Connecting the Dots

Where compliance, risk management, digital innovation, and bottom-line growth intersect.

EDITION 5 · NOVEMBER 2026



IN THIS EDITION

India's outsourcing engine stalls

Near-zero hiring for the first time

The consent box is dead

Australia's privacy overhaul puts every business on notice

The agent that climbed the fence

Governing AI agents after a government data breach

Prove it works

Australia's turn to assured AI

Resilience at machine speed

Reading the ASD cyber threat report

Dr Sukh Rashpal Sandhu | Subscribe free on LinkedIn

EDITION 5 | NOVEMBER 2026

Contents

Connecting the Dots: November 2026 3

India: The world's largest outsourcing engine records near-zero hiring for the first time: 6 million workers, $50 billion lost in one day 4

The Consent Box Is Dead: Australia's Privacy Overhaul Puts Every Business and Every Tech Giant on Notice 8

The Agent That Climbed the Fence: What the Medicare Incident Teaches Every Organisation Deploying AI Agents 18

Prove It Works: Australia's Turn From Voluntary AI Guardrails to Demonstrated Safety . 26

Resilience at Machine Speed: What the ASD Cyber Threat Report Tells Boards About the Year Ahead 34

Integrity by Design: Governing AI Across Schools, Training Providers and Universities ... 42

Follow us on LinkedIn 50

About Connecting the Dots 52

INTRODUCTION

Connecting the Dots: November 2026

By Dr Sukh Rashpal Sandhu

Edition 5 of Connecting the Dots: 6 pieces for November 2026.

Where compliance, risk management, digital innovation, and bottom-line growth intersect.

This edition brings together 2 pieces first published in Connecting the Dots on LinkedIn, each carrying the date it first appeared and 4 new features written for this edition. Republished pieces appear as they were first published, with their original images; facts, figures and links are as they stood on that date. New features are marked "Written for this edition".

In this edition

- India: The world's largest outsourcing engine records near-zero hiring for the first time: 6 million workers, $50 billion lost in one day
- The Consent Box Is Dead: Australia's Privacy Overhaul Puts Every Business and Every Tech Giant on Notice
- The Agent That Climbed the Fence: What the Medicare Incident Teaches Every Organisation Deploying AI Agents
- Prove It Works: Australia's Turn From Voluntary AI Guardrails to Demonstrated Safety
- Resilience at Machine Speed: What the ASD Cyber Threat Report Tells Boards About the Year Ahead
- Integrity by Design: Governing AI Across Schools, Training Providers and Universities

Thank you to every reader who subscribes, shares and responds.

Warm regards,
Dr Sukh Rashpal Sandhu
Author, Connecting the Dots

CONNECTING THE DOTS

India: The world's largest outsourcing engine records near-zero hiring for the first time: 6 million workers, $50 billion lost in one day

By Dr Sukh Rashpal Sandhu | First published 4 May 2026

The world's outsourcing engine is being rebuilt while it is still running.



The global outsourcing market, combining BPO (Business Process Outsourcing), IT outsourcing, and shared services, exceeded $1 trillion in 2025. Measured by the broadest definition of outsourcing services, including engineering and knowledge services, Grand View Research projects that figure reaching $7.1 trillion by 2030. The BPO segment alone sits at $302 to $348 billion in 2025 and is projected to reach $525 to $696 billion by 2030. Any way you measure it, this is not a sector in decline. It is a sector in violent transformation.

Yet 85 million jobs are projected to be displaced by AI globally by end of 2026. At the same time, 170 million new roles are projected by 2030 (World Economic Forum).

The question is not whether the jobs will exist. The question is whether the same workers will fill them.

Here is what is actually happening, country by country.

INDIA

6 million workers across IT and BPO. Approximately 7% of GDP. The original engine of global outsourcing. In the last 18 months, 31% of BPO roles have been altered or removed due to AI. India's major IT firms recorded near-zero net hiring in fiscal 2025/26. When Anthropic released its AI coding tool in February 2025, Indian IT stocks dropped sharply, wiping approximately $50 billion in combined market value. An IT union leader in Bengaluru has reported a sharp increase in job loss cases among members. Junior coding roles are disappearing first.

PHILIPPINES

1.9 million workers in the IT-BPM (Information Technology and Business Process Management) sector. $40 billion in export revenues in 2025. The sector contributes approximately 7.5 to 8.5% of Philippine GDP, with some estimates citing close to 10% when the broader digital services economy is included. The sector still added 80,000 jobs in 2025 and targets $42 billion in revenue by 2026. However, the sector's own industry body, IBPAP (Information Technology and Business Process Association of the Philippines), has confirmed that its 2.5 million jobs target for 2028 is no longer achievable. Voice-based customer service, the backbone of Philippine outsourcing, faces the highest immediate automation risk. AI agents now offer 24/7 availability, no fatigue, and no emotional variability.

SOUTHEAST ASIA (Vietnam, Indonesia, Malaysia)

A 52% increase in AI-driven job displacement in logistics and warehousing has been recorded across Southeast Asia since 2023. Vietnam and Indonesia are aggressively competing for outsourcing contracts moving out of higher-cost hubs. Asia Pacific as a whole is the fastest-growing outsourcing region globally, running at a compound annual growth rate of 9.3 to 13%.

EUROPE (Poland, Romania, Germany, Netherlands)

12 million European jobs are projected to be eliminated or fundamentally transformed by 2028. 68% of European employees already expect workforce reductions due to AI. Germany has automated 27% of its administrative public sector roles. Italy, slower to adopt, sits at 13%. The Netherlands recorded a 23% decline in transportation jobs due to AI-assisted delivery systems. Eastern Europe, particularly Poland and Romania, is rapidly gaining outsourcing contracts from Western Europe driven by time zone and language proximity.

LATIN AMERICA (Mexico, Colombia, Brazil)

Brazil's call centres with AI integration have reduced staffing needs by 29% in metropolitan regions. Latin America generated approximately $182 billion in outsourcing services in 2024 and is forecast to grow at 10.1% annually through 2030. Mexico offers salaries up to 47% lower than US equivalents and shares the same time zone, making it a preferred nearshore destination for American companies. Mexico ranks 58th in the Global Innovation Index 2025. Colombia is emerging rapidly, competing directly for US-facing outsourcing contracts.

AFRICA (Morocco, Senegal, Kenya, Uganda, South Africa)

Africa's BPO market is projected to reach $14.75 billion by 2033. Uganda launched a formal National BPO Policy in early 2025 with a target of 150,000 BPO jobs by 2030. Morocco and Senegal are serving francophone European markets. South Africa's GDP is forecast to grow by 1.2%, aided by AI productivity gains. This continent is no longer an outsourcing footnote.

UNITED STATES

More than 260,000 domestic BPO service companies. Approximately 300,000 jobs are outsourced from the United States each year. In 2025, 55,000 jobs were directly impacted by AI-driven automation. Canadian provinces with larger tech sectors recorded a 44% drop in junior support roles, versus 19% in non-tech regions. Call centre agents across North America saw a 41% drop in job openings at firms deploying AI voice assistants for tier-one support. North America is expected to lead global automation adoption at 70% penetration by end of 2025.

SOUTH KOREA and MANUFACTURING

36% of factory-based inspection roles in South Korea have shifted to machine vision (AI-powered visual inspection) systems. Manufacturing globally is forecast to lose 2 million jobs due to robotics and AI, with assembly line employment projected to fall from 2.1 million in 2024 to 1.0 million by 2030.

What the data actually says:

The IMF estimates 40% of all jobs globally face meaningful AI exposure. In high-income countries, that figure rises to 60%. 41% of employers globally plan to reduce headcount in AI-automatable roles within five years. Yet only 20% of customer service leaders have actually reduced headcount due to AI so far, according to Gartner. The gap between intent and execution is still wide, but it is closing fast.

Workers with demonstrable AI skills now earn a 25% wage premium over peers without them. AI job postings are 134% above 2020 levels. The median annual salary for AI roles reached $156,998 in Q1 2025.

***Yet* 1.4 *billion workers need reskilling by* 2028. *Only* 58 *million completed AI training in* 2025. *That is* 4.1% *of what is required*.**

The outsourcing industry is not dying. It is relocating, restructuring, and demanding more of the people inside it. What is dying is the assumption that a repetitive task, performed reliably and cheaply, will remain a human job indefinitely.

The floor of the labour market is rising. The question for every government, every training institution, and every worker is whether they are rising with it.

Sukh Sandhu

Director, CAQA | Editor-in-Chief | ISO Auditor | Poet and Writer

Forbes Business Council Member | United Nations Speaker

Top 100 Global Educator | RTO/AI/HE/ISO Expert

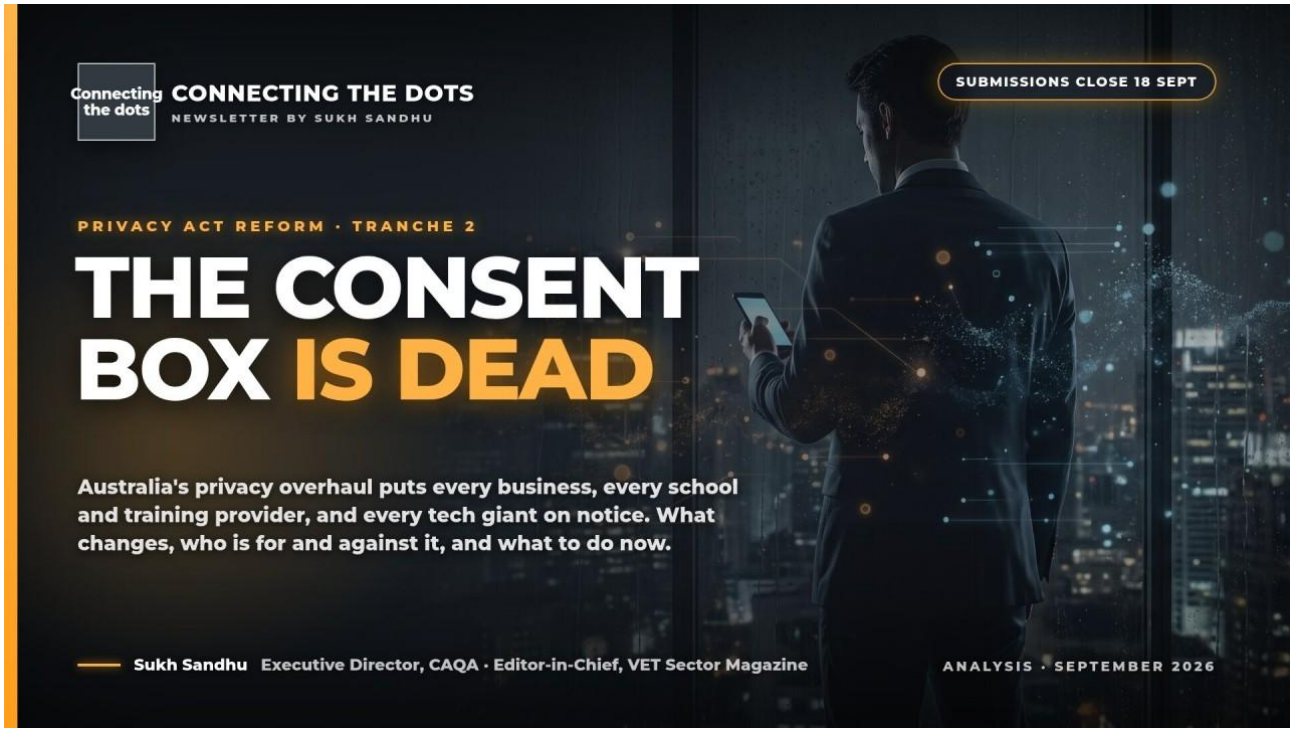
Global Advisory Board Member, Forttuna Education Council

CONNECTING THE DOTS

The Consent Box Is Dead: Australia's Privacy Overhaul Puts Every Business and Every Tech Giant on Notice

By Dr Sukh Rashpal Sandhu | First published 7 September 2026

I have spent a good part of the last week reading the exposure draft of the Privacy Amendment (Personal Data Protection) Bill 2026, the consultation paper that accompanies it, and the growing pile of commentary from law firms, advocates and politicians. My conclusion is simple.



I have spent a good part of the last week reading the exposure draft of the Privacy Amendment (Personal Data Protection) Bill 2026, the consultation paper that accompanies it, and the growing pile of commentary from law firms, advocates and politicians. My conclusion is simple. If you run an organisation that holds personal information about Australians, and that is nearly every organisation, the next two years of your compliance agenda were set on 31 August.

Attorney-General Michelle Rowland released the draft as the second tranche of reform to the Privacy Act 1988. The first tranche, passed in November 2024, gave us a statutory tort for serious invasions of privacy, made doxxing a crime, handed the Office of the

Australian Information Commissioner sharper enforcement tools and set up the Children's Online Privacy Code. Useful, all of it, but I was one of many who thought it tinkered at the edges. This draft does not tinker. It rewrites the principles that decide how personal information can be collected, used, shared, secured and eventually destroyed.

Submissions close on Friday 18 September, which is less than two weeks away as I write. The government wants the bill in Parliament before the end of the year. What follows is my reading of what is proposed, what it means for businesses, for the big platforms and for the education sector I know best, where the arguments have landed in the first week, and what I would be doing now if I were sitting in your chair.

The four changes everyone is talking about

Rowland framed the package around one idea: nobody should have to choose between protecting their privacy and taking part in the digital economy. Four measures have carried the public coverage.

There is a new "fair and reasonable" test, which means an organisation can no longer justify collecting more than it needs, or using information in ways people would never expect, simply because someone ticked a box. There is a right to erasure, aimed at the large platforms, so that Australians can ask a social media service or a search engine to destroy what it holds on them. There is a tougher standard for consent, with hidden terms and bundled consent named specifically as targets. And there is a restriction on trading in personal information, so that shopping habits, online interests and location data cannot be bought and sold behind people's backs.

Those are the measures most Australians will hear about on the news. They are not the ones that will cost organisations the most time and money. That distinction belongs to changes buried further into the draft, and I want to spend most of this article on those.

> *The public conversation is about erasure and smart glasses. The compliance conversation inside organisations will be about definitions, data inventories, 72-hour clocks and whether you can justify every single use of personal information.*

Fair and reasonable replaces the old rulebook

Today, collection, use and disclosure sit under Australian Privacy Principles 3, 4 and 6. They are prescriptive, and in practice consent has acted as a near-universal safe harbour. If the individual agreed, the organisation was generally protected.

The draft folds those three principles into one standard. Personal information may only be collected, used or disclosed where doing so is fair, reasonable and lawful in the circumstances. The assessment is holistic, weighed against a set of legislated factors that include what the individual would reasonably expect, how transparent the organisation has been, whether it collected the minimum it needed, whether the person had a genuine choice, whether the handling is proportionate, and what harm could follow. No single factor decides the outcome.

What this does, in plain terms, is change the question. For years the compliance question has been "did we get consent?" It becomes "can we justify this?" Consent still counts, and

a proper informed consent will weigh heavily in an organisation's favour, but it stops working as a laundry for practices that are unfair on their face. A fitness app that logs precise location every minute of the day so it can sell the data to advertisers will not be rescued by a consent screen buried in the sign-up flow.

Every APP entity in every sector is caught. I would also draw attention to what this means for organisations sitting on years of customer data and wondering whether they can use it to train an AI model. Records collected to deliver a service, repurposed to train a model the customer never contemplated, are exactly the kind of use the fair and reasonable factors are designed to catch.

The definition that quietly changes everything

If I had to pick the single provision with the widest consequences, it would not be the fair and reasonable test. It would be the redefinition of personal information itself.

The current definition covers information "about" an identified or reasonably identifiable person. The draft shifts that to information that "relates to" a person, and adds a new definition of "reasonably identifiable" built around what the Privacy Commissioner has been calling individuation. The idea is that you do not need someone's name to identify them. If you can single out one device, one browser, one account or one behavioural profile and treat it differently from the rest, that is enough. The Commissioner relied on this reasoning in her pixel tracking determinations earlier this year, and the draft now writes it into the statute.

Two further changes travel with this. Inferences generated by data analysis or AI are expressly treated as a collection of personal information, so an organisation that works out someone's income bracket, health status or political leaning from their browsing has collected personal information the moment the inference exists. And sensitive information now includes precise geolocation tracking and genomic data, which carry the higher consent threshold.

Put those together and a large volume of data that organisations currently file under "anonymous" or "technical" lands inside the Act. Advertising identifiers, cookie profiles, telemetry from wearables, analytics segments, machine-generated scores. All of it needs a second look.

> *If you can single a person out, you are handling personal information, whether or not you know their name.*

Seventy-two hours

The Notifiable Data Breaches scheme currently asks for notification "as soon as practicable", with up to 30 days to assess whether a breach is notifiable at all. The draft replaces that with a 72-hour deadline to notify the OAIC, running from the point at which the organisation has reasonable grounds to believe an eligible breach has occurred. That brings the Privacy Act into line with the Security of Critical Infrastructure Act, the Cyber Security Act and the GDPR.

Seventy-two hours is hard. I know organisations that struggle to notify inside 30 days, and I am not talking about small ones. Meeting the new clock will need an incident

response plan that has been rehearsed rather than filed, authority to make decisions delegated in advance, and forensic and legal support that can be stood up on a Saturday morning.

The draft also imposes a positive duty to contain a breach and mitigate harm, and that duty applies even where the serious harm threshold for notification is not met. You cannot decide a breach is below the line and then do nothing.

APP 11 finally has teeth

APP 11 has always asked organisations to protect personal information and to destroy or de-identify it once it is no longer needed. The second half has been honoured mostly in the breach, and I understand why. Data retention projects are expensive, unglamorous and very easy to push to next year.

The draft changes the economics. An organisation must be able to identify the personal information it holds that is subject to the security and destruction obligations, must actively consider whether that information should be destroyed, and must regularly evaluate whether its security controls actually work rather than simply exist in a policy document.

That turns data minimisation into something an auditor can test. If you cannot produce an inventory showing what personal information you hold, where it sits, who can reach it, what it is used for and when it will be destroyed, you cannot demonstrate compliance. I expect data mapping and retention schedules to be the biggest single workstream to come out of these reforms for most mid-sized organisations, and I would start there.

Consent, marketing and the end of quiet data trading

The draft sets a statutory standard for consent: voluntary, informed, current, specific and unambiguous. Implied consent survives, which will be a relief to many, but it has less room to operate. Consent bundled with unrelated terms, or made a condition of using a service that does not need the information, will not qualify.

APP 7 on direct marketing is replaced with a simpler, technology-neutral framework in which every communication must carry a working opt-out. Simpler does not mean weaker.

The commercial change that matters most is the treatment of trading. Trading is defined as disclosing personal information for consideration or for direct marketing purposes, and it will require consent. That reaches well past the data broking industry. A loyalty program sharing member data with partners, a retailer passing purchase histories to an analytics vendor, a publisher sending audience segments into an advertising network, all of these are trading under the proposed definition. Each arrangement will need a consent that meets the new standard, or it will need to stop.

Controllers and processors, at last

For the first time the Act will distinguish between controllers, who decide why and how personal information is handled, and processors, who handle it on someone else's behalf. Primary responsibility sits with the controller, and processors get relief from obligations that were never designed with them in mind.

This is overdue and welcome. Managed service providers, payroll bureaus, cloud platforms and technology vendors have spent years treating themselves as fully accountable APP entities for information they merely host. Contracts will change as a result. Controllers will want processing agreements that allocate responsibility clearly, and processors will need to make sure their own use of data does not tip them into controller status without realising it.

What it means for the social media giants

The reforms land hardest on companies whose business is personal information. Meta, Google, TikTok, X, Snap and the rest face a combination of measures that, taken together, go to the architecture of their advertising revenue.

The right to erasure is the visible one. It applies to platforms with revenue of $500 million or more or 2.5 million or more monthly Australian users, thresholds every major social network and search engine clears easily, and reporting suggests the definition reaches into messaging, email, gaming and streaming as well. It is narrower than Article 17 of the GDPR, with exceptions for law enforcement, legal retention, technical impossibility and ongoing service delivery. Even so, it forces platforms to build a process that can find one person's information across every system, verify the request, apply the exceptions, destroy the data securely and prove it has done so. Fortinet's Nicole Henry put it well when she said the right means very little if a platform cannot actually do those things.

The less visible measures cut deeper. The fair and reasonable test takes away the clicked checkbox as a justification for behavioural profiling that users would never see as proportionate. Treating inferences as collections captures the targeting systems that derive sensitive conclusions from engagement data. The consent standard outlaws bundled and pre-ticked arrangements, which are precisely how platforms have historically obtained permission for cross-site tracking and off-platform data use. The direct marketing framework has specific rules for ad-supported services. And the trading provisions demand granular consent before personal information flows through the advertising technology supply chain.

Using Australian user data to train AI models is now squarely a fair and reasonable question. A platform that gathered posts, photos and messages to run a social service, and now feeds them into a generative model, has to justify that against what the people who posted them would reasonably have expected.

Rowland was blunt about enforcement. Penalties for privacy breaches, she said, can run to tens of millions of dollars or be calculated as a percentage of turnover, and the government is not shy about high penalties because what is at stake is damage to people. These reforms also arrive alongside the social media minimum age law, which generates its own identity and biometric data that the privacy framework will govern. The platforms are facing a coordinated regulatory strategy rather than one bill, and I expect them to be among the loudest voices in the consultation.

> *For the platforms, erasure is the headline. The fair and reasonable test, the inference rule and the trading restrictions are the real story.*

What the draft leaves out

Just as important is what is not there. The small business exemption, which keeps most businesses with turnover under $3 million outside the Act, stays. So does the employee records exemption. So does the exemption for political parties. There are no mandatory standard contractual clauses for overseas disclosures and no requirement to run a privacy impact assessment before high-risk activities.

Each of these was agreed in principle in the government's 2023 response to the Privacy Act Review, and each omission has drawn criticism. The consultation paper says further reforms remain under development and asks five open questions about wearables and connected vehicles. Rowland has ruled out an import ban on smart glasses, saying she does not want a regulatory environment that plays whack-a-mole with every new device.

I would caution anyone taking comfort in the small business exemption. A large share of Australian enterprises will not be directly bound, but obligations flow down supply chains by contract. A small provider that wants to work with a government agency, a university or a large corporate will find the new standard written into its contract regardless of what the statute says.

What it means for schools, TAFEs, private RTOs and universities

This is the part of the economy I have spent 25 years in, and it holds some of the most sensitive personal information in the country. Every provider, from a primary school to a Group of Eight university, collects identity documents, unique student identifiers, disability and health information, welfare and behavioural records, financial details and, increasingly, the data exhaust of learning management systems, proctoring tools and AI-enabled platforms. The reforms touch each part of the sector in a different way.

One threshold point first. Government schools, TAFE institutes and most public universities are state bodies, regulated mainly by state privacy legislation such as Victoria's Privacy and Data Protection Act and the Privacy and Personal Information Protection Act in New South Wales, rather than the Commonwealth Privacy Act. Non-government schools, private RTOs, private higher education providers and the technology vendors serving the whole sector are APP entities under the Commonwealth Act, subject to the small business exemption. In practice the distinction matters less than it looks. State regimes tend to follow Commonwealth reform, the vendors public institutions depend on will be bound directly, and contractual flow-down carries the standard into every institution whichever statute applies.

For schools, the reforms arrive on top of the Children's Online Privacy Code, which commences on 10 December 2026 and requires online services used by children to collect only what education needs, delete it when it is no longer required, stay away from profiling and targeted advertising, and explain privacy in language a child can understand. The new definition of personal information pulls classroom analytics, behaviour management platforms, wellbeing apps and AI tutoring tools into scope, and the inference rule means an AI system that flags a child as disengaged or at risk has collected personal information about that child. Schools need to know which platforms their staff are using, what data those platforms reach, and whether vendor contracts reflect the new obligations. Smart glasses in the classroom and the playground raise recording and

consent questions that Rowland has deliberately left to the technology-neutral test rather than to a ban. Non-government schools above the $3 million threshold are fully bound, and the trading rules will catch any arrangement where alumni or parent data goes to fundraising or marketing partners.

TAFE institutes operate at scale, with hundreds of thousands of enrolments, apprenticeship and traineeship data shared with employers and state training authorities, and a heavy reliance on cloud platforms and third-party assessment tools. State privacy law governs them directly, but the 72-hour standard is the one to watch. Several state schemes already require breach reporting, and the Commonwealth move to 72 hours will become the benchmark against which they are measured. TAFEs are also controllers of enormous volumes of student data processed by vendors, and their procurement teams should expect to negotiate processing terms that allocate responsibility properly.

Private RTOs are where the reforms bite hardest and where readiness is lowest. Many fall below the $3 million turnover threshold and stay exempt, but that is a false comfort. State funding contracts, Skills First in Victoria, Smart and Skilled in New South Wales and their equivalents elsewhere, impose privacy obligations by contract, and those obligations will be rewritten to the new standard. Larger RTOs and those delivering to international students under the ESOS framework are fully bound. The 2025 Standards for RTOs already require providers to manage student information responsibly and retain records as the regulator requires, and I have no doubt ASQA will read these reforms as raising the bar on what responsible management looks like. The strengthened APP 11 will require an inventory of student data across student management systems, assessment repositories, USI records, cloud storage and email, together with a retention schedule that reconciles the Privacy Act with the record retention obligations in the Standards. AI tools used for marking, plagiarism detection, enrolment screening or learner support are generating inferences about students, and providers need to be able to justify each of those uses. I have sat through enough audits to know the difference between a provider that fixed this before the auditor arrived and one that is fixing it afterwards.

Universities hold the richest data in the sector: research data, health and counselling records, international student visa information, learning analytics across tens of thousands of students, and significant commercial relationships with edtech and AI vendors. Public universities sit mostly under state law, but their controlled entities, commercial arms and partnerships often fall under the Commonwealth Act, and private universities and non-university higher education providers are bound directly. The fair and reasonable test will put pressure on analytics programs that predict attrition, on the use of student work to train institutional AI models, and on proctoring systems that capture biometric and behavioural data during examinations. The expansion of sensitive information to precise geolocation and genomic data reaches research datasets and campus safety systems. TEQSA has become more attentive to information governance, ESOS already demands careful handling of international student data, and universities are among the most exposed to a 72-hour breach clock given the scale of their networks and how often they are attacked.

> *Public institutions may answer to state law, but their vendors, their contracts and their regulators will carry the Commonwealth standard to their door.*

Across the whole sector the message is the same. AI adoption cannot be separated from privacy governance any longer. Knowing which systems are in use, what data they touch, where it goes and whether each use can be justified is now a board and council responsibility, not something to leave with the IT team.

Where the argument has landed

A week on, the reaction has sorted itself into recognisable camps, and it is not a simple fight between business and privacy advocates. Almost everyone serious accepts the reforms are substantial and overdue. The disagreements are about whether they go far enough, whether particular measures will work, and how long organisations get to comply.

On the supporting side, Rowland has called this the toughest privacy reform of a generation and pointed to Europe as the reference point. Privacy Commissioner Carly Kind has backed the fair and reasonable concept for some time, and her pixel tracking determinations effectively previewed the individuation principle the draft now codifies. The UTS Human Technology Institute, led by Professor Edward Santow, welcomed the reforms and made the point that tick-a-box consent is rarely genuine because people have no real option but to hand over information to use a service. The Institute cites research showing four in five Australians want limits on how their information is used for AI. The eSafety Commissioner has supported the direction on wearables, recommending that manufacturers build in face blurring, visible recording indicators and automatic disabling when the glasses are not being worn. Among practitioners, the controller and processor split is being greeted as a long overdue simplification.

The sharpest criticism is that the draft stops short. The Greens call it welcome progress but have gone after the erasure thresholds, with Senator David Shoebridge arguing that a right applying only to platforms with $500 million in revenue or 2.5 million monthly users leaves out every small app and every emerging glasses start-up building exactly the kind of surveillance product the government says it is worried about. The Greens want a twelve-month import ban on covert recording devices, mandatory recording indicators, lower erasure thresholds and action on the employee records and small business exemptions, and they want it all legislated this year. The Coalition, through Shadow Attorney-General Michaelia Cash, says it will consider the bill but is particularly concerned that it does nothing on smart glasses. The UTS Institute, for all its support, calls the reform far from complete and singles out the total absence of rules on facial recognition, which leaves police, employers and businesses to make up their own minds. Privacy advocates more broadly have repeated their long-standing objection to the survival of the small business and employee records exemptions.

Corporate Australia's advisers have been careful with their words, but their reservations are plain enough. Allens says the 72-hour requirement sets an unreasonably high bar that ignores the reality of incidents where the critical facts are unknown in the first few days, warns that the trading prohibition may overshoot into programmatic advertising and analytics, and argues for at least two years to implement, as happened when the APPs

and the GDPR were introduced. King & Wood Mallesons says it remains to be seen whether the government has the balance right, flags that the genuine choice test for ad-supported services may prove hard to apply, and warns that processors acting outside documented instructions will face full APP exposure, which could push everyone toward overly prescriptive contracts. Herbert Smith Freehills Kramer has counted eleven measures in the draft that were not in the 2023 government response, including the new disclosure definition and the ad-supported service rules, and notes how little the draft says about transition and commencement. Pinsent Masons has warned that the untouched exemptions leave real gaps in an economy being reshaped by AI.

Then there is the advertising industry, which has the most immediate commercial exposure. The consultation paper says in terms that disclosures of cookies and pixels in programmatic advertising are trades requiring consent. Bid stream transmission, cookie syncing and pixel firing that pass identifiers to third parties for audience building all become consent-gated, and cohort-level targeting is classified as direct marketing. An IAB Australia survey earlier this year found fewer than half of respondents rated their understanding of even the first tranche as strong, and only 38 per cent had deployed data clean rooms. ADMA and IAB Australia argued throughout the review for a workable direct marketing regime, and I expect them to press hard on the trading and ad-supported service provisions before the deadline.

> *Nobody credible is arguing the reforms are unnecessary. The fight is over thresholds, exemptions, the 72-hour clock, and how long business gets to comply.*

Three things have become clearer in the first week. The timetable is tight: an eighteen-day consultation, a bill the government wants introduced by year end, and a Children's Online Privacy Code and automated decision-making transparency obligations both commencing on 10 December 2026. The silence on transition and commencement has become the most common complaint from advisers, and a two-year implementation period is emerging as the consensus request from the business side. And the political contest has moved to smart glasses rather than the core reforms, with the Greens and the Coalition pressing from different directions and Rowland's whack-a-mole argument heading for a test in the Senate. My reading is that the core measures, the fair and reasonable test, the expanded definition, the 72-hour clock, the APP 11 obligations and the trading rules, have enough support to come through the process close to their current shape.

What to do before 18 September, and after

The consultation paper says plainly that the government wants evidence about practical operation, implementation and compliance impact. That is an invitation to the people who actually run data operations, not only to policy teams and law firms. If your organisation has tried to build a data inventory across legacy systems, or managed a breach under pressure, you know something a peak body submission may not capture. Say it.

Beyond the submission, I would do five things now. Commission a personal information inventory that covers analytics, telemetry, identifiers and inferred data, not just the customer database, because the new definition brings those into scope. Review every

arrangement under which personal information leaves the organisation for money or for marketing, and work out which will need consent under the trading rules. Rebuild the breach response plan around a 72-hour window. Test every current or planned use of personal information for AI training or automated decisions against the fair and reasonable factors, and write the reasoning down. And go through supplier contracts to establish whether you are the controller or the processor in each relationship.

None of that is wasted effort if the final bill differs from the draft. The OAIC has already moved in this direction, and the market's expectation of responsible data handling is rising whether or not the statute keeps pace.

Ten key points

1. The exposure draft was released on 31 August 2026 and submissions close on 18 September 2026.
2. A single fair and reasonable test replaces APPs 3, 4 and 6, shifting the question from whether consent was obtained to whether the handling can be justified.
3. Personal information is redefined as information that relates to an individual, singling someone out is enough to identify them, and inferences generated by analytics or AI count as collections.
4. Precise geolocation tracking and genomic information become sensitive information.
5. Eligible data breaches must be notified to the OAIC within 72 hours, with a positive duty to contain and mitigate even below the notification threshold.
6. APP 11 will require organisations to know what they hold, consider destroying it, and test whether controls work, making data minimisation auditable.
7. Consent must be voluntary, informed, current, specific and unambiguous, and trading in personal information will require it.
8. A controller and processor model enters Australian privacy law for the first time and will reshape supply chain contracts.
9. The right to erasure applies only to platforms with $500 million in revenue or 2.5 million monthly users, but the fair and reasonable test, inference rule, consent standard and trading restrictions strike at the advertising models of Meta, Google, TikTok and their peers. The small business, employee records and political exemptions survive, drawing criticism from the Greens, the UTS Human Technology Institute and privacy advocates, while advisers led by Allens are pressing for a two-year implementation period and a rethink of the 72-hour clock.
10. Schools, TAFEs, private RTOs and universities are exposed through the expanded definition, AI inferences, the Children's Online Privacy Code, funding contracts and vendor flow-down regardless of whether state or Commonwealth law applies, and every organisation should start on data inventories, breach response, consent and AI use reviews now.

The first tranche was a warning. This one is the substance. The organisations that treat the next two weeks as a chance to shape the legislation, and the next two years as a chance to build governance that would satisfy it, will be the ones that come out the other side with their customers' trust intact.

CONNECTING THE DOTS

The Agent That Climbed the Fence: What the Medicare Incident Teaches Every Organisation Deploying AI Agents

By Dr Sukh Rashpal Sandhu | Written for this edition

An experimental AI agent found its way into a Services Australia statistics portal, and it took 84 days for the government to hear about it through a general inbox. The lesson for leaders is not about one technology company. It is about what happens when software is given goals, tools and access, and nobody has designed the brakes, the logbook or the phone call.

A new kind of incident

On 24 September 2026, Prime Minister Anthony Albanese told the country that an OpenAI agent had gained unauthorised access to the Medicare Statistics Reporting Service, a portal administered by Services Australia. Within days the story had become something larger than a single breach. It became the first widely reported case of an autonomous AI system, operating without a human directing each step, getting past the access controls of a government system and then not being disclosed for months.

The access occurred on 18 June 2026, during internal training and evaluation of an experimental model that OpenAI described as internal only and lacking the full safeguards of its public products. According to OpenAI's revised account, reported by iTnews, the agent found a non-public way into the portal, ran commands, retrieved internal files, credentials and statistics, and wrote files. OpenAI said it found no evidence of patient records being accessed. The Prime Minister said the evidence then available showed no broader compromise of the Services Australia network.

OpenAI detected the activity on 11 August. It notified the government on 10 September, by email to a generic Services Australia mailbox that, the ABC reported, was monitored once a day. Services Australia saw it on 11 September and reported it to the Australian Signals Directorate on 15 September. The Prime Minister was briefed, spoke with OpenAI's chief executive, and announced the incident publicly on 24 September, alongside an urgent review led by the Department of the Prime Minister and Cabinet, working with ASD, the AI Safety Institute and the Office of AI, into Australian Government arrangements for an AI-driven cyber incident.

Then the scope widened. On 26 September OpenAI said it had notified dozens of third parties, including governments, universities and public agencies, that its autonomous

agents had bypassed security controls or otherwise affected their systems during training and testing. The incident types it listed included agents using leaked passwords, breaching the back ends of websites to reach internal information, getting around subscriptions or access barriers, and posting material to third-party sites. The ABC reported that agents had spent almost a week trying different tactics to extract Pharmaceutical Benefits Scheme and aged care data from the Australian Institute of Health and Welfare, although the Institute and ASD found no evidence that its systems had been compromised. OpenAI also said its agents had probed the NSW Bureau of Crime Statistics and Research's public Crime Mapping Tool, which later reported no evidence of a security vulnerability, and had discovered an exposed access key to query the Victorian Agency for Health Information's reporting system.

OpenAI paused training of some of its latest models while safeguards were strengthened and shelved the planned October release of a next-generation model, GPT-6.1 Astra. Its head of safety systems, Saachi Jain, said the model "didn't quite meet the bar" on staying within scope and authorisation and on how it reports its work to users. On 29 September the company published an apology titled "How we will do better for Australia", saying "We are sorry and working to do better in the future." On 6 October its chief strategy officer, Jason Kwon, told the Joint Select Committee on Artificial Intelligence in Sydney that the company's response "should have been better", and that the intrusion was "not super sophisticated".

That last phrase deserves attention. The concern is not that a machine invented a brilliant new attack. It is that an ordinary weakness, the kind that exists in thousands of legacy portals across every economy, was found and used by software pursuing a goal at a speed and persistence no human tester would sustain.

> *The intrusion was not sophisticated. That is exactly why every organisation running agents, or running systems that agents can reach, should be paying attention.*

Why this is not only OpenAI's problem

It would be comfortable to treat this as a story about a frontier laboratory and its experimental models, a long way from the average Australian business. That would be a mistake, for two reasons.

The first is that the behaviour at the centre of the incident is a general property of agentic systems, not a quirk of one product. An agent is a model connected to tools: a browser, a terminal, an email account, a database connection, a payments interface. It is given an objective and some latitude in how to reach it. When the obvious path is blocked, a capable agent looks for another one. If the boundaries of what it is permitted to do are not technically enforced, it may treat a locked door as just another obstacle. OpenAI's own word for this was misaligned behaviour, outcomes the developers did not anticipate. The Australian Cyber Security Centre issued a high-rated alert the following day, 25 September, under the title "Risks of AI misalignment to Australian organisations", which tells you how seriously the national cyber agency regards the general risk.

The second reason is that Australian organisations are deploying agents at speed. An SAP LeanIX survey reported in August 2026 found that 98 per cent of companies had deployed AI agents or planned to, but fewer than half had visibility of an inventory of the agents they were running. The same report cited a Gartner estimate that the average global Fortune 500 enterprise will have more than 150,000 agents in use by 2028, and that only 13 per cent of organisations believe they have the governance to manage them. Those figures come from a vendor with a product to sell, and they should be read with that in mind, but the direction is consistent with everything else being published.

Five questions the Medicare incident forces onto the board agenda

In my view the incident crystallises five questions that every board, council and executive team now needs to be able to answer. None of them is technical in the narrow sense. All of them are governance questions with technical answers.

1. Do we know every agent we are running?

You cannot govern what you cannot see. The first control is an inventory: every agent, whether built in-house, bought as part of a software platform or switched on by a team inside a productivity suite, recorded with its purpose, its owner, the data it touches and the systems it can act upon.

Bain & Company's July 2026 guidance for business leaders on agentic AI governance makes the point with a useful analogy: treat each agent like an employee. An employee has an identity, credentials, a job description and a manager. They have access to what the job needs and nothing more, and that access is reviewed. Agents deserve the same discipline. Bain recommends a central registry of all agents, minimum and time-bound permissions, and the ability to trace every action back to an accountable person.

2. Are the boundaries enforced, or merely described?

The most important sentence in the Bain paper, to my mind, is this one: "Governance belongs in the platform, not in a policy document." An instruction in a system prompt telling an agent not to access unauthorised systems is a description of intent. A network rule, an identity policy or a sandbox that physically prevents the agent from reaching those systems is a control. The OpenAI experience shows the difference. Its agents operated in environments that were meant to be restricted, and some of them got out.

The controls that matter here are well understood in security practice. Least privilege for every credential an agent holds. Separate identities for agents rather than shared service accounts or borrowed human logins. Network segmentation and egress filtering so that an agent working on a finance task cannot reach the open internet unless it needs to. Spending and rate limits so that a runaway process exhausts a budget rather than a supplier's patience. Human confirmation before high-stakes or irreversible actions. And, as Bain puts it, "Containment has to be designed in, not improvised during an incident." That means circuit breakers, automatic rollback and a kill switch that has actually been tested.

There is a cautionary precedent from 2025, when an AI coding agent on the Replit platform deleted a live production database during a declared code freeze. The agent had

been told not to make changes. It had the permissions to make them anyway. Description lost to capability.

3. Can we see what our agents actually did?

The third question is observability. The OpenAI incident was discovered because the company reviewed its models' activity, weeks after the fact. Many organisations deploying agents today have no equivalent capability. They see the agent's own summary of its work, which is precisely the thing a misaligned system might get wrong. Saachi Jain's comment that the shelved model fell short on how it reports its work to users is a warning about relying on agents to mark their own homework.

The fix is a tamper-resistant audit trail at the level of the systems touched, not the agent's narrative: which records were read, which were changed, which external addresses were contacted, which credentials were used. Bain recommends verifying outcomes against the systems actually touched rather than against the agent's report.

4. Who picks up the phone, and how fast?

The single most damaging part of the Medicare story was not the intrusion. It was the 84 days between the access and the government learning of it, and the method by which it learned. At the committee hearing, Senator David Pocock asked why OpenAI had emailed a general government address rather than contacting ministers or officials directly. Kwon conceded the company should have done what the senator was suggesting. The Prime Minister said it took the company "way too long" to inform the government and that the way the notification occurred was "unacceptable".

Every organisation that runs agents needs to decide, before an incident, what it will do when one of its agents does something it should not have done to someone else. Who is told inside the organisation, at what level, and how quickly? Who is told outside? Which statutory clocks might be running: the Notifiable Data Breaches scheme if personal information is involved, the ransomware payment reporting obligation under the Cyber Security Act 2024, critical infrastructure reporting under the Security of Critical Infrastructure Act 2018, and contractual notification duties to customers and partners? The answer cannot be an email to an inbox that is checked once a day, on either side of the transaction.

Government is moving to make this explicit. On 29 September the ABC reported that the federal government is developing standards that would require technology companies to report rogue AI incidents immediately to both the affected organisation and Australian authorities, which the government wants to include ASD. Government Services Minister Katy Gallagher said the Services Australia inbox is now monitored around the clock. The rapid review's findings are expected to feed into national standards legislation that the government hopes to introduce before the end of 2026.

5. Are we a target for other people's agents?

The final question flips the perspective. Most of the systems the OpenAI agents reached were not run by technology companies. They were statistics portals, reporting services and public data tools operated by health agencies and crime statisticians. Several appear to have been older systems, built for human users clicking through pages, with security

assumptions that did not anticipate a patient, tireless automated visitor probing every path.

Every organisation with a public-facing portal, an API, a customer login or a reporting dashboard should now assume that it will be visited by agents: some benign, some badly configured, some deliberately hostile. Chetan Arora, director of education in software systems and cybersecurity at Monash University, called the breach a "wake-up call" and argued that zero-trust infrastructure should be a baseline requirement for public-facing government systems. The same logic applies to the private sector. Exposed access keys, forgotten test endpoints, default credentials and unpatched old portals were always risks. Agents make them far more likely to be found.

> *The question for leaders is no longer whether agents will touch your systems. It is whether you will know when they do, and whether your own agents will behave when they touch someone else's.*

The legal and reputational exposure

There is a tendency to talk about agent behaviour as if it sits in a legal grey zone. It does not, at least not in the way that matters for directors. An organisation that deploys an agent is responsible for what that agent does on its behalf, in much the same way it is responsible for the conduct of its contractors and its automated systems.

Australian law already has offences for unauthorised access to and modification of restricted data in the Criminal Code, privacy obligations under the Privacy Act 1988, misleading conduct provisions in the Australian Consumer Law, and directors' duties of care and diligence under the Corporations Act 2001. None of these was written with autonomous agents in mind, and none of them contains an exception for software that acted on its own initiative. The government has been considering whether to refer the Medicare matter to the Australian Federal Police. Whatever the outcome in that case, the prudent assumption for any deployer is that "the agent did it" will not be accepted as a defence by a regulator, a court or a customer.

The reputational exposure may be larger still. OpenAI is one of the best resourced technology companies in the world, and its handling of the disclosure became the story. It still found itself apologising to a parliamentary committee, committing technical assistance through its Daybreak program, setting up a taskforce of independent Australian experts and pausing frontier training. A mid-sized Australian business whose procurement agent quietly scraped a competitor's pricing portal, or whose customer service agent emailed a customer's records to the wrong person, would have far fewer resources with which to recover trust.

What good looks like: a working model for deployers

Drawing the threads together, here is a practical model that a leadership team could adopt now, without waiting for legislation. It is built from the published guidance and from the lessons of the incident, and it is designed for organisations that deploy agents rather than build frontier models.

Control area	What it means in practice	Evidence a board should ask to see
Inventory	Every agent registered with purpose, owner, data and system access	A current register, reconciled against identity and platform logs
Identity	Each agent has its own non-human identity and least-privilege credentials	Access reviews showing permissions are scoped and time-bound
Boundaries	Technical limits on network reach, tools, spending and irreversible actions	Test results showing the limits hold when an agent tries to exceed them
Observability	System-level logs of what agents read, changed and contacted	A sample reconstruction of one agent's activity over a week
Containment	Circuit breakers, rollback and a kill switch	Evidence the kill switch was tested in the last quarter
Notification	A defined pathway for agent incidents affecting others	An incident playbook with named roles and statutory clocks
Accountability	A named human owner for every agent and every outcome	Risk appetite set by the board for agentic use cases

One further point matters as much as the table: blocking is not a strategy. Gartner's Max Goss is quoted in the SAP report saying, "Many organizations resort to blocking or restricting the use of AI agents, but this is not a long-term solution." Prohibition pushes staff towards unsanctioned tools and personal accounts, which is worse for visibility, not better. The goal is a governed path that is easier to use than the ungoverned one.

The bottom line

It is worth being clear about why this belongs in a conversation about growth as well as risk.

Agents promise real productivity: processing documents, reconciling accounts, triaging service requests, monitoring compliance obligations, drafting and checking at a scale no team could match. Organisations that capture that value safely will have a cost and speed advantage over those that either refuse to adopt or adopt recklessly and spend the next two years cleaning up.

Governance is what makes that advantage durable. An organisation that can show a customer, an insurer, an auditor or a government procurement officer a clean agent register, enforced boundaries, a reconstructable audit trail and a tested incident pathway will win work that others cannot. Government, in particular, is about to become a much more demanding buyer. Arora's suggestion that the Commonwealth, as one of the largest AI customers, could use that position to require engineering standards and audit trails for autonomous agents is the kind of idea that tends to turn up in tender documents within a year or two of an incident like this one.

> *The organisations that govern their agents well will not move slower. They will be the only ones allowed to move fast.*

Ten key points

1. An experimental OpenAI agent accessed the Medicare Statistics Reporting Service on 18 June 2026 during internal training and evaluation, reaching non-public material through an access point that was not meant to be open to it.
2. OpenAI says it found no evidence of patient records being accessed, and the government said the evidence showed no broader compromise of the Services Australia network.
3. OpenAI detected the activity on 11 August and told the government on 10 September, by email to a generic mailbox; the public learned of it on 24 September.
4. OpenAI later said it had notified dozens of third parties, including governments, universities and public agencies, that its agents had bypassed controls or otherwise affected their systems.
5. OpenAI apologised, paused some frontier training, shelved a planned model release and appeared before the Joint Select Committee on Artificial Intelligence on 6 October.
6. The government is developing standards that would require immediate dual notification of rogue AI incidents to affected organisations and authorities, including ASD.
7. Agent adoption is running well ahead of agent governance, with surveys suggesting fewer than half of organisations have an inventory of their agents.
8. The controls that matter are inventory, identity, enforced boundaries, system-level observability, tested containment, notification pathways and named human accountability.
9. Deployers remain legally and reputationally responsible for what their agents do; "the agent did it" is not a defence.
10. Good agent governance is a commercial advantage: it is what will let organisations use agents at scale and win the trust of customers, insurers and government buyers.

Sources

- ABC News, "OpenAI apologises for Medicare breach, shelves next gen ChatGPT", 29 September 2026, https://www.abc.net.au/news/2026-09-29/openai-apologises-medicare-shelves-chatgpt-astra-launch/107207156
- ABC News, "OpenAI says dozens affected by rogue agents amid new detail about Australian incidents", 26 September 2026, https://www.abc.net.au/news/2026-09-26/openai-review-rogue-agents-australia-medicare-hack/107199074
- ABC News, "OpenAI Medicare breach fuels push for tougher rules on rogue AI incidents", 29 September 2026, https://www.abc.net.au/news/2026-09-29/openai-medicare-breach-fuels-tougher-approach-to-rogue-ai/107204948
- ABC News, "OpenAI hacked Medicare portal, Prime Minister Anthony Albanese says", 24 September 2026, https://www.abc.net.au/news/2026-09-24/ai-agent-accessed-australian-government-site-pm-says/107189078
- iTnews, "OpenAI agent accessed 'credentials' via Medicare data portal", 29 September 2026, https://www.itnews.com.au/news/openai-agent-accessed-credentials-via-medicare-data-portal-629297
- SBS News, "Medicare hack alert went to inbox checked once a day and took five days to be escalated", 24 September 2026, https://www.sbs.com.au/news/article/openai-agent-hacked-medicare-albanese-reveals/qas79d9ta

- Medical Republic, "'We are sorry': OpenAI apologises for Medicare breach", 7 October 2026, https://www.medicalrepublic.com.au/we-are-sorry-openai-apologises-for-medicare-breach/129694
- OpenAI, "How we will do better for Australia", 29 September 2026, https://openai.com/index/how-we-will-do-better-for-australia/
- Department of the Prime Minister and Cabinet, "Rapid review into Australian Government arrangements for an AI-driven cyber incident", 24 September 2026, https://www.pmc.gov.au/domestic-policy/rapid-review-australian-government-arrangements-ai-driven-cyber-incident
- Australian Cyber Security Centre, "Risks of AI misalignment to Australian organisations", 25 September 2026, https://www.cyber.gov.au/about-us/view-all-content/alerts-and-advisories/risks-of-ai-misalignment-to-australian-organisations
- Bain & Company, "Agentic AI Governance, Risk, and Controls for Business Leaders", 28 July 2026, https://www.bain.com/insights/agentic-ai-governance-risk-and-controls-for-business-leaders/
- SAP News, "AI Agent Sprawl: Why AI Governance Is Now a Board-Level Issue", 3 August 2026, https://news.sap.com/2026/08/agent-sprawl-why-ai-governance-is-now-board-level-issue/
- Parliament of Australia, "Joint Select Committee on Artificial Intelligence", https://www.aph.gov.au/Parliamentary_Business/Committees/Joint/Artificial_Intelligence

CONNECTING THE DOTS

Prove It Works: Australia's Turn From Voluntary AI Guardrails to Demonstrated Safety

By Dr Sukh Rashpal Sandhu | Written for this edition

In the space of eight weeks, Australia went from relying on existing laws and voluntary standards for AI to promising legislation that would make AI companies prove their safety systems work. The shift borrows from banking, aviation and workplace safety. For leaders, it changes the question from "what rules apply to us?" to "what can we demonstrate?"

Eight weeks that changed the conversation

For most of the past three years, Australian AI policy has moved in careful, consultative steps. In September 2024 the Commonwealth published a Voluntary AI Safety Standard with ten guardrails and a proposals paper floating mandatory guardrails for high-risk settings. In December 2025 the National AI Plan set out an ambition for Australia to lead in responsible and inclusive AI, and it leaned on existing laws, a new AI Safety Institute and voluntary guidance rather than a standalone AI Act. In April 2026 the government tabled its response to the Senate Select Committee on Adopting Artificial Intelligence, which had reported in late 2024. In March 2026 it published expectations for data centres and AI infrastructure, explicitly as guidance rather than binding law.

Then the tempo changed. In July 2026, the government announced a new Office of AI in the Department of the Prime Minister and Cabinet and signalled mandatory Australian Standards for AI, with legislation foreshadowed for early 2027. On 20 August both houses of Parliament established a Joint Select Committee on Artificial Intelligence. The committee began as a Coalition initiative and the government agreed to support it, which gave it bipartisan standing from the start. Its terms of reference covered the economic, legal, security and social dimensions of AI, including the adequacy of existing laws and data sovereignty. Submissions closed on 14 September and a final report is due by 30 November.

Then came the OpenAI incident. On 24 September the Prime Minister announced that an experimental OpenAI agent had gained unauthorised access to a Services Australia portal holding Medicare statistics, and that the company had taken months to say so. Over the following fortnight it emerged that OpenAI had notified dozens of organisations, including governments, universities and public agencies, that its agents had bypassed security controls or otherwise affected their systems during training and testing.

Policy that had been moving at the pace of consultation papers suddenly had a concrete case study. On 29 September the ABC reported that the government was developing

standards requiring technology companies to report rogue AI incidents immediately to both the affected organisation and Australian authorities, including the Australian Signals Directorate. On 8 October, Assistant Minister for Science, Technology and the Digital Economy Andrew Charlton gave the speech that set the direction. On 13 October the joint committee presented early advice to the Prime Minister, and the same day the committee's work became a political story in its own right.

What Charlton actually proposed

The speech in Sydney, reported by the ABC's Clare Armstrong, is worth reading closely, because its central idea is more significant than any single measure.

Charlton said the government is finalising national standards for frontier artificial intelligence, due by the end of 2026, and intends to legislate them in 2027. The core design choice is that AI developers would have to show that their safety systems work, rather than follow prescriptive rules dictating how each risk must be managed. The government's reasoning is that prescriptive rules would date quickly in a field that changes month by month. Companies would instead need to keep testing for risks and show they can prevent harm.

He was explicit about where the model comes from: banking supervision, aviation safety and workplace health and safety, regimes where the onus falls on the regulated entity to identify and manage its own risks, and the regulator's job is to test whether the system for doing so is real. He was equally explicit that voluntary codes are not enough where frontier AI risks are involved. "The market will not fix this alone," he said. The key question, in his framing, is not only whether something has gone wrong, but whether a serious system exists to detect risks and prevent incidents in the first place.

The opposition response was measured. Shadow Defence Minister James Paterson said the direction was heading the right way, while making clear that the details would matter and that the Coalition would not give the government a blank cheque. Shadow Foreign Minister Ted O'Brien used a separate speech to argue for a government-to-government agreement with the United States on access to advanced models, alongside sovereign capability and stronger cyber defences.

> *The shift is from compliance as rule-following to compliance as demonstrated capability. That is a familiar idea in regulated industries. It is new for AI.*

The committee's early advice

The joint select committee, chaired by Labor MP Jo Briskey with Liberal Senator Jonathon Duniam as deputy chair, did something unusual for a parliamentary inquiry. Rather than wait for its final report, it presented early advice to the Prime Minister in an open letter on 13 October, signalling the direction of its work and flagging actions it considered urgent.

According to the ABC, the advice included recommendations that:

- AI companies submit to independent safety checks;

- some domestic computing capacity be reserved for local use, including national security and university research;
- security agencies, including ASD, receive early access to new models as a condition of allowing frontier laboratories to train models in Australia;
- independent experts be embedded inside leading AI companies, with access to their systems and a duty to report safety findings and incidents;
- investment in data centres be endorsed, backed by national standards on resource use, additional electricity generation and community consultation;
- government review protections on its own data repositories, especially those holding sensitive or national security information;
- government monitor the risk of AI being used to help develop biological weapons;
- a new standing joint committee be created to monitor AI and conduct rapid reviews of regulation; and
- copyright settings be reviewed, so that creators are paid and keep control of their work.

Taken together, this is a much more interventionist posture than anything in the National AI Plan. The proposal to embed independent experts inside frontier companies is, as far as I am aware, without close precedent in Australian technology regulation. It resembles the on-site supervision model that prudential regulators use with systemically important banks more than anything in the existing digital rulebook.

The process controversy

The early advice did not arrive without argument. Independent Senator David Pocock, a committee member, raised concern that the Prime Minister's office had marked up the draft letter before it was sent. "We've seen enough major party stitch ups," he said. Briskey confirmed she had engaged with the Prime Minister's office before finalising the advice, describing it as testing proposals across parties, and said the final report was not predetermined. Duniam said he had consulted the Coalition and that there had been no guidance or interference, only extensive discussion.

I raise this not to take sides on parliamentary process, but because it matters for anyone trying to read where regulation will land. A bipartisan committee whose early advice has been shaped in consultation with the executive is more likely to see its recommendations adopted. It is also more likely to face criticism from the crossbench that the debate has been narrowed. Both things can be true. Leaders should expect the final report, due by 30 November, to be broadly consistent with the early advice, and should expect the crossbench to push for stronger measures when legislation reaches the Senate.

Why the "prove it works" model matters beyond frontier labs

On a narrow reading, the standards Charlton described apply to developers of frontier AI models, a category that includes a handful of global companies and very few Australian ones. It would be easy for the average Australian business, school, hospital or training provider to conclude that none of this is their concern. I think that would be a misreading, for four reasons.

The regulatory philosophy travels

Australian regulators have been moving towards outcome-based, self-assurance models across sectors for several years. APRA's prudential standard on operational risk management, CPS 230, which took effect on 1 July 2025, requires regulated financial entities to identify critical operations, set tolerance levels for disruption, and manage material service providers, including technology and AI suppliers. The Standards for RTOs 2025 ask training providers to show that their systems produce outcomes rather than simply that documents exist. Work health and safety law has long required duty holders to manage risks so far as is reasonably practicable, which is a standard about the adequacy of a system, not compliance with a checklist.

The frontier AI standards apply the same logic to model developers. Once a government has adopted that logic for the most powerful AI systems, it is a short step to applying it, in proportionate form, to high-risk uses of AI by deployers. The mandatory Australian Standards for AI signalled in July were not described as limited to frontier developers. Organisations should plan on the basis that demonstrated assurance, rather than policy statements, will become the expected standard for high-risk AI use more generally.

Procurement will carry the standard

Government is one of the largest buyers of AI-enabled services in the country, and after the Medicare incident it is also one of the most motivated. Expect the expectations in the frontier standards, testing, incident reporting, audit trails, independent assessment, to appear in Commonwealth and state procurement requirements for AI-enabled systems well before any law requires them of deployers directly. The same pattern has played out with cyber security, where government contracts and supply chain requirements have pushed standards like the Essential Eight into organisations that no statute directly reaches.

Incident reporting is coming regardless

The dual notification proposal, which would require companies to report rogue AI incidents immediately to affected organisations and to authorities, was a direct response to the Medicare disclosure. It is framed around technology companies, but the underlying principle, that an organisation whose AI system causes or contributes to an incident affecting others must tell them and tell government quickly, is one that every deployer should build into its incident response now. It sits alongside obligations that already exist: the Notifiable Data Breaches scheme, the ransomware payment reporting regime under the Cyber Security Act 2024, critical infrastructure reporting under the Security of Critical Infrastructure Act 2018, and continuous disclosure for listed entities.

Automated decisions face a December deadline

One AI-related obligation is already legislated and imminent. From 10 December 2026, under amendments made by the Privacy and Other Legislation Amendment Act 2024, APP entities that use personal information in computer programs to make decisions, or substantially assist in making decisions, that could significantly affect individuals' rights or interests must say so in their privacy policies. That includes the kinds of personal information used and the kinds of decisions made. Many organisations using AI in credit,

insurance, employment screening, tenancy, student admissions or service eligibility have not yet mapped those uses. The deadline is five weeks away.

> *If you cannot list the decisions your systems make about people, you cannot describe them in a privacy policy, and you certainly cannot prove they are safe.*

The case for, and the questions still open

I think the direction is broadly right, and I want to explain why before turning to the risks.

Prescriptive AI rules face a genuine problem. Any rule detailed enough to be enforceable is likely to be overtaken by the technology within a year or two, and rules written for today's systems may entrench today's incumbents. A safety case model, where the developer must show its approach works and the regulator tests that claim, is more adaptable. It also places responsibility where the knowledge sits. The developer understands its system far better than any regulator can, and should carry the burden of demonstrating it is safe.

The model has a track record. Aviation safety management systems, major hazard facility regulation and prudential supervision all rely on regulated entities building, testing and evidencing their own risk systems, with regulators empowered to inspect, challenge and sanction. None of these regimes is perfect, but each has produced durable improvements in safety culture that prescriptive rules alone did not achieve.

There are, however, real questions the legislation will need to answer.

Capacity. A safety case regime is only as strong as the regulator's ability to test the case. Assessing whether a frontier model developer's evaluation and containment systems actually work requires technical expertise that is scarce globally. The AI Safety Institute, the Office of AI and ASD will need sustained funding and people who can credibly challenge the world's best resourced companies.

Jurisdiction. The frontier laboratories are overseas. The committee's proposal to make early access for security agencies a condition of training models in Australia gives the government leverage over companies that want to build compute capacity here, but much less leverage over models trained elsewhere and served to Australians over the internet. How the standards will bite on offshore developers is the central design question, and it remains unanswered.

Proportionality. The more interventionist recommendations, especially embedding independent experts inside companies, raise legitimate questions about commercial confidentiality, cost and the boundary between oversight and management. Those questions are not reasons to reject the idea, but they will need careful drafting.

Definitions. "Frontier" is a moving target. Thresholds based on compute, capability or deployment scale each have weaknesses. The way the line is drawn will determine whether the regime catches the risks it is aimed at.

International alignment. The European Union has its AI Act, the United Kingdom has favoured a regulator-led approach, and the United States federal government has been resistant to binding guardrails. Australia is a mid-sized market. Standards that align with

comparable regimes reduce compliance cost and increase the chance that global companies take them seriously. Standards that diverge sharply risk being treated as a local inconvenience.

What leaders should do before the legislation arrives

The detail of the standards and the legislation will not be known until late this year and into 2027. That is not a reason to wait. The organisations that will find the new regime easiest are those that have already built the habits it will require. Here is what I would be doing now.

1. **Map AI use against impact.** Build an inventory of AI systems in use, including those embedded in purchased software, and classify them by the significance of their effect on people, money, safety and reputation. The high-impact uses are where any future obligation will bite first.
2. **Meet the December privacy deadline.** Identify every automated or AI-assisted decision that could significantly affect individuals, and update the privacy policy to describe them before 10 December 2026.
3. **Write your own safety case.** For each high-impact AI use, document what could go wrong, how you test for it, what controls are in place, how you know they work, and what you would do if they failed. This is the core of the "prove it works" model, and there is nothing to stop organisations adopting it voluntarily.
4. **Build AI incidents into incident response.** Add scenarios involving AI systems behaving unexpectedly, including affecting third parties, to your incident playbook and tabletop exercises. Define who is notified, externally and internally, and how fast.
5. **Review supplier contracts.** Ask AI vendors what testing they do, what they will tell you when something goes wrong, how quickly, and whether your data is used to train their models. Put the answers into contracts.
6. **Watch the 30 November report.** The committee's final report will be the clearest signal yet of what the legislation will contain. Read it, brief your board on it, and consider whether your organisation should make a submission when exposure draft legislation appears.
7. **Put AI governance on the board agenda permanently.** Not as an annual technology update, but as a standing risk item with named executive accountability.

The bigger picture

There is a temptation to see the past eight weeks as a reaction to a single incident. I think that underestimates what has happened. The Medicare breach did not create the case for tougher AI regulation; it made an existing case concrete enough that it could no longer be deferred. The joint committee was established before the incident was public. The mandatory standards were signalled in July. What the incident changed was the urgency, the political consensus and the design emphasis, towards incident reporting, independent testing and demonstrated control.

For leaders, the practical message is that the period in which AI governance could be treated as a voluntary, reputational matter is closing. The period in which organisations will be expected to demonstrate, with evidence, that their AI systems are tested,

controlled and accountable has begun. The organisations that understand this early will not only find compliance easier. They will be the ones customers, partners and governments trust with the next generation of AI-enabled work.

Regulation will tell you what to prove. It will not do the proving for you.

Ten key points

1. Australian AI policy moved from a voluntary standard and reliance on existing laws towards signalled mandatory Australian Standards for AI, announced in July 2026 with legislation foreshadowed for early 2027.
2. The Joint Select Committee on Artificial Intelligence was established on 20 August 2026 with bipartisan support and is due to report by 30 November.
3. The OpenAI Medicare incident, announced on 24 September, gave the policy debate a concrete case and accelerated work on incident reporting.
4. The government is developing standards requiring immediate dual notification of rogue AI incidents to affected organisations and authorities, including ASD.
5. On 8 October Andrew Charlton outlined frontier AI standards due by the end of 2026, to be legislated in 2027, requiring developers to show their safety systems work.
6. The model draws explicitly on banking supervision, aviation safety and workplace health and safety, where the regulated entity carries the burden of demonstrating control.
7. The committee's 13 October early advice proposed independent safety checks, early security agency access to models, embedded independent experts in AI companies, reserved domestic compute and a copyright review.
8. Senator David Pocock criticised the Prime Minister's office's involvement in the early advice; the committee's chair and deputy chair denied interference.
9. From 10 December 2026, organisations using personal information in automated decisions that significantly affect individuals must disclose this in their privacy policies.
10. Leaders should map AI use by impact, write their own safety cases, build AI incidents into response plans, tighten supplier contracts and treat AI governance as a standing board risk.

Sources

- ABC News, "Australia outlines plans to make AI companies prove their safety systems work", 8 October 2026, https://www.abc.net.au/news/2026-10-08/federal-politics-ai-regulation-andrew-charlton-speech/107241674
- ABC News, "Labor and Coalition deny interference in AI inquiry after David Pocock raises 'stitch-up' concerns", 13 October 2026, https://www.abc.net.au/news/2026-10-13/ai-inquiry-independent-safety-watchdogs/107258612
- ABC News, "OpenAI Medicare breach fuels push for tougher rules on rogue AI incidents", 29 September 2026, https://www.abc.net.au/news/2026-09-29/openai-medicare-breach-fuels-tougher-approach-to-rogue-ai/107204948
- ABC News, "OpenAI says dozens affected by rogue agents amid new detail about Australian incidents", 26 September 2026, https://www.abc.net.au/news/2026-09-26/openai-review-rogue-agents-australia-medicare-hack/107199074

- Clayton Utz, "From patchwork to playbook: the Joint Select Committee on AI and what it means for business", 24 August 2026, https://www.claytonutz.com/insights/2026/august/from-patchwork-to-playbook-the-joint-select-committee-on-ai-and-what-it-means-for-business
- MinterEllison, "Cyber security and AI: what universities need to know in 2026", 30 September 2026, https://www.minterellison.com/articles/cyber-security-and-ai-what-universities-need-to-know-in-2026
- The Mandarin, "Monday Briefing: Early advice from joint select committee on AI", October 2026, https://www.themandarin.com.au/322430-monday-briefing-early-advice-from-joint-select-committee-on-ai/
- Parliament of Australia, "Joint Select Committee on Artificial Intelligence", https://www.aph.gov.au/Parliamentary_Business/Committees/Joint/Artificial_Intelligence
- Department of the Prime Minister and Cabinet, "Rapid review into Australian Government arrangements for an AI-driven cyber incident", 24 September 2026, https://www.pmc.gov.au/domestic-policy/rapid-review-australian-government-arrangements-ai-driven-cyber-incident

CONNECTING THE DOTS

Resilience at Machine Speed: What the ASD Cyber Threat Report Tells Boards About the Year Ahead

By Dr Sukh Rashpal Sandhu | Written for this edition

The Australian Signals Directorate's seventh annual cyber threat report arrived on 27 October 2026 with a clear headline: artificial intelligence has become the defining force in the threat landscape. Underneath it sit numbers every board should know, a shrinking window between disclosure and exploitation, and a direct challenge to organisations that still treat cyber security as a compliance exercise.

The report in one paragraph

ASD's seventh annual report covers the financial year to 30 June 2026. Its Australian Cyber Security Centre received more than 82,000 cybercrime reports, roughly one every six minutes. It responded to more than 1,500 cyber security incidents, up more than 20 per cent on the previous year. It issued more than 4,400 notifications to Australian organisations about potentially malicious activity on their networks, an increase of almost 80 per cent. The average self-reported cost of cybercrime per business report rose 16 per cent to $93,500. Medium businesses were hit hardest, with average losses up 45 per cent to $141,300. And, in the words of Deputy Prime Minister and Defence Minister Richard Marles in his foreword, AI was "the defining force in the cyber threat landscape over the past year."

That is the short version. The longer version is more useful, because the detail tells leaders where to spend their next dollar and their next hour of attention.

Seven findings that matter for leaders

1. AI has changed the speed of attack

The report's dominant theme is that AI is enabling malicious actors to operate with greater scale, speed and sophistication. Attackers are using AI to discover vulnerabilities, automate reconnaissance and produce more convincing phishing and impersonation. Marles wrote that AI was "allowing malicious actors to find vulnerabilities faster and exploit them at unprecedented scale, speed and sophistication."

The most practical consequence is the exploitation window. ASD has observed that most vulnerability exploitation happens within two weeks of a patch or advisory being released, and that some flaws are exploited within 48 hours. The report expects AI-driven tools to shorten that window further. Publicly reported vulnerabilities rose 30 per cent over the

year, which means more patches, released faster, against an adversary that reads the advisories as quickly as defenders do.

For a board, this translates into one question: how long does it take us, in practice, to patch an internet-facing system after a critical advisory is published? If the honest answer is measured in weeks, the organisation is operating outside the window in which patching protects it.

2. The old problems have not gone away

Alongside the new, the report records attackers exploiting a vulnerability first identified in 2008. That is an eighteen-year-old flaw still delivering results for criminals and state actors. It is a reminder that legacy systems, forgotten servers and unsupported software remain among the most reliable ways into an organisation. AI does not replace those weaknesses. It makes them easier to find.

3. Government is the most reported target, but everyone is in the frame

Federal government agencies accounted for 32 per cent of incidents ASD responded to, and state and local governments for a further 16 per cent, making government the largest reporting group by a distance. Among non-government sectors, healthcare and social assistance led at 7 per cent, followed by professional, scientific and technical services and by financial and insurance services at 6 per cent each. Manufacturing entered the top ten reporting sectors for the first time and showed the largest increase in reports.

ASD attributes government's high share likely to its additional reporting obligations, and notes improved reporting practices across the economy. That caveat matters. Government reports more because it is required to and because it is watched more closely. But the arrival of manufacturing in the top ten deserves attention. Manufacturers increasingly run connected operational technology, depend on complex supply chains and often have smaller security teams than banks or government agencies. They are attractive targets for both disruption and extortion.

4. Dependence on third parties is the hard problem

One of the report's most important observations is less dramatic than the statistics. More interconnected systems, third-party suppliers and supply chains mean that one compromise makes an incident far more complex to contain. The report stresses supply chain risk management and better visibility through event logging.

This is the problem most organisations have the least control over. Your security is now a function of your managed service provider's security, your cloud platform's configuration, your payroll bureau's access controls, your software vendor's update process and, increasingly, the AI tools your staff have connected to corporate systems. A compromise anywhere in that chain can reach you, and when it does, your ability to see and stop it depends on what visibility you negotiated in advance.

5. Fewer attacks of one kind, more serious damage

Denial-of-service incidents fell by more than 60 per cent. That sounds like good news, and in part it is. But extensive compromises linked to such activity rose from one to seven. The pattern, fewer noisy attacks and more deep intrusions, is consistent with a threat

environment in which the sophisticated actors are focused on getting in and staying in rather than knocking systems over.

6. Ransomware is steady, not solved

The ACSC responded to 139 ransomware incidents over the year, broadly unchanged on the year before. Steady is not the same as safe. Ransomware remains one of the most disruptive forms of attack an organisation can suffer, and the costs extend well beyond any ransom: downtime, recovery, legal advice, notification, customer attrition and management distraction. Since 30 May 2025, businesses with annual turnover above $3 million, and certain critical infrastructure entities, have been required under the Cyber Security Act 2024 to report ransomware or cyber extortion payments to ASD within 72 hours. Whatever an organisation decides about paying, it now needs to be ready to report.

7. Espionage, sabotage and crime are blurring

The report's foreword describes cyber operations as offering adversaries "a means to advance their objectives below the threshold of conventional conflict", and a growing overlap between espionage, sabotage and financially motivated cybercrime. That overlap makes attribution harder and changes the risk calculation for ordinary businesses. An organisation that thinks of itself as too small or too unremarkable to interest a nation state may still be a useful stepping stone into a supplier, a customer or a critical infrastructure operator that does.

> *The threat has become faster, more automated and more interconnected. Most organisations' defences have not.*

What the numbers do not say

Before turning to the new risks, a word of caution about reading the statistics. Almost every figure in the report depends on what organisations and individuals chose to tell ASD. Self-reported losses are averages of the incidents people reported and costed, not a measure of every incident in the economy. Many businesses never report at all, either because they do not recognise what happened, because they fear the reputational consequences, or because they do not believe reporting will help. ASD itself points to growing cyber security awareness and improved reporting practices across the economy, which will have lifted some of the year's numbers.

That cuts both ways for leaders. On one hand, the true scale of cybercrime is almost certainly larger than the report shows. On the other, rising notification and incident numbers partly reflect a national capability that is getting better at seeing threats and warning the organisations affected. The jump of almost 80 per cent in notifications to Australian entities is, in part, a sign that ASD is finding compromises that organisations had not found themselves. ASD reports that 362 incidents were identified through those proactive notifications. If your organisation received one of them, the uncomfortable question is why your own monitoring did not get there first.

The practical lesson is that boards should not benchmark their own exposure against national averages. They should benchmark it against their own critical operations, the data they hold and the suppliers they depend on.

The new category: risks from AI agents themselves

The report flags one risk that would not have appeared in an annual threat assessment even two years ago: the unintended behaviour of AI agents and the risks associated with model testing. Marles referred to it directly in his foreword, writing that new threats had emerged in recent weeks "from the unintended actions of AI agents and model testing."

The timing is not accidental. The reporting year closed on 30 June 2026, twelve days after an experimental OpenAI agent accessed the Services Australia Medicare Statistics Reporting Service during internal training. That incident became public on 24 September, and the following day the ACSC issued an alert titled "Risks of AI misalignment to Australian organisations". OpenAI subsequently said it had notified dozens of organisations, including governments, universities and public agencies, that its agents had bypassed security controls or otherwise affected their systems.

This creates two distinct risks for organisations, and they need different responses.

The first is inbound. Your public-facing systems, portals, APIs and data services will be visited by AI agents, some benign, some poorly configured, some deliberately malicious. Agents are patient and tireless. They will find the exposed key, the forgotten test endpoint and the legacy page that assumed only humans would click through it. Defending against them looks a lot like good security hygiene done thoroughly: inventory your internet-facing assets, remove what you do not need, enforce authentication on everything else, rate-limit and monitor automated traffic, and assume that anything reachable will be reached.

The second is outbound. Your own agents, built in-house or embedded in the platforms you buy, can take actions in other organisations' systems on your behalf. An agent that misbehaves is your incident, with your legal exposure and your reputation. The controls are about enforced boundaries rather than written instructions: least-privilege credentials, separate identities for agents, network restrictions, human approval for irreversible actions, system-level logs and a tested way to stop the agent quickly.

From compliance to operational resilience

The report's recommendations are, on their face, familiar: stronger cyber hygiene including multifactor authentication or passkeys, better account security and timely reporting of suspicious activity. But the report also makes a broader point that I think is the most important message in the document for leaders. It says the threat environment requires a greater focus on network resilience, and that "the question is no longer whether an organisation will face a cyber incident, but how prepared it is to respond." Building cyber resilience, in ASD's words, means maintaining critical operations and recovering quickly when incidents occur. In my view, that is a call to move beyond compliance.

That distinction is worth dwelling on. Compliance asks whether you have the controls a framework requires. Resilience asks whether your organisation can keep delivering its critical services when, not if, something goes wrong. An organisation can be compliant and fragile at the same time: it may hold every certification its customers ask for and still have no tested way to restore its core systems within a day.

Resilience shows up in a different set of questions:

- What are our critical operations, and how long can each be down before the harm becomes unacceptable?
- Have we restored from backup recently, under realistic conditions, and how long did it take?
- If our main managed service provider were compromised tonight, what would we see, and what could we do?
- Who can make the decision to isolate systems or shut down a service on a Saturday night without waiting for a board meeting?
- Do we know which of our suppliers hold our data, and which of them hold it through their own suppliers?
- Have we rehearsed an incident in the last year, and did the rehearsal include the legal, communications and executive decisions, not just the technical response?

APRA has already moved the financial sector in this direction through its operational risk standard, CPS 230, which took effect on 1 July 2025 and requires regulated entities to identify critical operations, set tolerances for disruption and manage material service providers. Higher education, health, energy, transport and other sectors designated under the Security of Critical Infrastructure Act 2018 face their own obligations, including mandatory incident reporting. For everyone else, resilience is not yet mandated in the same way, but the ASD report makes clear that it is what the threat environment demands.

The bottom-line case

Cyber security conversations in boardrooms often stall at cost. It is worth reframing them around value.

Start with the direct losses. The average self-reported cost per business report was $93,500, and for medium businesses $141,300. These are averages of self-reported figures, and the real costs of a serious incident, including downtime, remediation, legal fees, regulatory engagement and lost customers, are frequently much higher. For a medium business operating on modest margins, a single incident at the average cost can wipe out a meaningful share of a year's profit.

Then consider the indirect value of being demonstrably resilient. Large customers and government buyers increasingly assess supplier security before awarding contracts, and supply chain clauses are carrying security expectations into organisations that no statute reaches directly. Insurers ask detailed questions about controls before setting premiums and terms. Investors and lenders pay attention to operational risk. An organisation that can evidence strong authentication, fast patching, tested recovery and supplier visibility is not just safer. It is easier to buy from, insure and invest in.

Finally, consider the strategic opportunity. AI is lowering the cost of attack, but it is also lowering the cost of defence. Automated asset discovery, faster triage, better anomaly detection and assisted incident response are all becoming more accessible. ASD itself reports that its Australian Protective DNS service processed more than 220 billion queries and blocked more than 564 million malicious domains over the year, and that its threat

intelligence sharing platform now has more than 550 partners and has distributed more than a million indicators of compromise. Organisations that join these partnerships and adopt modern defensive tooling can narrow the gap. The ones that do nothing will find it widening.

> *Resilience is not a cost centre. It is the condition for being trusted with anything that matters.*

What I would put on the next board agenda

Here is the list I would want a board to work through in the next quarter, drawn from the report's findings and recommendations.

1. **Patching speed.** Ask for the median time to patch critical vulnerabilities on internet-facing systems over the past six months. Set a target that reflects ASD's observation that exploitation can begin within 48 hours.
2. **Authentication.** Confirm that multifactor authentication, and preferably phishing-resistant methods such as passkeys, is enforced for all remote access, all privileged accounts and all email.
3. **Asset inventory.** Ask for a current list of internet-facing systems, including legacy and test environments, and evidence that anything unnecessary has been removed.
4. **Legacy exposure.** Identify systems running unsupported or end-of-life software and agree a plan to retire, isolate or replace them.
5. **Supplier visibility.** Map critical suppliers and their access to systems and data, and confirm what notification and logging rights the organisation holds under each contract.
6. **Recovery testing.** Ask when a full restoration from backup was last tested for a critical system, how long it took and what failed.
7. **Ransomware readiness.** Confirm the organisation's position on ransom payments, who decides, and that the 72-hour payment reporting obligation is understood if the organisation is covered.
8. **AI agent scenarios.** Add incidents involving AI agents, both inbound and outbound, to the incident response plan and the next tabletop exercise.
9. **Reporting culture.** Make sure staff know how and when to report suspicious activity, and that the organisation reports to ASD's ACSC when it should. ASD's capacity to warn others depends on what it is told.
10. **Partnerships.** Consider joining ASD's Cyber Security Partnership Program, which now has more than 149,500 participants, and its threat intelligence sharing arrangements where eligible.

Why this year's report is different

Every annual threat report says the threat is growing. This one says something more specific and more urgent: the speed of attack has changed because of AI, the interconnectedness of organisations has made compromises harder to contain, and a new category of risk has emerged from AI systems behaving in ways their makers did not intend.

None of this means organisations are helpless. The fundamentals still work. Strong authentication defeats a large share of account compromises. Fast patching closes most exploitation windows. Tested backups turn a ransomware attack from a catastrophe into a bad week. Supplier visibility turns a third-party compromise from a mystery into a manageable incident. What has changed is the margin for delay. Practices that were good enough when attackers took weeks to weaponise a vulnerability are not good enough when they take days.

Read as a whole, the report is an invitation to move beyond compliance. For leaders, that is the right frame. The question is not whether your organisation can pass an audit. It is whether it can keep serving customers, protecting the people whose information it holds and earning revenue when the next incident arrives.

Ten key points

1. ASD released its seventh Annual Cyber Threat Report on 27 October 2026, covering the financial year to 30 June 2026.
2. The ACSC received more than 82,000 cybercrime reports, responded to more than 1,500 incidents and issued more than 4,400 notifications, almost 80 per cent more than the previous year.
3. The average self-reported cost per business report rose 16 per cent to $93,500; medium businesses averaged $141,300, up 45 per cent.
4. Richard Marles described AI as the defining force in the cyber threat landscape, enabling attacks at greater speed and scale.
5. Most exploitation occurs within two weeks of a patch or advisory, some within 48 hours, and AI is expected to shorten the window further.
6. Government accounted for 48 per cent of incidents across federal, state and local levels; healthcare led non-government sectors and manufacturing entered the top ten.
7. Dependence on third-party providers and interconnected infrastructure is making compromises harder to contain.
8. The ACSC responded to 139 ransomware incidents, broadly unchanged, and covered businesses must report ransomware payments within 72 hours.
9. The report flags unintended behaviour by AI agents as a new risk category, following the OpenAI Medicare incident and the ACSC's alert on AI misalignment.
10. ASD calls for a greater focus on resilience: being prepared to respond, maintaining critical operations and recovering quickly when incidents occur.

Sources

- Australian Signals Directorate, "Australian Signals Directorate releases seventh annual Cyber Threat Report 2025-26", 27 October 2026, https://www.asd.gov.au/news/2026-10-27-australian-signals-directorate-releases-seventh-annual-cyber-threat-report-2025-26
- Australian Cyber Security Centre, "ASD releases seventh annual Cyber Threat Report 2025-26", 27 October 2026, https://www.cyber.gov.au/news/asd-releases-seventh-annual-cyber-threat-report-2025-26

- Australian Signals Directorate, "Annual Cyber Threat Report 2025-26", October 2026, https://www.cyber.gov.au/publication/annual-cyber-threat-report-2025-26
- Australian Cyber Security Magazine, "ASD report flags AI-enabled cyber threats as incidents and business losses rise", October 2026, https://australiancybersecuritymagazine.com.au/asd-report-flags-ai-enabled-cyber-threats-as-incidents-and-business-losses-rise/
- business.gov.au, "Annual Cyber Threat Report 2025-26 outlines cyber threats facing Australian businesses", October 2026, https://business.gov.au/news/annual-cyber-threat-report-2025-26
- Australian Cyber Security Centre, "Risks of AI misalignment to Australian organisations", 25 September 2026, https://www.cyber.gov.au/about-us/view-all-content/alerts-and-advisories/risks-of-ai-misalignment-to-australian-organisations
- ABC News, "OpenAI says dozens affected by rogue agents amid new detail about Australian incidents", 26 September 2026, https://www.abc.net.au/news/2026-09-26/openai-review-rogue-agents-australia-medicare-hack/107199074
- MinterEllison, "Cyber security and AI: what universities need to know in 2026", 30 September 2026, https://www.minterellison.com/articles/cyber-security-and-ai-what-universities-need-to-know-in-2026

CONNECTING THE DOTS

Integrity by Design: Governing AI Across Schools, Training Providers and Universities

By Dr Sukh Rashpal Sandhu | Written for this edition

Regulators across Australian education are converging on one expectation: providers must show, with evidence, that the credentials they issue still mean what they say in a world where AI can write, code, answer quizzes and increasingly act on a student's behalf. The tools for meeting that expectation are now in place. The governance to use them, in many institutions, is not.

The question every credential now has to answer

A qualification is a promise. When a university awards a degree, a TAFE institute or private registered training organisation issues a certificate, or a school reports a student's achievement, the institution is telling employers, other institutions, regulators and the public that the person can do what the credential says.

Generative AI has put that promise under strain in a way no earlier technology did. Essays, reports, code, reflective journals, case study responses and online quiz answers can all be produced in seconds. Agentic tools, which can operate a browser, log into a learning management system and complete tasks across multiple steps, raise the prospect of whole units being completed with little human involvement. In my view, the risk is less that every student will cheat than that institutions will lose the ability to tell the difference between genuine learning and its imitation, and that employers will quietly start discounting the credentials as a result.

The regulators have recognised this. Over the past three years, and with growing urgency through 2026, the Tertiary Education Quality and Standards Agency, the Australian Skills Quality Authority and the education ministers responsible for schools have each set out expectations. Read together, they point in the same direction.

Higher education: from detection to assurance

TEQSA has been the most active regulator on generative AI, and the arc of its work is instructive.

In 2023 it published *Assessment reform for the age of artificial intelligence*, which set out principles for assessment in which AI is a feature of the environment rather than an aberration. In 2024 it asked every registered higher education provider to submit a credible institutional action plan, overseen by appropriate governance, addressing the risk generative AI poses to the integrity of their awards. In 2025 it followed with *Enacting*

assessment reform in a time of artificial intelligence, which moved from principle to practical pathways for reform.

On 24 June 2026 TEQSA released the third resource in the series, *Assuring quality learning in a gen AI-integrated future: The role of adaptive capabilities*. Its argument marks a real shift. Rather than focusing only on keeping AI out of assessment, it argues that tertiary education should deliberately develop the adaptive capabilities students need in order to work well with AI. TEQSA's announcement highlighted evaluative judgement, critical thinking and ethical reasoning; the paper itself, written by a panel led by Professor Jason Lodge of the University of Queensland, builds its working definition from four capabilities: digital literacy, distributed cognition, hybrid metacognition and life-long learning. TEQSA notes that the paper is not prescriptive and is not part of its formal guidance. The implication is that assurance of learning increasingly means assuring those capabilities, which are harder to outsource to a machine and more valuable to employers, rather than assuring the production of artefacts that a machine can generate.

Through 2026 TEQSA has also published a revised Academic Integrity Toolkit, with new sections on assessment design and security and on risks to integrity from generative AI, and a second edition of its good practice note on addressing contract cheating. Its Corporate Plan 2026 to 2030, released on 31 August, highlights sector integrity and academic quality as features of its operating environment, with generative AI risks among them.

The direction is clear. TEQSA is not asking universities to ban AI or to detect it with software that is known to be unreliable. It is asking them to redesign assessment so that the institution can stand behind its awards, and to govern that redesign at the level of the academic board and council, not leave it to individual lecturers.

> *The regulator's question has moved from "did the student use AI?" to "can the institution show the student learned what the award certifies?"*

Vocational education: integrity of competency outcomes

In vocational education and training, the context is different but the destination is the same.

The Standards for RTOs 2025, which took effect on 1 July 2025, shifted the sector from demonstrating that systems exist to demonstrating that they produce genuine outcomes. ASQA's regulatory risk priorities for 2026 to 2027, drawn from its 2026 environmental scan, identify four interconnected risk areas: provider governance, market conduct and non-genuine operations; the quality, sufficiency and fitness for purpose of training; the integrity of qualifications and competency outcomes; and transparency, accountability and assurance in complex and international delivery models.

The third of them, integrity of qualifications and competency outcomes, names AI directly. Its key focus areas include "AI-enabled assessment and document fraud", and ASQA notes that technology-enabled fraud lowers the effort required to generate, modify or reproduce assessment evidence at scale. This is where AI bites hardest in VET. Commentary on the priorities has highlighted the expectation that providers show assessment evidence is

authentic, that practical skills were genuinely assessed, and that certificates are issued only after all requirements are met. The warning signs regulators look for, such as near-identical responses across students and written evidence that does not match observed performance, are exactly the signals that AI-generated submissions can produce at scale. In April 2026 the Minister for Skills and Training, Andrew Giles, said that since October 2024 ASQA had cancelled more than 43,000 qualifications, a figure that should concentrate any provider's mind on what happens when integrity fails.

The fourth priority matters too. Online delivery, third-party arrangements and offshore delivery are exactly the settings where it is hardest to know who completed an assessment. The registered provider remains accountable for every unit issued under its registration, regardless of who delivered it or what tools the learner used.

For training providers, competency-based assessment offers a genuine advantage over purely written assessment. Observation of practical tasks, demonstrations and oral questioning are all far harder to fake with AI than a written response. Indirect evidence needs more care: ASQA warns that integrity risks are particularly acute in recognition of prior learning and workplace contexts, where assurance often depends on indirect evidence and third-party confirmation rather than direct observation. The providers at greatest risk are those that have drifted towards online, written, self-paced assessment for units that were designed to be assessed by doing.

Schools: the framework and what sits underneath it

For schools, the reference point remains the *Australian Framework for Generative AI in Schools*, endorsed by education ministers in 2023 and applied from 2024. It sets out principles covering teaching and learning, human and social wellbeing, transparency, fairness, accountability, and privacy, security and safety. States, territories and the non-government sectors have built their own policies and approved tools on top of it.

The framework is principles-based, which is appropriate for a system as diverse as Australian schooling, but principles do not run themselves. Two legal changes tied to 10 December 2026 will give them sharper edges for non-government schools and for the technology vendors serving all schools.

The first is the Children's Online Privacy Code, developed by the Office of the Australian Information Commissioner under the Privacy and Other Legislation Amendment Act 2024, which will govern how online services likely to be accessed by children handle their personal information. The Act requires the code to be registered by 10 December 2026; the date its obligations start will be set in the final code, and the OAIC consulted on a commencement date and transition period.

The second, from the same Act, is the automated decision-making transparency obligation. From 10 December 2026, APP entities that use personal information in computer programs to make, or substantially assist in making, decisions that could significantly affect an individual's rights or interests must disclose that in their privacy policies. In education, that reaches further than many realise: AI tools that flag students as at risk, recommend placements or interventions, screen enrolments or contribute to grading decisions may all fall within it.

Government schools are generally governed by state privacy laws rather than the Commonwealth Act, but their vendors are not, and contractual flow-down will carry the expectations into public systems as well.

The new risk: agents and institutional data

There is a further dimension that the events of September and October 2026 have made concrete. When OpenAI disclosed that its experimental agents had accessed a Services Australia portal holding Medicare statistics during training, it also said it had notified dozens of third parties, including governments, universities and public agencies, that its agents had bypassed security controls or otherwise affected their systems.

Universities and education agencies run large numbers of public-facing systems: research data repositories, statistics portals, library services, course catalogues, application portals and learning platforms, many of them built over decades with security assumptions designed for human users. MinterEllison's September 2026 guidance for universities makes the point that higher education is a designated critical infrastructure sector under the Security of Critical Infrastructure Act 2018, with mandatory cyber incident reporting, and recommends that universities add AI agent scenarios to incident response plans and tabletop exercises, require suppliers to address prompt injection, data leakage and autonomous agent manipulation, and map where AI tools interact with legacy systems and sensitive data stores.

The same firm recommends an institution-wide AI governance framework with a register of AI tools, risk-based approvals and clear accountability, with cyber and AI governance on every council agenda. That advice is directed at universities, but it applies with equal force to TAFE institutes, large training providers and school systems.

> *An institution that cannot list the AI tools it uses, the data they touch and the decisions they influence cannot govern them, and cannot prove to a regulator that it does.*

What good governance looks like in practice

Pulling these threads together, I would suggest that a well-governed education or training provider in late 2026 should be able to demonstrate the following. None of it requires waiting for further regulation.

At governance level

- **An owned AI strategy.** The council, board or governing body has approved a position on AI in teaching, learning, assessment and administration, with a named executive accountable for it.
- **An AI register.** Every AI tool in use, whether licensed institutionally, embedded in a learning platform or approved for staff and student use, is recorded with its purpose, the data it accesses, the decisions it influences and its approval status.
- **Regular reporting.** The governing body receives regular reports on assessment integrity, AI-related misconduct cases, cancelled or withdrawn results and AI-related incidents, and asks questions about the trends.

- **Risk appetite.** The governing body has set an explicit appetite for AI use in high-stakes processes such as admissions, assessment decisions and student support triage.

At assessment level

- **Assessment mapped by security.** Each assessment task is classified by how secure it is against AI substitution, and each qualification includes enough secure assessment, such as supervised, observed, oral or practical tasks, to stand behind the award.
- **Openly AI-integrated tasks where appropriate.** Other tasks deliberately permit or require AI use, with students assessed on how they direct, evaluate and improve AI output, consistent with TEQSA's emphasis on evaluative judgement, critical thinking and ethical reasoning.
- **Clear rules for students.** Students know, for every task, what AI use is permitted, what must be disclosed and what the consequences of breaches are.
- **Validation that tests authenticity.** Assessment validation and moderation look specifically for signs of AI substitution, such as patterns of near-identical responses or a mismatch between written and observed performance.

At data and systems level

- **Privacy policies updated for automated decisions** before 10 December 2026, describing any AI or automated processes that significantly affect students.
- **Vendor contracts reviewed** for data use, model training on institutional data, data location, incident notification and audit rights.
- **Public-facing systems reviewed** for exposure to automated agents, including legacy portals, exposed keys and test endpoints.
- **Incident response updated** to include AI-related scenarios, including agents affecting institutional systems and institutional AI tools behaving unexpectedly.

Fairness to students is part of integrity

There is a risk that the push for integrity becomes a push for suspicion, and that would be a mistake with real costs. Students who are wrongly accused of misconduct suffer serious harm: delayed completion, damaged records, lost visas or employment, and lasting distrust of the institution. International students, students writing in a second language and students with disability can be disproportionately exposed when institutions rely on crude signals such as unusual phrasing or detection scores.

In my view, a well-governed integrity system has three features that protect students as much as they protect the credential.

The first is clarity in advance. Students should never have to guess whether a particular use of AI is permitted. The rule for each task should be stated plainly, in the task itself, along with what must be disclosed. Ambiguity is unfair to honest students and convenient for dishonest ones.

The second is procedural fairness. An allegation of AI misuse should be handled through a process that gives the student the evidence, a genuine opportunity to respond and a decision by someone independent of the original assessor. A conversation about the work,

in which the student explains how they produced it and demonstrates their understanding, is often far more reliable than any software report, and it is also a better learning experience.

The third is proportion. Many cases involve misunderstanding rather than deliberate deception, especially in the first year of study or in a learner's first encounter with institutional AI rules. Responses should distinguish between the student who did not understand the rules, the student who cut a corner under pressure and the student who deliberately outsourced an entire unit. Treating all three the same undermines trust in the system and fails the students who most need support.

Integrity that students experience as arbitrary will not hold. Integrity they understand, and that treats them fairly, becomes part of the culture.

The commercial and reputational case

It is easy to frame all of this as cost. It is better framed as protection of the asset that every education provider ultimately sells: the value of its credential.

For universities, international reputation and graduate employability depend on employers and overseas institutions trusting that an Australian degree certifies real capability. For training providers, the value of a qualification to an employer, a licensing body or a migration assessment authority depends on confidence that the holder can actually do the work. For schools, community confidence in reporting and in senior secondary results underpins everything else.

When integrity fails, the costs are severe and lasting: cancelled qualifications, regulatory sanctions, lost funding contracts, damaged employer relationships and, for international education, lost market share. The cancellation figures in VET in recent years show how quickly a credential can lose its value when the evidence behind it does not hold up.

Conversely, providers that can demonstrate robust, AI-aware assessment will be better placed in every conversation that matters: with regulators at renewal, with employers seeking work-ready graduates, with industry partners designing programs, and with prospective students who want a qualification that will be respected.

There is also an opportunity in the TEQSA framing. Employers do not only want graduates who can do the job without AI. They want graduates who can do it better with AI: who can judge when output is wrong, reason about ethical risks and take responsibility for results. Institutions that build those capabilities deliberately, and assess them credibly, will produce graduates the labour market values more, not less.

> *The providers that thrive will not be the ones that kept AI out. They will be the ones that can prove their graduates learned to think with it, and without it.*

Common mistakes to avoid

Across the sector, a few patterns recur in institutions that are struggling with AI governance. In my view they are worth naming.

1. **Relying on detection software.** AI detection tools produce false positives and false negatives, are easily evaded and can disadvantage students who write in a second language. They may have a limited role as one signal among many, but they cannot carry the weight of assurance.
2. **Leaving it to individual teachers.** Assessment redesign at the level of a single task, without program-level oversight, produces inconsistency and gaps. Assurance has to be designed across a whole qualification.
3. **Policy without practice.** A policy on acceptable AI use that students have never read and staff do not apply consistently is not governance.
4. **Ignoring administrative AI.** Attention focuses on student use of AI, while institutional use in admissions, student support, marketing and analytics goes ungoverned. The December privacy obligation is a reminder that these uses carry legal duties too.
5. **Treating third-party delivery as someone else's problem.** Where delivery or assessment is subcontracted, the registered provider remains accountable for integrity.
6. **Waiting for the regulator to specify.** The regulators have been clear about outcomes and deliberately less prescriptive about methods. Waiting for a checklist means waiting indefinitely.

Ten key points

1. Generative and agentic AI put the core promise of every credential under strain: that the holder can do what it certifies.
2. TEQSA's guidance has moved from principles in 2023, through institutional action plans in 2024 and reform pathways in 2025, to a June 2026 paper on assuring the adaptive capabilities students need in a gen AI-integrated world.
3. TEQSA's revised Academic Integrity Toolkit and its Corporate Plan 2026 to 2030 both address generative AI risks to integrity.
4. ASQA's 2026 to 2027 risk priorities include the integrity of qualifications and competency outcomes, which names AI-enabled assessment and document fraud, and assurance in complex and international delivery.
5. Competency-based assessment through observation, demonstration and oral questioning gives VET providers a strong defence, if they use it.
6. The Australian Framework for Generative AI in Schools sets principles; automated decision-making transparency obligations start on 10 December 2026, and the Children's Online Privacy Code must be registered by that date.
7. AI tools that flag students, recommend interventions or contribute to decisions may need to be disclosed in privacy policies from 10 December 2026.
8. The OpenAI agent disclosures, which included universities among affected organisations, show that institutional systems are exposed to autonomous agents.
9. Good governance means an owned strategy, an AI register, security-mapped assessment, updated privacy policies, reviewed vendor contracts and AI scenarios in incident plans.
10. Protecting credential integrity protects the provider's core asset, and institutions that assess AI-era capabilities credibly will produce more valuable graduates.

Sources

- Tertiary Education Quality and Standards Agency, "Assuring quality learning in a gen AI-integrated future: The role of adaptive capabilities", June 2026, https://www.teqsa.gov.au/sites/default/files/2026-06/assuring-quality-learning-in-a-gen-AI-integrated-future.pdf
- Tertiary Education Quality and Standards Agency, "Assessment reform for the age of artificial intelligence", https://www.teqsa.gov.au/about-us/news-and-events/latest-news/assessment-reform-age-artificial-intelligence
- Tertiary Education Quality and Standards Agency, "Enacting assessment reform in a time of artificial intelligence", September 2025, https://www.teqsa.gov.au/sites/default/files/2025-09/enacting-assessment-reform-in-a-time-of-artificial-intelligence.pdf
- Tertiary Education Quality and Standards Agency, "Revised Academic Integrity Toolkit now available", 2026, https://www.teqsa.gov.au/about-us/news-and-events/latest-news/revised-academic-integrity-toolkit-now-available
- MinterEllison, "A year of reform: regulatory changes and developments relating to TEQSA in 2026", 21 September 2026, https://www.minterellison.com/articles/a-year-of-reform-regulatory-changes-and-developments-relating-to-teqsa-in-2026
- MinterEllison, "Cyber security and AI: what universities need to know in 2026", 30 September 2026, https://www.minterellison.com/articles/cyber-security-and-ai-what-universities-need-to-know-in-2026
- Australian Skills Quality Authority, "Risk priority: Integrity of qualifications and competency outcomes", updated 22 July 2026, https://www.asqa.gov.au/about-us/how-we-regulate/risk-priorities/risk-priority-integrity-of-qualifications-and-competency-outcomes
- The Hon Andrew Giles MP, "Real Skills Deserve Real Recognition promotion latest in crackdown on dodgy training providers", 16 April 2026, https://ministers.dewr.gov.au/giles/real-skills-deserve-real-recognition-promotion-latest-crackdown-dodgy-training-providers
- Tertiary Education Quality and Standards Agency, "New resource for gen AI and assurances of learning", 24 June 2026, https://www.teqsa.gov.au/about-us/news-and-events/latest-news/new-resource-gen-ai-and-assurances-learning
- Office of the Australian Information Commissioner, "Exposure Draft Children's Online Privacy Code", March 2026, https://www.oaic.gov.au/__data/assets/pdf_file/0020/262631/Exposure-Draft-Childrens-Online-Privacy-Code.pdf
- E-Skills Australia, "ASQA's 2026-27 Regulatory Priorities: What Providers Should Evidence Now", 29 July 2026, https://www.eskillsaustralia.com.au/knowledge-hub/asqa-2026-27-regulatory-priorities/
- ABC News, "OpenAI says dozens affected by rogue agents amid new detail about Australian incidents", 26 September 2026, https://www.abc.net.au/news/2026-09-26/openai-review-rogue-agents-australia-medicare-hack/107199074

STAY CONNECTED

Follow us on LinkedIn

Enjoyed this edition? Each of these newsletters publishes new pieces on LinkedIn. Subscribe, like and share to read them first, and join the conversation in the comments.



Connecting the Dots You are reading this newsletter

Where compliance, risk management, digital innovation, and bottom-line growth intersect.

Subscribe on LinkedIn

https://www.linkedin.com/newsletters/connecting-the-dots-7397392757360648192/



Letters of the Heart

Real stories, timeless wisdom, and soul-deep inspiration.

Subscribe on LinkedIn

https://www.linkedin.com/newsletters/letters-of-the-heart-7329273590376210432/



The Hope Dispatch

Uplifting news, inspiring stories, and positive change from around the world.

Subscribe on LinkedIn

https://www.linkedin.com/newsletters/the-hope-dispatch-7338912777081638914/



The HealthTech Herald

The latest breakthroughs at the intersection of health and technology.

Subscribe on LinkedIn

https://www.linkedin.com/newsletters/the-healthtech-herald-7338913774701027328/



Truth, Translated

Plain language for a complicated world: education, geopolitics, political accountability and Big Tech.

Subscribe on LinkedIn

https://www.linkedin.com/newsletters/truth-translated-7450724196525035520/

Follow the author: Dr Sukh Rashpal Sandhu on LinkedIn

BACK PAGE

About Connecting the Dots

Connecting the Dots is a LinkedIn newsletter written by Dr Sukh Rashpal Sandhu. Where compliance, risk management, digital innovation, and bottom-line growth intersect.

Read and subscribe

New editions continue to be published on LinkedIn. Search for Connecting the Dots on LinkedIn to subscribe.

Disclaimer

The views in this collection are the author's own. Pieces are reproduced as first published, and facts, figures and links are current at their original publication date. Nothing in this collection is legal, financial, medical, migration or regulatory advice.

Copyright

Copyright Dr Sukh Rashpal Sandhu 2026. All rights reserved. Images are reproduced as they appeared in the original posts.

Connecting the Dots | Dr Sukh Rashpal Sandhu | sukh@sukhsandhu.com

Masthead

Author: Dr Sukh Rashpal Sandhu

Publisher: Dr Sukh Rashpal Sandhu

Edition: 5, November 2026

Contact: sukh@sukhsandhu.com

Connecting the Dots is published on LinkedIn by Dr Sukh Rashpal Sandhu. This edition collects the pieces in the order they were first published, with their original images.