including in AI.



NVIDIA. Acknowledged as the world's leading AI computing company, NVIDIA presented insights into its GPUs, systems and software stack including its future roadmap.



XENON Systems. An Australian systems integrator focused on high-performance computing and AI. XENON is NVIDIA'S ANZ Enterprise Partner of the Year.



Stanford University. One of the world's leading research and teaching universities, located in the heart of Silicon Valley and globally recognised for its track record in AI venture creation.




Sharon AI. An AI Neocloud with compute infrastructure designed to keep data and AI processing onshore. Sharon AI is launching Australia's first Cisco Secure AI Factory with NVIDIA.



University of San Francisco. A specialist in applied data science and AI education, recognised for its industry partnerships and whole-of-institution approach to embedding AI across disciplines.



VAST Data. A global technology company providing a unified data platform purpose-built for the AI era, unifying storage, database and compute into a



3. Major themes from the Study Tour

AI architectures and use cases are evolving for education, research and operations

AI is becoming genuinely more valuable. Models are more reliable, costs continue to fall, and the range of questions that researchers and educators can now ask – and answer – is expanding rapidly. This is not incremental improvement; it is a broadening of what is possible. The study tour surfaced three architectural shifts that will define how universities need to think about AI infrastructure in the coming years:

- **From conversational AI to agentic AI.** The most consequential shift underway is the move from AI that responds to prompts to AI that acts autonomously. As an example, Stanford University has used automated agents to double the success rate of molecule discovery and dramatically speed up the drug discovery pipeline. Agents require continuous inference – potentially running for months – rather than the periodic, spike-based demand of chatbots. High bandwidth and low latency are no longer optional upgrades; they are the baseline requirement for an agentic environment. Infrastructure that was designed for the chatbot era will not be adequate for the agentic one.
- **Security and trust as an enabler, not a barrier.** Organisations that have moved furthest with AI treat security as a foundation, not an afterthought. Agentic AI raises the stakes: an agent that can move funds, deploy code or alter policies demands more rigorous controls than a chatbot. The encouraging news is that security frameworks – such as the principle of least privilege, giving agents only the access they need – are maturing rapidly.

- **Unlocking the next wave of AI value through machine data.** The first wave of AI was powered by human-generated content (text, audio, video), and the next will be driven by machine-generated data including metrics, logs, traces and telemetry flowing from networked infrastructure. This transition will produce a different kind of intelligence: not just answering questions but predicting where problems will emerge before they occur. Organisations that learn to harness this data will be more resilient, efficient and dynamic.

“Australian universities know AI isn't going away. The question is whether we embrace it or get left behind. There's a growing recognition that the institutions that hesitate too long won't get to shape what comes next.”

- Study Tour Delegate

Safe, secure and responsible access to AI is a priority

Earlier waves of technology came with a genuine choice: adopt early, follow later, or opt out entirely. AI is different – it is opt-in only, and the cost of delay is not neutrality, it is falling behind. Stanford's view was striking: the greater risk is not moving too fast but too slowly – a deliberate choice made sustainable by keeping ethics as a live consideration throughout, rather than a checkpoint at the end.

The institutions moving most confidently are taking action across several fronts:

- **Creating sanctioned environments that enable safe experimentation with AI.** Without these, staff and students will inevitably turn to external tools outside institutional oversight, leaving universities blind to how AI is being used. Centrally governed platforms that route consumption through institutional infrastructure transform AI adoption from an unmanaged risk (including intellectual property leakage, non-compliance with government data requirements, and privacy violations) into a visible, governable capability.
- **Establishing governance frameworks that sit alongside, rather than obstruct, adoption.** At the University of San Francisco, all AI investments pass through centralised evaluation covering data controls, security and privacy – but alongside a university-wide push to embed AI across every discipline. Stanford dedicated AI advisory committees coordinate across academic integrity, technology and research, while faculty retain genuine latitude to innovate.
- **Gathering intelligence needed to drive responsible uptake.** Understanding how AI is being used across the organisation is a precondition for managing it well. Cisco's experience is instructive: tracking daily active users across AI tools provides a real-time proxy for organisational readiness, giving leadership the evidence base to push further investment. For universities, this visibility is the foundation for knowing where to invest, where to intervene, and how to demonstrate return.



Creating safe environments for AI experimentation

Stanford's AI Playground provides free access to multiple leading AI models through a single, Stanford-hosted interface. Rather than restricting access to external tools, Stanford channels demand into an environment it can see, manage and continuously improve. An API gateway sits alongside the Playground for developers and researchers who need direct model access for building applications.



Enabling safe access to AI agents

Cisco's CIRCUIT extends Stanford's approach into the agentic era. Designed to give employees safe, governed access to AI agents, CIRCUIT allows staff to interact with a registry of approved agents, connect their own agents to the platform, or have CIRCUIT build one for them using natural language.

Queries are automatically routed to the most appropriate agent and model, with data provenance maintained throughout.

What these examples have in common is they both provide a safe playground for AI experimentation with centralised visibility and governance, while not gating access to state-of-the-art AI models.

AI literacy is a pre-condition for effective adoption

AI literacy is frequently treated as a personal initiative – something individual staff and students pursue on their own time, using their own means. The study tour made clear that AI literacy needs to be built deliberately and systematically across the entire organisation – from researchers and faculty through to professional staff. Delegates observed that the unevenness of AI literacy – from enthusiastic early adopters through to active sceptics – created a two-speed organisation.

For many faculty and staff, AI remains synonymous with ChatGPT and Microsoft CoPilot – a framing that significantly underestimates both the technology's scope and what engaging with it seriously requires. AI is reshaping research workflows, professional practice, administrative operations and the nature of graduate employment simultaneously. Building genuine literacy means helping people understand what AI is, what it can and cannot do, and how to use it responsibly – not as a one-off training exercise, but as institutional capability built across every discipline and every level of the organisation.



Stanford's multi-pronged approach to building AI literacy

Stanford University has brought together a suite of interconnected resources designed to build AI fluency across the entire university community:

- **Resources for students** including an AI Guide that helps students think through when and how to use AI in their learning, building the critical thinking needed to evaluate AI tools as they continue to evolve.
- **Resources and training for instructors** including a searchable library of pedagogical examples drawn from real course designs across the university and structured disciplinary workshops where faculty explore the role of AI in their specific field.
- **Purpose-built AI tools** including a chemistry chatbot that guides students through problem-solving processes and a voice-mode tool for practising difficult conversations
- **AI Catalyst Team**, a dedicated consulting service that supports academics to incorporate AI into their teaching or research workflows




AI as a core graduate competency at the University of San Francisco

The University of San Francisco (USFCA) has embedded AI across courses in almost every discipline with a whole-of-institution commitment to ensuring every curriculum reflects how professionals actually work today.

The School of Law offers the most instructive example. Students use AI for mock client interviews and practice sessions with AI judges. They are also taught to critically evaluate AI-generated content, including how to identify and manage “hallu-citations” when using AI for legal citation.

USFCA also operates a practicum program through its Data Institute. Students spend nine months embedded within a partner company, working on real data and AI challenges alongside industry professionals. The outcomes are striking: 96% of graduates are employed full-time in data science and AI within six months of completing the program. The practicum also functions as a live intelligence feed back into curriculum design – the observation that “AI engineer” has become the most sought-after title in partner hiring directly shaped the development of a new AI engineering course.



4. The need for a Sovereign AI Cloud

One of the most striking aspects of the study tour was the sheer breadth and volume of AI workloads that US universities are running. For the universities we visited, AI is no longer an experimental add-on; it is embedded across the full spectrum of university activity.

The challenge for Australian universities in catching up is just not awareness of what is possible. It is having the infrastructure to do it. Three barriers are holding Australian institutions back:


- 1. Fragmented and insufficient GPU access.** Most universities procure compute in small, isolated batches and rely on HPC systems optimised for batch processing rather than iterative AI development. This creates long queues, unpredictable availability and limited ability to support peak demand.
- 2. IP and sensitive data at risk.** Many faculties lack secure, compliant environments for working with regulated datasets. Researchers turn to public AI tools out of necessity, exposing university and industry IP and undermining commercialisation pathways.

- 3. High costs from fragmented procurement.** Without pooled demand, universities face higher per-GPU costs, duplicated infrastructure and ongoing budget inefficiency.

NIIN AI Cloud

The study tour included a dedicated discussion on the proposed NIIN AI Sovereign Cloud – a shared, Australian-hosted environment that would provide universities and their industry partners with:

- **Sovereign GPU infrastructure** – on-demand compute configured for AI workloads, supporting model development, secure simulation and advanced analytics
- **A common services and applications layer** – including AI playgrounds, model access and development tools available to researchers, educators and students across the network
- **Embedded governance and responsible AI safeguards** – designed to host sensitive and regulated datasets in compliance with Australian requirements.



The NIIN AI Cloud would sit between Tier-1 HPC systems, which are built for batch-processed scientific workloads, and commercial AI cloud platforms, which are optimised for real-time services but are not always suitable for Australian sovereignty requirements.

NIIN AI Skills Challenge

The study tour was equally clear that capability gaps – in how staff, researchers and students understand and work with AI – are as much a constraint on adoption as infrastructure.

The establishment of a NIIN AI Cloud will be augmented by the NIIN AI Skills Challenge which is intended to deliver a continuum of training opportunities aligned to real-world needs. These include micro-credentials, workplace-based learning, digital apprenticeships, sector-specific short courses and collaborative projects with industry.

The program (which helps deliver on the Australian Government’s drive for tertiary harmonisation and connection between vocational and higher education) is designed to support both foundational AI literacy and the emerging skills required for the next wave of AI capabilities including agentic and physical AI. It will also develop the skills needed for management and operation of the NIIN AI Cloud itself.

Where to from here

The study tour made clear that the barriers holding Australian universities back are not primarily about awareness or appetite – they are structural. No single institution can build the

infrastructure, develop the skills pipeline, or establish the governance frameworks that the AI era requires on its own. The institutions that will lead are those that move from individual experimentation to coordinated, sector-wide action.

The NIIN AI Cloud has been in development since 2025 as one expression of that coordination, progressing through a structured sequence of activity designed to ground the initiative in sector feedback. The study tour is the most recent milestone in that journey, and its findings directly inform the next steps:

- **University Information Sessions (April 2026).** A series of information sessions with participating universities will share the current thinking on the AI Cloud prototype design and give institutions the opportunity to indicate their interest in participating.
- **The Prototype Phase (mid 2026).** The NIIN will move into a prototype phase to establish the foundational governance, architecture, and commercial constructs before committing to full-scale deployment. Two nodes will be identified during this phase, with all NIIN universities participating as users of the platform.

<ul style="list-style-type: none"> • NIIN launched 	<ul style="list-style-type: none"> • NIIN identifies AI as national priority in response to Australian Government’s Strategic Examination of R&D (Apr 2025) • NIIN AI Adoption Moonshot announced (Nov 2025) • RFI for NIIN AI Cloud released (Dec 2025) 	<ul style="list-style-type: none"> • NIIN AI Adoption Moonshot report released in response to Australian Government’s National AI Plan (Feb 2026) • NIIN roundtable with Cisco CEO Chuck Robbins (Feb 2026) • Cisco AI Study Tour to the US (Mar 2026) • University information sessions on NIIN AI Cloud (Apr 2026) • Commencement of prototype phase (mid 2026)
<p style="text-align: center;">2020</p>	<p style="text-align: center;">2025</p>	<p style="text-align: center;">2026</p>

5. Acknowledgement

The AI Study Tour represented something rare: ten universities setting aside institutional boundaries to learn together, ask hard questions, and think collectively about what the future demands of their institutions.

The delegation left the United States with a shared sense of both the scale of the challenge and the strength of the opportunity.

Australian universities are not starting from zero – they bring research excellence, strong institutional governance, and a genuine commitment to public good. What the tour reinforced is that translating those strengths into AI leadership requires accelerating from individual action to coordinated effort, and from careful observation to deliberate investment.

The findings in this report reflect the generosity of the organisations that hosted and briefed the delegation – Cisco, NVIDIA, VAST Data, XENON Systems, Sharon AI, Stanford University, and the University of San Francisco – who shared their experience, infrastructure, and thinking openly and without reservation. Also reflect the quality of the delegation itself. Many of the conversations that took place on the side were as valuable as the formal sessions. That spirit of openness and collaboration is the foundation on which the next steps will be built.



NIIN

niiin.com.au