



SDA SUBMISSION TO

The Joint Select Committee on Artificial Intelligence

Date Submitted: 11 September 2026

Submitted by: Gerard Dwyer
National Secretary-Treasurer



About the SDA

The Shop, Distributive and Allied Employees' Association (the SDA) is one of Australia's largest trade unions with over 200,000 members. Our members work in retail, warehousing, fast food, hairdressing, beauty, pharmacy, online retailing, and modelling.

The majority of SDA members are women (60%), under 35 years (57%), and low-income. Retail and food services are two of the three industries with the lowest median weekly earnings. The SDA has a long history of advocating on behalf of members. We do this through enterprise bargaining; submissions on Awards and the NES to maintain a relevant safety net; and submissions to parliamentary and government inquiries and other important reviews.

The SDA has 10 policy principles that guide our engagement in these reviews. For a list of these, see Appendix A.

Contents

About the SDA	- 1 -
Contents	- 2 -
Executive Summary.....	- 3 -
Recommendations	- 5 -
Shop, Distributive and Allied Employees' Association Submission to the Joint Select Committee on Artificial Intelligence	- 9 -
Psychosocial hazards, work intensification and the AI enabled workplace	- 10 -
Biometric surveillance, consultation and social licence	- 16 -
Opportunities for resilience, productivity, competitiveness and living standards	- 17 -
Spreading the benefits of AI across Australia's suburbs and regions	- 18 -
Innovation, new industries and export opportunities, including defence industry	- 20 -
The rate and extent of AI adoption, and barriers facing small and family businesses	- 21 -
Sovereign AI capability: building, retaining and commercialising domestic capacity	- 22 -
Data sovereignty and domain-specific AI models trained on Australian data	- 24 -
Workforce and skills implications.....	- 26 -
Government adoption of AI, and transparency and accountability	- 28 -
The adequacy of Australia's existing laws and regulatory frameworks	- 30 -
Intellectual property, copyright, and Australian creative and cultural content.....	- 33 -
Risks and harms: fraud, deepfakes, and the safety of vulnerable Australians	- 34 -
National security, cyber security and critical infrastructure resilience	- 35 -
Strategic resilience: regulators, the AI Safety Institute and the security community	- 37 -
Conclusion	- 39 -
Appendix A: Principles underpinning SDA policy positions	- 42 -

Executive Summary

This submission is grounded first in the lived experience of SDA members. Our members work on the shop floor, at the checkout, in the warehouse and behind the counter, and they are already telling us what artificial intelligence looks like in practice: algorithmic rostering that removes control over their working hours, automated task allocation systems built on assumptions about how work is done rather than the reality of doing it, and workplace change introduced with little or no consultation before decisions are finalised.

Our own research reflects the scale of this problem. Among the more than 11,000 retail workers surveyed for the SDA's 2025 *Under Pressure* report, 72 per cent reported experiencing burnout and 77 per cent were assessed at high risk of psychosocial harm, with inadequate consultation on workplace change identified as one of the most significant and persistent contributing hazards. These are not abstract risks on the horizon; they are conditions our members are working under today, and they are intensifying as AI enabled systems become more deeply embedded in how retail, fast food and warehousing work is organised, monitored and evaluated.

Artificial intelligence is already reorganising jobs, reshaping decision-making, and redefining where economic, institutional and strategic power sits. Its impacts span every domain of national life. AI can lift productivity, rebuild sovereign industrial capability, and create secure, well-paid jobs across Australia. But if left to diffuse through markets alone, it also risks entrenching inequality, accelerating workplace insecurity and psychosocial harm, and deepening Australia's dependence on foreign-governed systems that underpin defence, critical infrastructure and democratic stability. None of these outcomes is predetermined. They are the product of political choices.

This submission draws on two complementary reports commissioned by the SDA from the John Curtin Research Centre (JCRC). *For All of Us: Making Artificial Intelligence Work for Working People* shows how AI is already transforming job quality, wages and workplace power. It documents the spread of algorithmic management, automated rostering, intrusive surveillance and opaque decision-making across retail, warehousing and distribution, and demonstrates that without enforceable worker voice and institutional oversight, technological change will intensify insecurity rather than lift capability. The report proposes a worker-centred national AI strategy grounded in Fair Work Commission governance, universal access to AI skills, and infrastructure that augments rather than automates.

Defending Australia: A National Artificial Intelligence Sovereign Capability and Security Agenda for the 21st Century examines the other side of the same problem: Australia's growing reliance on foreign-governed AI systems across defence, critical

infrastructure and the democratic information environment. Frontier models, cloud platforms and data pipelines are concentrated in a handful of foreign firms subject to foreign political authority. As the report shows, access to essential AI capability can be revoked, restricted or repriced by decisions made outside Australia. Sovereignty in the AI era is therefore determined not by where data centres sit, but by whether a nation can design, host, govern and secure the systems that underpin its essential functions.

Together, these reports and the SDA's own evidence make a single argument: sovereignty and fairness are inseparable. A nation that cannot govern the AI systems running its defence, energy grid or hospitals is not sovereign. A workforce that cannot see, contest or influence the algorithms that set its rosters, evaluate its performance or determine its employment, and whose health and safety are put at risk when that algorithm is wrong, is not empowered. Sovereign capability is built through people: through a workforce with the skills, authority and institutional standing to govern the systems it operates, not merely to use them.

This submission responds to each element of the Committee's Terms of Reference. It recommends that Australia:

- address the psychosocial hazards and work intensification caused by poorly implemented AI systems as a Work Health and Safety obligation, as well as an industrial relations question
- embed AI governance within the Fair Work system through a specialist Expert Panel on AI
- extend mandatory consultation, transparency and human oversight obligations to workplace AI systems, consistent with existing WHS duties to eliminate or minimise hazards before they arise
- build sovereign compute, cloud and model-assurance infrastructure across defence and essential services
- establish a National AI Sovereignty Compact and a National AI Skills Academy as tripartite institutions of government, industry and unions
- move the Commonwealth's current voluntary AI guardrails onto a mandatory, enforceable footing

Australia faces a choice. AI can strengthen national resilience, rebuild capability and expand opportunity, or it can deepen dependency, insecurity and harm. The path we take will determine whether AI works for all of us, or only for a powerful few.

Recommendations

Recommendation 1: Recognise safe staffing, sustainable workloads and algorithmic management decisions as engaging Work Health and Safety obligations where AI systems are used

– Where staffing, rostering, task allocation, productivity monitoring, recruitment, progression, promotion, discipline or termination decisions are generated or materially informed by AI systems, this should be explicitly recognised as engaging the WHS Act's primary duty to eliminate or minimise psychosocial hazards. Staffing levels should no longer be treated by industrial tribunals as a matter of unreviewable managerial discretion where the evidence shows understaffing is a foreseeable and identified psychosocial hazard. To give this duty practical effect, legislators and regulators should require that AI driven systems used for these purposes be built on realistic, repeatable and long-term sustainable measurements of work (not benchmarks derived from short, peak-period time-and-motion studies), that they factor in required break times, and that they not allocate work or make decisions based on discriminatory attributes. WHS entry permit holders should have appropriate rights to access and inspect relevant digital work systems when investigating suspected contraventions. SafeWork Australia and state and territory regulators should update relevant codes of practice to reflect these requirements specifically for AI and algorithmic systems.

Recommendation 2: Guarantee retraining, redeployment and consultation during AI-driven transition

– Industrial agreements should be required to include advance notice of AI-driven organisational change, rights to explanation for automated decisions affecting employment, and guaranteed access to retraining and redeployment pathways, with redundancy treated as a last resort rather than a default response to automation.

Recommendation 3: Establish independent monitoring of AI-attributable psychosocial harm

– SafeWork Australia and state and territory WHS regulators should collect and publicly report data on workers' compensation claims, burnout and psychosocial hazard exposure specifically attributable to AI-enabled workplace systems, so that future policy responses are grounded in sector wide evidence. The tripartite Retail Employee Safety Council model should be examined as a template for extending this monitoring function to other AI-exposed sectors.

Recommendation 4: Embed AI governance in the Fair Work system

– A specialist Expert Panel on AI should be established within the Fair Work Commission, tasked with setting standards and overseeing the use of AI in workplace relations, including rostering, surveillance, performance management, hiring and termination. Employers should be required to establish statutory AI Implementation Agreements with workers and their representatives before deploying high impact AI systems, covering consultation, job security, retraining, transparency, and privacy and data use, operating as a precondition to deployment rather than a matter left to employer discretion or bargained away through a general enterprise agreement. This framework should be underpinned by national AI legislation and supported by a National AI Authority, so that

these obligations apply consistently rather than depending on the bargaining strength of individual workplaces.

Recommendation 5: Require WHS grounded consultation before AI driven workplace change, at every affected site – The WHS Act's duty to consult, so far as reasonably practicable, before workplace decisions are made should be applied explicitly to the introduction or material alteration of AI and algorithmic systems, with risk assessments and genuine engagement of workers, health and safety representatives and unions conducted at each individual workplace rather than a single pilot site whose sign-off is then applied chain-wide.

Recommendation 6: Mandate collaborative, worker informed design of AI workplace systems – Employers deploying AI systems that materially affect task allocation, monitoring or performance evaluation should be required to engage workers and their representatives at the design and testing stage, not only at implementation, so that systems reflect the reality of how work is actually done rather than assumptions about how it is imagined to be done, reducing the introduction of unworkable or unsafe systems.

Recommendation 7: Guarantee worker data and surveillance rights – Privacy, surveillance and workplace technology protections should be updated to give workers enforceable rights to know when AI is used in decisions affecting them, to receive an explanation of automated decisions affecting their employment, and to challenge those decisions, consistent with the approach recently legislated in New South Wales' Work Health and Safety Amendment (Digital Work Systems) Act. Where an automated system contributes to a disciplinary or performance outcome, the standard it applies must be disclosed to the worker in advance, and a human decision-maker must meaningfully assess individual circumstances before any disciplinary consequence follows. Deployment of biometric surveillance technology, including facial recognition, must be preceded by genuine consultation with workers and their union, use rapid-matching systems that delete non-matching data almost instantly, and be subject to independently verifiable retention, destruction and audit requirements.

Recommendation 8: Invest in sovereign compute, cloud and data governance – The Commonwealth should invest in sovereign compute and cloud infrastructure for defence and essential services, insulated from foreign jurisdictional claims, and adopt a tiered data sovereignty classification system distinguishing data that must never leave Australian-governed infrastructure from data subject to preferred or standard hosting arrangements, with appropriate carve-outs for AUKUS and Five Eyes obligations.

Recommendation 9: Establish a National Model Assurance Facility – A sovereign, standing technical body should be established to conduct adversarial testing, continuous verification of model behaviour, and training-data auditing for AI systems used in defence and critical infrastructure, so that Australia is not required to trust foreign-governed systems it cannot independently verify.

Recommendation 10: Establish a National AI Sovereignty Compact – Government, industry and unions should establish a standing tripartite institution – a National AI

Sovereignty Compact – chaired by the Department of the Prime Minister and Cabinet, with binding rather than advisory authority. The Compact should align sovereign compute, model assurance, workforce and R&D investment across departments, set sovereign capability standards for defence, essential services and critical democratic systems, and govern procurement conditions so that firms deploying AI in critical infrastructure must demonstrate compliance with sovereign capability and labour standards as a condition of Commonwealth eligibility.

Recommendation 11: Establish a National AI Skills Academy – Government, unions, research institutions and industry should establish a National AI Skills Academy to design and deliver AI and cybersecurity training pipelines, embedding organised labour in curriculum design so that training reflects the conditions in which workers actually operate AI systems. This Academy should build on and scale the work already underway through the Future Skills Organisation, including its Skills Accelerator-AI and existing training package reforms recognised in the Commonwealth's National AI Plan, rather than duplicating this infrastructure, with formal tripartite governance giving government, industry and unions equal standing over its future direction. This should include portable, stackable micro-credentials accessible to casual and part-time workers and to workers in regional Australia.

Recommendation 12: Move the National AI Plan's guardrails onto a mandatory footing – The ten voluntary guardrails set out in the Commonwealth's Guidance for AI Adoption should be made mandatory for AI systems that materially affect employment, safety or consumer welfare, rather than remaining negotiable guidance. Voluntary compliance regimes have consistently under-delivered in comparable digital transitions, and workers should not bear the cost of a second retrofitted-regulation cycle of the kind seen with social media.

Recommendation 13: Legislate human accountability for AI generated and AI published content – Legislation should make clear that responsibility for AI generated content, including content automatically published without human moderation, rests with the platform, developer or deployer, not with the technology itself or with an end user who did not control publication.

Recommendation 14: Support regional AI adoption and small business capability – Government should fund regionally distributed AI and automation hubs linked to TAFEs, universities and industry clusters, together with targeted support to help small and family businesses overcome the cost, skills and infrastructure barriers that currently concentrate AI adoption among large firms and metropolitan centres.

Recommendation 15: Fund sovereign, domain specific AI research and development – Public R&D investment, through the National Reconstruction Fund, the Defence Industry Cooperative Research Centres program, CSIRO and university partnerships, should target domain specific sovereign AI systems trained on Australian data in areas of sectoral advantage such as energy, water and defence, with sovereign capability weighting embedded in grant conditions so that publicly funded research generates Australian-owned capability rather than subsidising foreign platforms.

Recommendation 16: Strengthen regulator and AI Safety Institute capacity – The Australian AI Safety Institute, sector regulators and the Fair Work Commission should be properly resourced to independently assess model assurance claims and identify emerging risks, rather than relying on the assurances of the firms whose systems they are assessing.

Shop, Distributive and Allied Employees' Association

Submission to the Joint Select Committee on Artificial Intelligence

The SDA welcomes the establishment of the Joint Select Committee on Artificial Intelligence and the opportunity to build on our 2024 submission to the House Standing Committee on Employment, Education and Training's Inquiry into the Digital Transformation of Workplaces.¹

Two years on from that earlier submission, AI has moved from a topic of speculative concern to a lived reality in Australian workplaces, in Australia's critical infrastructure, and in the geopolitics of technology itself. The Committee's Terms of Reference reflect that shift: they span economic opportunity and productivity; the spread of benefits across suburbs and regions; sovereign capability and data sovereignty; workforce and skills; government use of AI; the adequacy of Australia's laws; intellectual property; risks and harms including fraud, deepfakes and child safety; and national security, cyber security and critical infrastructure resilience. Few Committees have been asked to consider a technology with implications so wide.

This submission argues that these apparently distinct matters share a single underlying logic. Whether the question is who controls the algorithm that sets a warehouse worker's roster, or who controls the compute and cloud infrastructure behind Australia's defence systems, the answer determines who holds power over decisions that materially affect Australians' lives. Sovereignty and worker voice are not separate policy tracks to be traded off against one another; they are the same principle applied at different scales. A nation that has ceded governance of its essential systems to foreign platforms it cannot audit, and a workplace that has ceded governance of its labour process to an algorithm its workers cannot see or contest, have made the same mistake for the same reason: mistaking the convenience of adopting a system for the harder work of building the capability to govern it.

This submission is grounded first in the lived experience of SDA members, who are on the front line of AI, automation and technological change in Australian workplaces. That experience directly informed the SDA's own member research, *Under Pressure: The Hidden Cost of Retail* (August 2025), which surveyed more than 11,000 retail workers and found that inadequately managed technological and workplace change is a leading driver of the psychosocial harm now affecting Australian retail employees. It also informed two recent JCRC policy reports (can be read online at

¹This submission was prepared with the assistance of the John Curtin Research Centre.

the relevant footnote), both prepared with SDA input: *For All of Us: Making Artificial Intelligence Work for Working People* (April 2026),² which examines the impact of AI on jobs, wages and workplace power, including detailed case studies from retail and warehousing sectors; and *Defending Australia: A National Artificial Intelligence Sovereign Capability and Security Agenda for the 21st Century* (June 2026),³ which examines AI as a challenge to national security, critical infrastructure resilience and democratic integrity. Together, this evidence base keeps the submission anchored in the day-to-day reality of AI's introduction into Australian workplaces, not only in its long-term strategic and economic implications. Both JCRC reports converge on the same conclusion: sovereign AI capability is built not solely through infrastructure but through people. A workforce that can only operate a system, not govern it, has not achieved sovereignty in any meaningful sense.

The remainder of this submission addresses each element of the Committee's Terms of Reference in turn.

Psychosocial hazards, work intensification and the AI enabled workplace

Artificial intelligence is not only reshaping which jobs exist; it is reshaping the day-to-day psychological experience of the jobs that remain, and for many SDA members that experience has become one of chronic, escalating pressure. This is not a speculative risk for the future. It is a documented, measurable harm occurring in Australian retail workplaces now, and it is a Work Health and Safety issue with its own legal weight, distinct from and complementary to the industrial relations protections addressed elsewhere in this submission.

The scale of the problem is significant. In the SDA's 2025 survey of more than 11,000 retail members, 72 per cent reported experiencing burnout, roughly double the cross-industry benchmark, and 77 per cent were assessed at high risk of psychosocial harm, with an overall hazard exposure score of 69 out of 100, well above the threshold for high risk. 58 per cent reported very low wellbeing. Inadequate staffing was identified as the single highest scoring psychosocial hazard across every retail sub-sector surveyed, followed closely by work overload, low job control and poor change management; hazards directly connected to how technology and algorithmic systems are introduced into the workplace, and how much control workers retain once they are. 62 per cent of members rated consultation processes at their workplace as inadequate, and 42 per cent said they

² John Curtin Research Centre, *For All of Us: Making Artificial Intelligence Work for Working People*, Policy Report No. 12 (April 2026) <https://curtinrc.org/wp-content/uploads/2026/04/For-All-of-Us-Policy-Report-No12-DIGITAL-2.pdf>.

³ John Curtin Research Centre, *Defending Australia: A National Artificial Intelligence Sovereign Capability and Security Agenda for the 21st Century*, Policy Report No. 14 (July 2026) <https://curtinrc.org/wp-content/uploads/2026/07/Defending-Australia-Policy-Report-No.-14-v3.2-2.pdf>.

did not receive adequate training when new technology or working processes were introduced. These are not isolated complaints; they are a consistent pattern across supermarkets, department stores, hardware, and every other retail sub-sector surveyed.

The economic cost of this harm is also escalating, and an increasing share falls on the public system. Mental health conditions now account for 11 per cent of serious workers' compensation claims nationally, a 97 per cent increase over the past decade. The average time lost per claim has grown from 13 weeks to 34.7 weeks, and the median cost per claim has roughly tripled from \$18,600 in 2009 to \$58,715 by 2021, according to SafeWork Australia data cited in the SDA's *Under Pressure* report. Work related mental ill-health now costs the Australian economy an estimated \$12.8 billion a year, based on Mental Health Australia and KPMG analysis also cited in that report.

This escalating cost is already being reallocated rather than absorbed. Victoria's 2024 WorkCover reforms narrowed eligibility for psychological injury claims, and NSW's Workers Compensation Legislation Amendment Act 2025 introduced higher impairment thresholds from November 2025, both justified by government on the basis that psychological claims were financially unsustainable within the existing scheme. As workers' compensation tightens, the underlying need for income support does not disappear; it shifts, and superannuation-based TPD and income protection insurance are already absorbing a growing share of it, prompting insurers to redesign products, in some cases tightening the tests applied to mental health claims. The net effect is that the cost of poorly managed technological change is shifting onto workers' compensation schemes, super insurance and ultimately the public purse, rather than being borne by the businesses generating it through AI adoption. This is a genuine externality, and any assessment of AI's productivity benefits in retail is incomplete if it ignores these downstream costs. A technology that lifts measured productivity while increasing psychosocial harm has not delivered a net community benefit if it has simply relocated the cost of that harm onto public systems to which the adopting firm does not contribute in proportion to the harm it causes.

These figures predate the more recent, more intensive deployment of AI driven scheduling and monitoring systems into retail, warehousing and distribution work and there is no reason to expect this trajectory to reverse without deliberate intervention.

A recent dispute at ALDI's New South Wales distribution centres, an individual disciplinary case involving a supermarket checkout worker, and a national retailer's rollout of a new digital work system each illustrate a different facet of the same underlying problem: systems built from aggregate data or imagined workflows, deployed without genuine consultation with the people who do the work.

Case Study 1: Algorithmic benchmarking without consultation – ALDI distribution centres, NSW

ALDI implemented a national warehouse management system deployment and optimisation project, introducing new 'Picking Efficiency' productivity benchmarks generated through data analysis of warehouse operations, without consulting the workers affected or their union. Management has reported the project increased throughput by 16 per cent and reduced labour costs by 26 per cent. An SDA survey of affected workers found a different picture: significant and escalating levels of musculoskeletal strain, extreme stress and mental exhaustion, extreme fatigue, longer recovery periods, and workplace injury, coinciding directly with the new benchmarks' introduction. Workers who had met or exceeded productivity requirements for years, in some cases, found themselves unable to meet the new targets despite maintaining the same or greater effort, not because their performance had declined, but because the benchmarks were generated from data analysis rather than from consultation with the people doing the work, and consequently failed to account for order complexity, replenishment delays, traffic and stock congestion, manual handling demands, fatigue, or operational requirements such as splitting oversized loads across multiple pallets. Workers also reported the system's picking order allocation created genuine safety risks, in some cases sequencing the heaviest items last onto a pallet in a way that risked load instability during transport. This is the work-as-imagined versus work-as-done gap in concrete form: a system built from aggregate data about how warehouse work should proceed is deployed without reference to how the work is actually done, leaving workers to absorb the resulting physical and psychological cost.

The consequences compounded further with reductions to support roles such as cleaners and caretakers, introduced over the same period, creating a further layer of safety issues, spanning from basic cleanliness to the movement of empty pallets, issues that will over time erode the very efficiency gains the project was designed to produce. Overlaying all of this was an intensified monitoring and disciplinary regime where workers reported that failing to meet the new targets could trigger performance management discussions, formal warnings, or termination, reinforced by fortnightly rate-review meetings and constant performance scrutiny. A recurring theme among workers was that the system, as implemented, incentivises speed at the expense of safety.

What this shows: a productivity system built and deployed without consultation could not detect its own failings until workers began reporting harm, at which point the employer faces a costlier retrofit than genuine consultation would have required at the outset. This is entirely foreseeable, and entirely preventable.

Case Study 2: Undisclosed standards and automated discipline – supermarket checkout monitoring

A supermarket employee with more than 30 years' service recently received a final written warning for scanning inaccuracy, generated by an AI checkout monitoring system that uses overhead cameras to compare items moved through the checkout against items scanned, flagging alerts wherever it detects a possible mismatch. Those alerts are automated and do not appear to account for transaction complexity, customer interaction, or ordinary human error, and in practice they trigger an escalation cycle where the system flags a possible issue, head office visibility is triggered, store management is expected to respond, and disciplinary action becomes the easiest defensible outcome for a manager under scrutiny from above. The employee's actual performance told a different story. Her error rate was approximately 0.18 per cent (roughly one error per 533 items scanned), consistent with ordinary human variation in a high-volume, fast-paced checkout environment, and comparable to the accuracy rates advertised for AI driven robotic systems used in major distribution centres. At no point was she told what error rate the system considered acceptable, what tolerance applied, what benchmark the system used, how many flagged incidents would trigger discipline, or how alerts were assessed, verified or given context before being escalated. She was in effect disciplined against a standard that was never disclosed to her. This sits alongside a further, direct contradiction in how performance is managed: stores simultaneously track, publicise and reward scan speed, encouraging employees to work quickly, while an undisclosed monitoring system implicitly demands near-perfect accuracy, creating an unresolved conflict between the behaviour employees are incentivised toward and the standard they are actually held to.

What this shows: poor decision-making input and low job control, identified among the ten most common psychosocial hazards in retail, cannot be treated as separate from the question of algorithmic management. A worker who cannot know the standard they are being measured against, and whose disciplinary outcome is generated by a system rather than meaningfully reviewed by a person before it is acted on, has no realistic opportunity to contest, correct or improve. The hazard and the human agency failure are the same event, viewed from two angles. This is not a case for abandoning automated monitoring; it is a case for ensuring a human decision maker meaningfully assesses individual circumstances, with the applicable standard disclosed in advance, before any disciplinary consequence follows an automated alert.

Case Study 3: Where industrial consultation and WHS consultation diverge – digital work system rollout

A large national retailer introduced a new digital work system that fundamentally changed how stock processing and replenishment work was performed, moving employees away from holding responsibility for a section of the store, including stocking, presentation, restocking and customer assistance, toward completing narrowly defined tasks within system generated timeframes, with task completion

scanned and monitored for productivity. The union was initially consulted only on the industrial dimensions of the change, rostering arrangements and the timing of work, and this consultation did not identify that the changes were being driven by a new digital work system at all. It was only after the union separately became aware that the same system had been associated with significant work health and safety concerns at another retailer that work health and safety specific consultation began, by which point the decision to implement the system had already been made and the opportunity to meaningfully influence its design, implementation method or risk controls had been substantially diminished.

Following implementation, workers reported feeling micromanaged, experiencing reduced job control, and feeling overwhelmed by the new methods of work, alongside concerns about work pace and workload expectations. Long serving employees described self-doubt and a loss of confidence in their ability to perform work they had done successfully for decades. Outcomes varied significantly by store, and the pattern of that variation is itself instructive. The most positive outcomes occurred where experienced managers communicated well, listened to concerns, and exercised discretion to make practical adjustments when the system generated unrealistic expectations. Stores with poorer communication and less consultative management experienced greater difficulty, more pressure and more uncertainty, and were often less efficient despite placing greater emphasis on strict adherence to the system's requirements.

What this shows: an employer can consult in good faith on rostering and hours, genuinely believe it has met its consultation obligations, and still never trigger the specific work health and safety consultation the change actually required, because the two consultation processes ask different questions and are triggered by different people at different points in an organisation. The pattern of variation by store confirms rather than complicates the case for consultation at the design stage. Relying on individual managers to notice and informally correct a badly designed system is not a substitute for designing the system properly in the first place, and it produces exactly the uneven, manager dependent outcomes a national retailer with hundreds of stores cannot reasonably rely on as its safety strategy.

These three cases share a common thread, and this is where the Work Health and Safety framework, not only industrial relations law, must apply. Under the Work Health and Safety Act, employers hold a primary duty to eliminate or, where that is not reasonably practicable, minimise risks to workers' health and safety, and this duty explicitly extends to psychosocial hazards such as excessive workloads, unsafe staffing levels, and poorly managed change. Staffing levels have traditionally been treated by industrial tribunals as a matter of managerial prerogative, sitting outside the reach of enforceable regulation. That position is no longer tenable once staffing and workload are recognised as psychosocial hazards in their own right: an employer's WHS obligations are not satisfied by consulting on a rostering algorithm's

implementation if the staffing outcomes that algorithm produces are themselves unsafe. This distinction matters because it changes who has standing to act, and on what basis. A WHS duty to eliminate or minimise a foreseeable hazard cannot be bargained away, applies regardless of whether a workplace has an enterprise agreement in place, and is enforceable by SafeWork regulators independently of the industrial relations system. There is also a critical difference in timing between the two mechanisms. The WHS duty to consult applies before a decision to introduce or materially alter a system is made, not merely before it is implemented, and it continues as the system's impacts become apparent in practice, since a hazard that only emerges once a system is running in a live workplace still falls within the ongoing duty to eliminate or minimise it. This differs from the industrial right to consultation, which is typically triggered only once a definite decision to change has already been made, by which point the system's design, procurement and specifications are often locked in and consultation is reduced to managing the terms of a change that workers had no role in shaping. This timing difference is not only better for workers. It is also better for business and for genuine productivity, since a system designed with early input from the people who actually perform the work is more likely to reflect how work is done on the shop floor rather than how it is imagined in head office or in a consultant's report, reducing the costly cycle of deploying a system, discovering it does not fit real operating conditions, and then retrofitting or reversing it. Early consultation is not a constraint on productive AI adoption. It is a condition of it.

Both tracks are necessary and neither is a substitute for the other. Nor can an employer's own account of its staffing levels be taken at face value. In the same digital work system case, workers reported that allocated labour hours gradually decreased after rollout despite ongoing operational challenges, with salaried managers absorbing the resulting workload gap by performing additional shop floor work themselves. This masked the true labour requirement of the system and risked presenting an inaccurate picture of implementation success to anyone relying on the employer's own reporting, which is precisely why independently verifiable oversight of staffing outcomes, not employer self-assessment, is necessary to give a WHS duty of this kind any real effect.

The consultation failures documented in this evidence base are also a specific and recurring pattern, not a general grievance. Chain retailers frequently trial and consult on major technological change at a single site, then roll the resulting system out across all locations without further consultation, on the assumption that sign-off at one store constitutes sign-off everywhere. This approach is inconsistent with the basic premise of psychosocial risk management: each workplace has its own staffing profile, physical layout, customer base and operational pressures, and a system tested safe in one context is not thereby proven safe in another. Consultation that stops at implementation also leaves workers exposed when a system fails rather than when it runs as intended. At one distribution centre, workers reported that when

the automated logistics system went down, staff had no ability to switch to manual operations and were instead required to submit an IT ticket and wait for the system to be restored, with no contingency process in place for continuing operations in the meantime. A consultation process genuinely covering procurement, design, piloting and deployment would need to account for failure modes as well as normal operation, not treat system downtime as someone else's problem to fix. The WHS Act's duty to consult, so far as reasonably practicable, before workplace decisions are made was not designed with this practice in mind, but it provides the legal basis to prohibit both gaps.

Addressing this requires treating AI driven workplace change with the same seriousness the law already affords other major organisational change and extending existing WHS mechanisms rather than inventing new ones to explicitly cover algorithmic and AI enabled systems. This submission's recommendations reflect that approach, recognising unsafe staffing and unsustainable workloads generated by AI systems as an engaged WHS duty rather than a matter of managerial discretion; requiring genuine, site specific consultation before AI driven change is implemented, not after; mandating collaborative design so systems are built with workers rather than imposed on them; and establishing independent monitoring of AI attributable psychosocial harm, so that policy responses are grounded in evidence as this technology's presence in Australian workplaces continues to grow, rather than developed retrospectively once the harm is already entrenched.

Biometric surveillance, consultation and social licence

While some technologies raise legitimate workplace concerns, this submission does not oppose technology adoption in principle. Facial recognition technology illustrates why consultation, rather than blanket resistance, is the right frame for assessing new AI enabled systems in retail workplaces. The SDA supports the use of facial recognition technology in retail stores, but only for the purpose of protecting workers and customers from aggressive and violent incidents, and only where its use is governed by robust, enforceable guardrails protecting the right to freedom of association, informed consent to its use, restrictions on the use of data and the privacy of both workers and customers. Used well, the technology has a legitimate protective role, supporting Workplace Protection Orders by enabling the identification of offenders for the purpose of seeking an order, and by identifying breaches of an existing order without requiring a worker to directly confront the person responsible.

The risk is not the technology itself, but deployment without consultation, and the consequences of getting this wrong are already evident. In February 2026, the Administrative Review Tribunal ruled that Bunnings could use facial recognition technology to protect staff and customers from violence and serious crime, setting

aside some, but not all, aspects of an earlier Office of the Australian Information Commissioner finding against the company. That earlier privacy ruling and the resulting suspension of the technology's use while the company worked to rebuild trust in its deployment, were substantially avoidable. Had Bunnings consulted its workers and their union before implementation, many of the consent and transparency issues that drove the OAIC's original findings would likely have been identified and resolved in advance, rather than surfacing only after deployment and regulatory intervention.

This is a genuine social licence problem, not only a legal compliance one. A technology can be lawful and still lose the trust of the workers and customers it affects if introduced without engagement, and once that trust is lost, an employer faces a slower and more difficult path to reinstating a legitimate use of the technology than genuine consultation would have required at the outset. Protecting that social licence requires enforceable guardrails specific to biometric data, such as rapid-matching systems that delete non-matching data almost instantly rather than retaining it, regular independent auditing of how matched and retained data is used, and, as with any significant workplace technology change, genuine consultation with workers and their union before deployment, not after a regulator has already had to intervene.

Opportunities for resilience, productivity, competitiveness and living standards

Artificial intelligence offers Australia a set of opportunities that are both economic and strategic, but these opportunities will only be realised if AI is governed deliberately, embedded in strong institutions and aligned with national and worker interests.

AI can strengthen national resilience by reducing dependence on external systems and ensuring continuity of essential services. Developing sovereign capability in compute, cloud and assurance infrastructure would allow Australia to operate defence, critical infrastructure and democratic systems under Australian authority, reducing exposure to external shocks and ensuring that essential national functions remain secure and controllable.

AI can also lift productivity across the economy when deployed in ways that augment human labour rather than displace it. This opportunity is urgent: after more than a decade of shocks through a pandemic, global supply disruptions, war, inflation and now an energy supply crisis, Australia's productivity growth has stalled to its lowest level in generations, while real wages have grown even more slowly. Realising that opportunity depends on how systems are designed and deployed, not only on their

existence. Genuine productivity gains are more likely where design and deployment are cooperative and consultative from the outset, so that systems are built to improve how work actually happens on the shop floor rather than how it is imagined to happen in head office or in a consultant's report. Transparent systems, enforceable consultation and human oversight enable improvements in safety, task design and decision-making that translate into higher wages, safer workplaces and more secure jobs, rather than into deskilling and insecurity.

Australia's long-term competitiveness depends on sovereign innovation rather than passive adoption. Industry commissioned modelling suggests AI and automation could contribute up to \$2.2 trillion to the Australian economy by 2030, a scale of opportunity that will only be captured domestically if Australia builds the capability to design, adapt and govern its own systems. The current pattern points the other way: analysis of domestic data centre economics suggests that for every \$100 currently invested in Australian data centres, as much as \$70-\$80 leaves the country almost immediately through imported hardware, proprietary software licences and foreign-controlled service provision. Investment in domestic R&D, model assurance and Australian trained domain specific systems, designed to retain operational control, intellectual property and governance authority onshore, would shift the nation from being a downstream user of foreign platforms to one capable of designing, adapting and governing its own systems, turning that \$2.2 trillion opportunity into captured domestic value rather than further leakage.

AI can improve living standards when access to skills, infrastructure and capability is universal. Distributed training, portable entitlements and regionally anchored capability hubs ensure that benefits are not concentrated among large firms or metropolitan centres but shared across small businesses, regional communities and workers in all forms of employment.

Across all of these domains, the central insight is that sovereignty and fairness are inseparable. A nation that cannot govern the systems underpinning its economy and security is not sovereign, and a workforce that cannot see, contest or influence the algorithms shaping its working life is not empowered. Building sovereign capability and empowering workers are therefore mutually reinforcing strategies that create resilience, lift productivity, strengthen competitiveness and improve living standards.

Spreading the benefits of AI across Australia's suburbs and regions

Artificial intelligence can expand opportunity across Australia only if its benefits reach suburbs, towns and regions rather than remaining concentrated in capital cities or large firms. AI has the potential to lift productivity, strengthen local capability and improve service delivery in communities that have historically been underserved by technological investment. But this will not occur automatically. Without deliberate policy, AI adoption will cluster where infrastructure, capital and specialist skills are

already concentrated, reinforcing existing geographic inequalities and leaving many workers and small businesses behind.

Spreading the benefits of AI requires ensuring that people in suburbs and regions have access to the skills, infrastructure and institutional support needed to participate fully in the technological transition. Universal access to training is essential. Workers in regional Australia, outer-suburban logistics hubs and smaller retail centres must be able to acquire AI literacy and digital capability through accessible, portable and stackable learning pathways. This includes embedding AI skills in TAFEs, local training providers and community-based institutions so that workers do not have to relocate or leave employment to participate.

Infrastructure also matters. Regionally distributed AI and automation hubs, linked to local education providers and industry clusters, can give small, medium and family businesses access to tools, expertise and support that would otherwise be out of reach. These hubs can help firms adopt AI safely and productively, lowering barriers to entry and enabling local industries to innovate rather than fall behind metropolitan competitors. When combined with sovereign digital infrastructure, they ensure that regional communities are not dependent on distant or foreign governed systems for essential services.

Fair deployment of AI in workplaces is equally important. Evidence from retail, warehousing and distribution, sectors that anchor local economies and employ large numbers of workers, shows that algorithmic scheduling, automated fulfilment and intensified labour monitoring are becoming the norm, impacting regional and non-metropolitan Australia at greater scale, where a single large employer often dominates the local labour market and workers have few alternative options. Ensuring transparency, consultation and human oversight protects job quality and prevents AI from becoming a driver of insecurity in communities already facing economic pressure. When workers have a genuine voice in how AI is introduced, the technology can improve safety, reduce low-skill repetitive work and support better decision-making across retail, warehousing, logistics and regional service industries.

Finally, spreading the benefits of AI strengthens Australia's economic resilience at a regional level. AI enabled industries like logistics, energy infrastructure, advanced manufacturing, precision components and data centre operations, are already distributed across regional Australia, and directing workforce development investment to service regional labour markets converts sovereign capability building into a genuine mechanism for regional economic diversification and resilience. Strengthening local capability strengthens national capability.

The core insight is simple: AI can support broad based prosperity only when its benefits are deliberately distributed. When workers, small businesses and communities across Australia's suburbs and regions have the skills, infrastructure and authority to shape how AI is used, the technology becomes a driver of inclusion

rather than concentration. Ensuring this outcome is essential to building a fair, resilient and competitive Australia in the AI era.

Innovation, new industries and export opportunities, including defence industry

Artificial intelligence can be a catalyst for innovation across the Australian economy, creating new industries, expanding export opportunities and strengthening the nation's industrial base, including in defence. AI enables the development of advanced manufacturing, autonomous systems, logistics optimisation, data centre operations and new forms of digital services that can position Australia as a producer of high-value capability rather than a consumer of foreign platforms. When combined with sovereign compute, cloud and assurance infrastructure, AI becomes a foundation for domestic innovation rather than a channel through which value is captured offshore.

The emergence of new industries depends on Australia's ability to build and retain upstream capability. Investment in sovereign research and development, domain specific models trained on Australian data, and domestic systems integration capacity allows Australian firms to innovate in areas where the nation holds natural or strategic advantages, including energy and critical infrastructure, advanced manufacturing, precision components and defence. Existing mechanisms such as the National Reconstruction Fund and the Defence Industry Cooperative Research Centres program already provide channels for long-term capital into sovereign AI R&D. Expanding these instruments, with sovereign capability weighting built into grant criteria, would convert Australia's existing research talent into operational, exportable capability.

AI also opens new export opportunities in defence and national-security related industries, most immediately through AUKUS. The partnership's Pillar II advanced capabilities stream offers a structured pathway to co-develop sovereign defence AI systems across Intelligence, Surveillance, and Reconnaissance (ISR) analytics, autonomous platform navigation, logistics optimisation and cyber defence automation, provided co-development embeds Australian organisations at the model development layer rather than as downstream integrators of finished allied systems. Beyond AUKUS itself, democratic partners across the Indo-Pacific seeking alternatives to both US commercial dependence and Chinese platform risk represent a genuine export market for Australian capability in AI audit, assurance and sovereign cloud systems. Anchoring this to exportable governance standards and interoperable assurance frameworks would position Australia not merely as a capable AUKUS partner, but as a trusted AI governance provider to the region, strengthening collective security while building domestic industrial scale and long-term employment.

Realising these opportunities requires deliberate policy. Sovereign capability must be treated as a strategic national asset, with procurement frameworks that prioritise Australian-owned innovation and ensure that public investment generates domestic intellectual property rather than subsidising foreign platforms. Workforce development is equally important: innovation depends on a pipeline of workers with advanced technical skills, supported by institutions capable of training, accrediting and retaining talent across the economy.

The central insight is that innovation, new industries and export opportunities arise when Australia builds and governs its own technological capability. AI can be a driver of national prosperity, industrial renewal and strategic strength, but only if the systems that underpin innovation are sovereign, the workforce is skilled and empowered, and public policy ensures that the value created in Australia remains in Australia.

The rate and extent of AI adoption, and barriers facing small and family businesses

Artificial intelligence is being adopted across the Australian economy at a rapid but uneven pace. Large firms in metropolitan centres, particularly in finance, telecommunications and major retail chains, are deploying AI at scale through automated rostering, inventory optimisation, customer-service triage, algorithmic management and data-driven decision systems. But this scale of adoption is not automatically a national capability gain: evidence suggests AI adoption in large enterprises is frequently accompanied by outsourcing and offshoring of technical functions to foreign providers, hollowing out domestic technical expertise in sectors such as banking, telecommunications and public administration even as adoption accelerates. In contrast, adoption among small and family businesses is slower and more fragmented, reflecting structural barriers that limit their ability to participate in the technological transition on their own terms.

The most significant barrier is cost, and the risk is not unique to small firms. Despite between \$30-40 billion in enterprise investment in generative AI, 95 per cent of organisations had received little to no measurable increase in profits, with only around 5 per cent of firms capturing outsized gains. If well-capitalised enterprises face this level of investment uncertainty, the financial risk facing small businesses operating on thin margins, who cannot absorb the upfront expense of AI tools, cybersecurity upgrades or ongoing subscription and integration costs, is proportionately far greater. This creates a widening capability gap between large firms that can afford to absorb uncertain returns and smaller enterprises that cannot.

Skills and capability are equally important. Small and family businesses often lack access to staff with the technical expertise needed to evaluate AI tools, integrate

them safely, or adapt workflows to make effective use of new systems, and owners frequently carry responsibility for both operational and strategic decisions without institutional support. A Universal AI Skills Guarantee, ensuring publicly supported AI and digital literacy training, including portable training rights for part-time and casual workers and dedicated support via the states for sole traders and small business, would directly address this gap. Without accessible, regionally distributed training pathways, AI adoption will continue to cluster where specialist labour is already available.

Infrastructure constraints also slow adoption. Reliable connectivity, secure data environments and access to compute resources remain unevenly distributed across suburbs and regional Australia, and businesses outside major centres often face higher costs and weaker protections in deploying AI systems that rely on cloud platforms or real-time data processing. These gaps reinforce geographic inequality and limit the ability of regional businesses to compete with metropolitan firms.

Finally, governance and trust shape adoption. Small businesses are likely to be wary of opaque AI systems that affect hiring, rostering, pricing or customer engagement without clear accountability, particularly where vendors provide limited transparency or regulatory protections remain unclear. Without strong, enforceable standards for transparency, human oversight and accountability, many small businesses may continue to avoid AI tools that could otherwise support productivity and growth.

Sovereign AI capability: building, retaining and commercialising domestic capacity

Australia's ability to benefit from artificial intelligence depends on whether it can build, retain and commercialise domestic capability rather than relying on foreign governed systems. Sovereign AI capability is not simply a technological aspiration; it is an industrial, economic and strategic requirement. Nations that cannot design, host, govern and secure the systems that underpin their essential functions are exposed to external decisions they cannot influence. Building sovereign capability ensures that Australia can operate critical systems under Australian authority, maintain continuity during crises and shape the direction of technological development in its own interests.

Developing sovereign capability begins with upstream investment. Australia needs domestic compute, cloud and model-assurance infrastructure that is insulated from foreign jurisdictional claims and capable of supporting defence, critical infrastructure and high-value industries. A dedicated Sovereign Cloud and Compute Program, co-invested by government in coordination with relevant agencies, prioritising Australian operational control, security assurance and long-term supply resilience for strategic workloads, would provide the foundation for local innovation, allowing Australian

researchers, firms and public institutions to build and test advanced systems without relying on external platforms whose access, pricing or functionality can be altered by foreign political decisions.

Retention of capability requires a skilled workforce and institutions that can support long-term development. In the near term, this will likely require targeted skilled migration pathways for cleared AI engineers, systems integrators and model assurance specialists, industry partnership programs that embed Australian workers in sovereign capable domestic firms and streamlined security clearance pipelines. Over the longer term, a national sovereign AI research network coordinating investment across public research agencies, university AI institutes and relevant government science and technology bodies would create the institutional spine for sustained R&D, one designed to produce not just research output but trained sovereign practitioners. Workers need the authority and institutional backing to govern the systems they operate, ensuring that capability is embedded not only in infrastructure but in people. Without a domestic workforce capable of understanding, challenging and adapting AI systems, sovereign capability cannot be sustained.

Commercialisation is equally important. Australia must be able to turn domestic capability into exportable products and services, particularly in sectors where the nation holds strategic or comparative advantages, including energy and critical infrastructure, advanced manufacturing and defence. Trusted AI systems for autonomous platforms, logistics optimisation, cyber resilience and situational awareness are increasingly sought by democratic partners, including through AUKUS, seeking reliable suppliers. Realising this requires intellectual property and data governance to be embedded in R&D funding conditions from the outset. Commonwealth funded AI research should retain ownership of model weights, training data and architectural innovations within Australia, and where partnerships with international firms are necessary, agreements must preserve Australia's right to audit, adapt and redeploy model outputs without seeking foreign permission. Commercialising these capabilities strengthens Australia's industrial base, creates high-skill jobs and contributes to allied security.

Achieving sovereign AI capability requires deliberate policy choices. As previously stated, existing instruments such as the National Reconstruction Fund and the Defence Industry Cooperative Research Centres program offer channels for long-term capital that should be expanded to fund AI capability development, with sovereign capability weighting built into grant criteria. This public investment should not be treated as an industrial question alone: consistent with the government's own approach through Future Made in Australia, sovereign AI funding should be tied to social aims, ensuring new technologies are held to the same worker-focused standards as established industries rather than allowing capability gains to come at the cost of job quality or workplace consultation. Regulatory institutions must have the resources and authority to assess and verify AI systems independently.

The central insight is that sovereign capability is built through people, infrastructure and institutions working together. When Australia can build, retain and commercialise its own AI systems, it strengthens national resilience, expands economic opportunity and ensures that the value created in Australia remains in Australia. Sovereign AI capability is therefore not only a technological goal but a foundation for economic security, strategic autonomy and shared national prosperity.

Data sovereignty and domain-specific AI models trained on Australian data

Artificial intelligence can only strengthen Australia's economy and national resilience if the data that underpins it is governed, stored and processed under Australian authority. Data sovereignty is fundamental to ensuring that essential systems remain secure, reliable and aligned with national interests. Foreign-hosted AI systems can be subject to extraterritorial legislation such as the US CLOUD Act or China's National Security Law, both of which allow overseas authorities to compel access to data stored or processed through foreign owned infrastructure, regardless of where that data physically resides. Critically, this exposure is not resolved simply by storing data onshore, as even where operators store data in Australian facilities under Australian law, the underlying cloud platforms and model-hosting environments may remain governed by foreign legal frameworks that can compel disclosure or access without Australian consent.

Breaking this dependency requires enforceable legal protections, not just physical data location. This is best achieved through a tiered data sovereignty classification system. Data supporting defence, critical infrastructure or essential services should be classified as requiring Australian governed infrastructure and must never leave Australian jurisdiction under any circumstances. A second tier of sensitive but lower risk data should be subject to preferred hosting arrangements, requiring Australian government approval but allowing limited offshore processing under strict conditions. A third tier of standard commercial data can continue under existing arrangements. This classification system should include appropriate carve outs for data sharing obligations under AUKUS and Five Eyes arrangements, since Australia's alliance commitments require some degree of interoperability with partner intelligence and defence systems that a blanket restriction on offshore data would make impossible to honour. The purpose of these carve outs is not to weaken the classification system but to ensure it reflects the reality that sovereignty and alliance cooperation are not the same question, and conflating them risks either compromising genuine sovereign control or breaching commitments Australia has already made to its closest security partners.

Data sovereignty is also essential for building domain specific AI models that reflect Australian conditions, institutions and priorities. In energy, this means sovereign models for grid balancing, distributed energy resource management and demand forecasting, trained on Australian grid data and maintained under Australian jurisdictional control. In water, it means sovereign anomaly detection and treatment optimisation systems designed for Australian infrastructure conditions. In defence, it means the ISR analytics, logistics models and cyber defence automation tools that underpin sovereign military capability. Generic global models cannot reliably support these functions, and reliance on them risks embedding assumptions, behaviours and vulnerabilities that are misaligned with Australian conditions. Domain specific models trained on Australian data allow the nation to develop capability that is accurate, trusted and tailored to local requirements. These systems do not require frontier scale compute, but rather domain expertise, sovereign data pipelines and sustained institutional commitment.

Developing these models requires sovereign infrastructure capable of hosting training data, running high-performance compute workloads and conducting independent assurance. Sovereign cloud and compute environments ensure that training pipelines, model weights and operational data remain under Australian jurisdiction, protected from foreign access or interference. Assurance frameworks allow Australia to test, verify and audit model behaviour without relying on vendor claims or opaque foreign processes. Together, these capabilities form the foundation for trusted AI systems that can be deployed in critical national functions.

Commercialisation depends on the same foundations. When Australia can build domain specific models trained on Australian data, it can export capability in areas where it holds natural or strategic advantages, including energy, water infrastructure and defence. Trusted AI systems for autonomous platforms, logistics optimisation and cyber resilience are increasingly sought by partners who require reliable, transparent and democratically governed technology. Sovereign data and sovereign models allow Australia to participate in these markets as a producer of high-value capability rather than a downstream consumer of foreign platforms.

The central insight is that data sovereignty and domain specific models are not technical preferences; they are strategic necessities. They ensure that Australia can govern the systems that underpin its economy and security, build capability that reflects local needs, and commercialise innovation in ways that strengthen national resilience. Without sovereign data and sovereign models, Australia remains dependent on external systems it cannot fully control. With them, Australia can build a secure, competitive and future ready AI ecosystem.

Workforce and skills implications

Artificial intelligence is reshaping the skills required across the Australian economy, altering job design, changing workplace power dynamics and creating new demands for training, governance and institutional support. The workforce implications are profound, and evidence suggests this transition differs from previous waves of automation. Unlike earlier technological shifts that primarily displaced routine or low-skill labour, research indicates AI poses significant risk to workers with higher levels of education and specialised skills. This was reflected in this Parliament's own prior scrutiny, with the Select Committee on Adopting Artificial Intelligence's November 2024 final report identifying a consultation gap in AI implementation, expanding workplace surveillance with insufficient privacy protection, a critical AI skills shortage with regional and lower-income workers most at risk of displacement, and a regulatory vacuum around AI driven hiring, firing and performance decisions. AI is already influencing how work is organised in retail, warehousing, logistics, transport, finance, telecommunications and essential services, and these findings indicate the changes required are structural, not incremental.

The most immediate implication is the need for universal access to digital and AI skills. Workers across all industries including casual, part-time and regional workers, must be able to acquire foundational AI literacy and advanced technical skills through accessible, portable and stackable training pathways. A Universal AI Skills Guarantee, ensuring publicly supported AI and digital literacy training, portable pro-rata training rights for part-time and casual workers, paid time for employer-directed training, and dedicated support for sole traders and small business, would deliver this directly. Delivering this guarantee requires an institution capable of designing and running it, not just a funding commitment. Much of this foundation already exists.

The Future Skills Organisation, recognised in the Commonwealth's National AI Plan, is already delivering national AI skills infrastructure through its Skills Accelerator-AI and related training package reforms, and this work should be built on and scaled rather than duplicated. Government, unions, research institutions and industry should establish a National AI Skills Academy to coordinate and scale this existing work under formal tripartite governance, with organised labour given equal standing in shaping how it develops so that training reflects the conditions in which workers actually operate AI systems, rather than a generic technical curriculum designed without reference to real workplace practice.

Without a body of this kind providing that governance, training risks becoming fragmented across individual employers and providers, each teaching a different version of AI literacy with no consistent standard and no guarantee that casual, part time or regional workers can access it on equal terms. Without a functioning skills pipeline, AI adoption will deepen inequality by concentrating opportunity among workers who already have access to training, while leaving others exposed to deskilling, displacement or diminished bargaining power. Skills development must be

embedded in institutions that workers trust, including TAFEs, unions, community-based providers and industry partnerships that reflect real workplace conditions, with the Academy coordinating this network, building on FSO's existing infrastructure, rather than duplicating it.

AI also changes the nature of workplace governance. As automated systems increasingly influence hiring, rostering, performance evaluation and termination decisions, workers need enforceable rights to transparency, explanation and human oversight. A requirement for employers to establish AI Implementation Agreements with workers and their representatives before deploying high-impact AI systems, covering hiring, rostering, performance management, surveillance, discipline or termination, would formalise this, guaranteeing early notification, worker input into system design, clear information on data inputs and decision-making parameters, and meaningful human review of significant automated decisions. Consistent with this approach, the SDA's memorandum of understanding with Microsoft already recognises the right of union delegates to represent workers and to receive information relevant to that representation. An AI Implementation Agreement should extend this principle as a statutory baseline across the retail sector, rather than leaving it dependent on the terms individual employers are willing to negotiate. An AI Implementation Agreement should operate as a statutory precondition to deployment, sitting alongside rather than replacing any existing enterprise bargaining agreement. This means an employer cannot deploy a high-impact AI system merely because a general enterprise agreement is silent on the matter and cannot use enterprise bargaining as a vehicle to negotiate away the specific protections an AI Implementation Agreement is designed to guarantee. Where an enterprise agreement already contains AI specific terms, the AI Implementation Agreement process should build on and strengthen those terms rather than duplicate or undercut them, but the statutory requirement to consult and agree before deployment should apply regardless of whether an enterprise agreement exists or what it says. Without clear rules of this kind, AI can become a mechanism for opaque decision-making that undermines job security and erodes workplace agency.

This is not a hypothetical concern. A Kmart Australia worker described an automated resume filtering system used in recruitment for around two years, an automated rostering system that schedules staff during hours they have specifically marked as unavailable, requiring managers to manually correct the roster afterwards, and a digital stock system that can fail to flag missing products for weeks at a time, leaving workers to identify and correct the gap manually. The worker also reported that, despite the company's public partnerships with named generative AI providers, staff are rarely told directly when AI is being used in the systems that govern their work, and that technological changes are generally implemented through top-down decisions with limited opportunity for shop floor feedback and uncertain union involvement. None of this reflects malicious intent on the part of the employer. It reflects the absence of any legal requirement to disclose when AI is in use, to consult

before it is introduced, or to ensure the systems deployed actually match the reality of how the store operates day to day. An AI Implementation Agreement, properly enforced, would require each of these gaps to be addressed before deployment rather than discovered by workers afterwards.

The workforce implications extend to national capability. Australia needs a pipeline of advanced technical talent through engineers, data scientists, cybersecurity specialists, systems integrators and technicians, capable of building, operating and securing sovereign AI systems. This requires long-term investment in education, apprenticeships, research pathways and mid-career transitions that allow workers to move into high-value roles as industries evolve, alongside targeted skilled migration and clearance ready training pipelines to address acute near-term shortages of cleared AI engineers and model assurance specialists. Workforce capability is not only an economic issue but a strategic one, as without domestic expertise, Australia cannot maintain sovereign control over critical systems or participate effectively in emerging industries.

Finally, AI driven change must be managed through institutions that support fair transition. A dedicated AI in Industrial Relations Taskforce, jointly led by government with formal participation from unions, employer representatives and sole-trader bodies, should oversee this transition, including a comprehensive review of the Fair Work Act to address the risks posed by AI, algorithmic management and automated decision-making, with terms of reference designed to prevent the process becoming a vehicle for employer-led deregulation. These protections should not be left to individual negotiation or employer discretion. Industrial agreements should be required to include advance notice of AI driven organisational change, rights to explanation for automated decisions affecting employment, and guaranteed access to retraining and redeployment pathways, with redundancy treated as a last resort rather than a default response to automation. Making these terms a required inclusion, rather than a matter for case-by-case bargaining, ensures workers in workplaces with weaker bargaining power are not left with lesser protection than those in workplaces where a union has successfully negotiated these terms already.

Government adoption of AI, and transparency and accountability

Artificial intelligence is becoming embedded across government operations, from service delivery and regulatory enforcement to infrastructure management and national security. As public agencies adopt AI systems to improve efficiency, manage complex datasets and support decision-making, the stakes for transparency and accountability increase. Government use of AI carries unique responsibilities because public institutions exercise coercive power, administer rights and entitlements, and operate systems that citizens cannot opt out of. The consequences of getting this wrong are not hypothetical: the Robodebt scheme illustrated what

happens when automated decision-making in high-stakes public administration operates without adequate human oversight, explainability or accountability, and despite the Robodebt Royal Commission's 2023 recommendations for legislative reform and an independent audit body, no such framework has since been legislated. The pattern has already repeated with the Commonwealth Ombudsman finding that the Department of Employment and Workplace Relations unlawfully terminated income support payments for nearly a thousand jobseekers using an incorrectly coded automated system. Ensuring that AI strengthens rather than undermines public trust requires clear governance, independent oversight and enforceable transparency obligations.

The most immediate requirement is visibility. When AI systems influence decisions about welfare eligibility, compliance actions, visa processing, public housing allocation, emergency-response coordination or critical-infrastructure operations, affected individuals must be able to understand how those decisions were made. This is not a new concern for this Parliament: the Joint Committee of Public Accounts and Audit's *Proceed with Caution* inquiry into AI governance across Commonwealth departments recommended whole-of-government working groups, consistent governance frameworks and clearer accountability and risk-mitigation mechanisms as AI systems are deployed in public service operations, recommendations that should inform this Committee's own findings. Opaque or unexplainable systems erode procedural fairness and can produce errors that are difficult to detect or challenge. Transparency must therefore extend beyond technical documentation to include mandatory plain-language explainability at the point of decision, and meaningful avenues for review.

Accountability is equally essential. Government agencies must remain responsible for the outcomes of the AI systems they deploy, including those procured from external vendors. Automated decisions cannot be treated as neutral outputs of technology; they are administrative actions with legal and social consequences. Human-in-the-loop mechanisms must function as substantive safeguards rather than formalities, ensuring public servants retain the authority and capacity to interrogate, override or contextualise algorithmic outputs, backed by independent auditing of model performance, bias and outcomes and enforceable appeal and redress rights. Clear liability frameworks are needed to determine responsibility when harm arises, particularly in systems involving multiple actors such as vendors, developers and government agencies.

Government adoption also requires independent assurance specific to public administration. Public institutions must be able to test, verify and audit the AI systems they deploy without relying solely on vendor claims or proprietary processes, with oversight shared between the Attorney-General's Department, the Commonwealth Ombudsman and the Office of the Australian Information Commissioner. This assurance function should sit alongside, and draw on, whatever

independent technical audit capability Australia establishes to close its broader AI regulatory gaps, a question addressed in the following section on the adequacy of Australia's existing laws and regulatory frameworks.

Finally, transparency and accountability are foundational to democratic integrity. Election administration, parliamentary communication, public service announcements and government digital services must operate on infrastructure governed by Australian law as foreign governed cloud platforms introduce exposure to foreign legal claims and political decisions that cannot be mitigated through contractual or technical safeguards alone. As synthetic media becomes more sophisticated, sovereign verification capability becomes essential: national digital signature infrastructure allowing public officials to authenticate their own statements, AI enabled detection tools for synthetic media, and public facing verification portals would allow citizens to distinguish authentic government communication from fabricated content. Public institutions must model best practice in AI governance, setting standards that apply across the economy.

The central insight is that government adoption of AI must be held to a higher standard than private-sector deployment. Public trust depends on transparency, accountability and human oversight. When government systems are explainable, contestable and independently verified, AI can improve service delivery, strengthen institutional integrity and support democratic resilience. Without these safeguards, government adoption risks repeating past failures and undermining confidence in public administration.

The adequacy of Australia's existing laws and regulatory frameworks

Australia's existing laws and regulatory frameworks were built for an earlier technological era and are not adequate for the scale, speed and structural impact of artificial intelligence. The government's own December 2025 National AI Plan is a start, but its market neutral approach is not sufficient: markets do what they can get away with, not what the public interest demands, and without firm guardrails the economic and social costs of inadequate regulation will fall on workers and consumers first. While current legislation provides partial coverage of issues such as privacy, workplace relations, consumer protection and critical infrastructure security, these frameworks were designed around human-centred decision-making and conventional digital systems. They do not provide the transparency, accountability or sovereign control required when automated systems influence employment, safety, essential services and national security.

The most significant gap is not a total absence of consultation obligations, but that the obligations which do exist are limited, inconsistent and often triggered too late to be meaningful. Employers already have legal consultation obligations under existing Termination, Change and Redundancy (TCR) provisions in modern awards and

enterprise agreements, which are engaged where a decision is made that is likely to have a significant effect on employees, including changes to how work is performed. Where AI systems change jobs or the way work is done, this would generally fall within the scope of these existing change consultation obligations, commencing once a decision to introduce the change has been made and before it is implemented. In practice, however, these obligations are rarely applied with any rigour to AI and automated systems. Their coverage depends on which award or agreement applies to a given workplace, meaning protection is inconsistent across the retail sector rather than universal, and because the duty to consult is triggered only once a decision to change has already been made, workers and their representatives have no legal right to be consulted on whether or how a system is designed before that decision is locked in. Existing workplace laws also do not guarantee workers the right to know when AI is used in a decision affecting them, to receive an explanation of an automated decision, or to challenge that decision, protections that fall outside what TCR provisions were designed to address. As discussed earlier, a dedicated review of the Fair Work Act, with terms of reference designed to prevent the process becoming a vehicle for employer-led deregulation, is needed to assess and close these gaps, both by ensuring existing change consultation obligations are actually applied to AI driven change in practice, and by legislating the transparency and explanation rights that fall outside their current scope. Without these protections, AI can be used in ways that undermine job quality, workplace safety and procedural fairness.

Regulation of AI generated content and platform operated systems shows the same enforcement pattern repeating across sectors. Australia's media regulator can request information from platforms but cannot independently audit algorithmic systems, access model architectures, or compel technical evidence of compliance, and its own oversight reports have found that under the voluntary disinformation code, platforms self-report against metrics they have largely designed themselves, with no independent verification of whether commitments are being met. The Misinformation and Disinformation Bill 2024, which would have introduced enforceable co-regulatory obligations, was abandoned before passing. Voluntary guardrails and industry self-regulation have consistently under-delivered in comparable digital transitions, and there is no evidence they will be sufficient for AI. The same problem applies to the ten voluntary guardrails set out in the Commonwealth's Guidance for AI Adoption. As currently designed, compliance is a matter for each organisation to determine for itself, with no regulator empowered to assess whether a given deployment actually meets the guardrails, no penalty for a deployment that fails to meet them, and no independent mechanism to verify the self-assessments organisations report. Converting these guardrails into enforceable requirements does not require inventing new regulatory machinery. For AI systems that materially affect employment, safety or consumer welfare, compliance with the guardrails should become a legal obligation enforced by the regulator with existing

jurisdiction over that sector, such as the Fair Work Commission for workplace deployments or the Australian Competition and Consumer Commission for consumer facing systems, with the Australian AI Safety Institute providing the technical assessment those regulators currently lack the capacity to perform themselves. Deployments found not to meet the guardrails should be subject to the same order making and penalty powers the relevant regulator already holds for comparable non-compliance in its sector, rather than a new and separate enforcement regime being built from scratch.

Critical infrastructure and national security legislation provides important safeguards but does not fully address the risks created by dependence on foreign governed AI systems. Existing frameworks such as the Security of Critical Infrastructure Act were not designed with AI model assurance in mind. Aligning sovereign AI obligations with these frameworks, rather than assuming they already cover AI-specific risk, is an immediate task, not a completed one. This leaves Australia vulnerable to external jurisdictional claims, vendor lock-in and supply-chain risks that cannot be mitigated through traditional regulatory tools alone.

Finally, regulatory institutions themselves lack the resources, technical capability and statutory authority needed to oversee AI effectively. Bodies such as the Fair Work Commission and the Australian Competition and Consumer Commission face rapidly expanding responsibilities without corresponding increases in funding or technical expertise, a gap other jurisdictions are already addressing; in Canada, for instance, courts provide a fund helping parties in AI related cases access expert evidence, so that courts are properly informed rather than reliant on whichever party can afford it. An Artificial Intelligence Accountability Act, establishing statutory definitions of developers, deployers and catastrophic risk, mandatory human oversight for high-risk systems, and annual independent reviews by the Australian AI Safety Institute, paired with a dedicated sovereign technical audit capability modelled on France's Pôle d'Expertise de la Régulation Numérique, where regulators retain legal authority while a dedicated forensic body performs the technical work, would convert regulatory intent into demonstrable, enforceable accountability across both platform governance and public administration.

Well-designed, clearly scoped regulation introduced early is also more likely to support investment than to deter it. Investors and firms planning long-term deployment of AI systems in workplaces or critical infrastructure need to know what standards they will be held to, so they can build compliant systems from the outset rather than retrofitting them once a regulatory gap is eventually closed. The absence of clear rules does not leave firms free to invest with confidence; it leaves them exposed to the uncertainty of eventual, potentially more disruptive regulatory correction, as has occurred in other digital transitions where voluntary regimes were replaced only after harm had already accumulated. A predictable, enforceable framework, settled early, gives industry a stable basis on which to plan capital

investment, workforce development and system design, rather than deferring decisions in anticipation of regulation that is known to be coming but not yet defined.

The central insight is that Australia's current legal and regulatory frameworks are not fit for purpose in an AI enabled economy, and no existing overseas model can simply be imported wholesale: the US model's light-touch approach risks entrenching market concentration and labour precarity, while the EU's dense regulatory architecture is difficult to replicate within Australia's federal system and smaller technology sector. What Australia needs is a distinctively Australian framework, one that leaves workers with enforceable protections, consumers with genuine transparency, regulators with real technical authority, and the nation with sovereign control over critical systems. Modernising these frameworks is essential to ensuring that AI strengthens Australia's economy, democracy and national resilience rather than undermining them.

Intellectual property, copyright, and Australian creative and cultural content

Artificial intelligence poses significant challenges for Australia's intellectual property system, the protection of creative work, and the sustainability of Australian cultural industries. Existing copyright and IP frameworks were designed for human created content and conventional digital technologies, not for generative systems capable of producing text, images, audio, video and design outputs at industrial scale. As AI becomes embedded across creative, cultural and commercial sectors, Australia must ensure that its laws protect creators, uphold cultural sovereignty and maintain the economic viability of domestic creative industries.

The most immediate issue is the use of Australian creative content in training data. Many AI systems are trained on large, uncurated datasets that include copyrighted material, journalism, literature, music, film and other creative works. Without clear rules governing how Australian content can be used, creators risk having their work absorbed into training pipelines without consent, compensation or attribution. This undermines the economic foundations of creative labour and erodes the cultural distinctiveness of Australian content in global markets. Ensuring that training data is collected lawfully, transparently and with enforceable rights for creators is essential to maintaining a fair and sustainable creative ecosystem.

Generative AI also challenges traditional notions of authorship and ownership. When AI systems produce content that resembles or directly draws from existing works, questions arise about who owns the output, who is responsible for its publication, and how derivative or infringing content should be regulated. Clear liability rules are needed to ensure that responsibility rests with the platform, developer or deployer, not with end users who did not control training data or publication. Without such

rules, creators face uncertainty about how their work can be protected, and consumers face risks associated with misleading or unauthentic content.

Australian cultural sovereignty is also at stake. Creative and cultural content plays a central role in shaping national identity, democratic discourse and social cohesion. If AI systems trained primarily on foreign data dominate Australian cultural spaces, local voices, stories and perspectives risk being overshadowed or homogenised. Supporting Australian trained models, Australian datasets and domestic creative industries ensures that AI systems reflect Australian values, histories and cultural priorities rather than reproducing external narratives or biases.

Economic implications are equally significant. Creative industries contribute billions to Australia's economy and employ large numbers of workers across publishing, media, design, music, film, gaming and digital production. If AI systems can replicate or replace creative outputs without compensating creators, the viability of these industries is threatened. Strong IP protections, transparent licensing frameworks and enforceable rights to attribution and remuneration are essential to ensuring that AI supports rather than undermines creative work.

Finally, regulatory institutions must have the authority and capability to oversee AI's impact on intellectual property and cultural content. This includes the ability to audit training datasets, enforce copyright compliance, regulate automated content publication and protect Indigenous cultural knowledge from unauthorised use. Without robust oversight, Australia risks losing control over how its cultural assets are used, reproduced and commercialised.

The central insight is that intellectual property and cultural content are not peripheral issues in the AI transition; they are foundational to Australia's cultural sovereignty, economic resilience and democratic integrity. Ensuring that creators are protected, that Australian content is respected, and that AI systems operate within a clear and enforceable legal framework is essential to building a fair, secure and culturally confident AI ecosystem.

Risks and harms: fraud, deepfakes, and the safety of vulnerable Australians

Artificial intelligence introduces a new class of risk to Australia's information environment, particularly through deepfakes and synthetic media capable of undermining trust in institutions, public communication and democratic processes. Synthetic audio, video and fabricated documents can be used to impersonate public figures, distort civic discourse and undermine confidence in authentic communication. Personalisation and automated content generation allow this material to be produced and distributed at a scale and speed that outpaces the ability of government and platforms to respond. In high-stakes contexts such as

elections, emergency communications and public health messaging, this creates a genuine risk of manipulated or fabricated content displacing authentic information before it can be identified and corrected.

Australia currently lacks sovereign capability to address this threat. Detection of deepfake video, synthetic audio, AI generated images and coordinated synthetic identity activity cannot rely solely on foreign platforms, which may prioritise their own commercial or political interests over Australian public safety. Dedicated sovereign detection capability is required including cross-platform monitoring for coordinated manipulation, rapid response mechanisms for disinformation incidents, and public-facing verification tools, so that Australia is not dependent on offshore systems to identify harms occurring within its own information environment.

Authenticity infrastructure is central to closing this gap. Public officials and institutions must be able to authenticate their own statements, video, audio and public announcements through mechanisms such as cryptographic signatures and secure identity verification, paired with tamper resistant communication channels. Without this, impersonation and synthetic attacks can and will undermine institutional credibility, at a pace that current verification and correction processes are not equipped to match.

The central insight is that AI enabled deepfakes and synthetic media are not a distant or hypothetical risk but an active, accelerating one, and Australia's exposure stems as much from a lack of sovereign detection and verification capability as from the technology itself. Building that capability rather than relying on foreign platforms to self-police is a precondition for maintaining institutional trust and democratic resilience in an AI enabled information environment.

National security, cyber security and critical infrastructure resilience

Artificial intelligence is rapidly becoming embedded in the systems that underpin Australia's national security, cyber security and critical infrastructure. Defence platforms, intelligence analysis, emergency management, energy distribution, transport networks and communications systems increasingly rely on AI enabled tools to coordinate complex operations in real time. Built and governed under Australian authority, this same capability represents a genuine opportunity to protect Australians: AI enabled anomaly detection, behavioural analysis and automated threat containment can operate at the machine speed now required to counter threats that themselves move at machine speed, offering faster detection, more resilient infrastructure coordination and stronger protection of essential services than traditional, human paced security models allow. Realising that opportunity, rather than being exposed by the same technology, depends on how deliberately Australia governs the systems involved.

The most fundamental challenge is structural dependence on foreign governed AI systems, and this is not a hypothetical risk. In April 2026, the US Department of Defense sought unrestricted access to Anthropic's frontier AI models for applications including mass domestic surveillance and autonomous weapons development. When the company refused, the US Government designated it a supply chain risk and ordered all federal agencies to cease using its systems immediately, while firms that complied retained their federal integration pathways. The episode demonstrated a structural truth with direct relevance to Australia: access to frontier AI capability can be revoked, restricted or repriced by foreign political decision, with no recourse for dependent states. Sovereign control over compute, cloud and model assurance infrastructure is therefore essential to ensuring that Australia can operate critical systems under Australian authority, maintain continuity during crises and protect sensitive national information.

Foreign interference is a distinct and escalating risk in its own right. State-based cyber and intelligence operations already target Australian systems: Russian and Chinese cyber espionage networks have repeatedly penetrated government systems, critical infrastructure and private organisations for intelligence collection and strategic positioning, while North Korea and Iran use cyber enabled theft and disruption to offset conventional capability gaps. AI accelerates all of these activities, enabling adversaries to analyse large datasets, generate highly targeted phishing campaigns, craft malware that adapts to defensive measures, and conduct large-scale coordinated attacks at a speed and scale that overwhelms traditional perimeter-based security. Defending against this requires sovereign detection and response capability that does not itself depend on the foreign governed platforms the threat is often directed through.

Cyber security risks are further amplified by the increasing integration of the critical systems AI coordinates. Critical infrastructure sectors are becoming more tightly coupled: energy grids depend on telecommunications networks, water systems depend on energy, and transport systems depend on both, with AI intensifying these interdependencies by coordinating operations across domains in real time. A vulnerability in one sector can now cascade into others, turning a localised failure into a multi-sector disruption. Defensive AI systems are also, themselves, potential targets: if adversaries compromise the models underpinning cyber defence, they can blind detection systems, manipulate alerts or trigger false positives that disrupt operations rather than prevent disruption. Ensuring resilience requires embedding cyber security into infrastructure planning, mandating independent verification of AI systems used in essential services, and building a national AI red-team capability able to participate in allied red-teaming and threat testing programs under AUKUS Pillar II.

Critical infrastructure resilience also depends on the integrity of the data and models that govern operational systems, and Australia currently has limited visibility into this.

Frontier AI systems are built on globally dispersed supply chains involving closed training data, proprietary model architecture and opaque update cycles, and Australia cannot independently verify how these systems are trained, secured or updated. Manipulated training data, compromised model weights or adversarial inputs can distort decision-making in ways that are difficult to detect or attribute. Without a national model assurance facility capable of adversarial testing, continuous verification and model behaviour auditing, Australia remains reliant on vendor claims rather than independent evidence that the systems governing its critical infrastructure are behaving as intended.

AI also affects national security through the information environment. Automated content generation, deepfakes and synthetic media can be used to manipulate public discourse, undermine trust in institutions, disrupt emergency communication and interfere with democratic processes. These risks escalate during crises, elections or geopolitical tension, when rapid and reliable information is essential. Workers in telecommunications, media and platform adjacent industries have direct operational visibility into how content is moderated and how platform governance obligations are met in practice rather than on paper, and structured channels for reporting algorithmic anomalies or manipulation patterns to relevant authorities would convert that frontline knowledge into a genuine early-warning function that no regulatory body currently has equivalent access to.

The central insight is that national security, cyber security and critical infrastructure resilience are now inseparable from sovereign AI capability. Australia must be able to design, host, govern and secure the systems that underpin its essential functions. Without sovereign infrastructure, independent assurance and a skilled domestic workforce, AI becomes a vector for strategic dependency, foreign interference and systemic vulnerability. With them, AI can strengthen national resilience, enhance defence capability and ensure that critical systems remain secure, reliable and under Australian control, turning AI from a source of exposure into a genuine tool for protecting Australians.

Strategic resilience: regulators, the AI Safety Institute and the security community

Artificial intelligence is now embedded across the systems that underpin Australia's economy, democracy and national security. As these systems become more complex, more automated and more interdependent, strategic resilience depends on institutions capable of governing AI with authority, technical depth and national coordination. Regulators, the AI Safety Institute and the broader security community form the backbone of this capability. Their effectiveness determines whether Australia can anticipate emerging risks, respond to crises, and maintain sovereign control over critical systems.

Regulators face rapidly expanding responsibilities without a corresponding expansion of authority. In workplace governance, the Fair Work Commission would need to be expanded to integrate AI oversight within Australia's industrial relations system directly, supported by a dedicated Expert Panel on AI with statutory authority to provide independent advice, oversee complaints and exercise binding dispute-resolution powers over workplace AI and algorithmic management. In platform and information governance, ACMA's mandate would need to extend beyond its current information request powers to require platforms to disclose how recommendation systems operate, allow independent audit of high-risk algorithms, and meet mandatory, timely and enforceable reporting obligations covering coordinated inauthentic behaviour and foreign interference attempts. In critical infrastructure, regulators must ensure that AI systems used in energy, water, transport and communications are safe, verifiable and resilient. Without statutory authority and technical capability to match these responsibilities, regulators cannot assess advanced models, audit training data or challenge vendor claims, and Australia will remain dependent on voluntary disclosures shaped by foreign commercial and political incentives.

The AI Safety Institute plays a central role in national resilience, and its role should be formalised rather than assumed. The Institute should be legislated to conduct annual independent reviews of developers' risk management systems and the efficacy of human oversight mechanisms, with findings reported publicly and with provisions enabling workers to review and provide input on system performance and impacts – a function that gives the Institute both technical depth and a direct line to how AI systems actually behave in workplaces, not just in laboratory testing. Complemented by mandatory annual reporting from developers and deployers under a legislated accountability framework, this would allow the Institute to function as a genuine national early warning system, identifying emerging risks before they escalate into systemic failures, rather than relying on vendor assurances and foreign processes that may not align with national interests.

The security community, including defence, intelligence, cyber agencies and emergency-management organisations, must integrate AI governance into their operational frameworks. AI enabled threats such as automated cyber intrusion, synthetic media manipulation, supply-chain compromise and model level interference require coordinated national responses. Security agencies need the ability to verify the integrity of AI systems, detect adversarial manipulation, and maintain operational continuity during crises, supported by a national model assurance facility with institutional capacity to conduct adversarial testing, continuous verification and model behaviour auditing across the systems they depend on. They also need mechanisms to share information rapidly across jurisdictions. A cross-portfolio register of AI dependencies, vulnerabilities and emerging threats, coordinated centrally, would ensure that insights from one domain

strengthen resilience across the whole system rather than remaining siloed within individual agencies.

Strategic resilience depends on collaboration, and this requires a standing institution, not informal coordination between existing bodies. Government, industry and unions should establish a National AI Sovereignty Compact, chaired by the Department of the Prime Minister and Cabinet, with binding rather than advisory authority. Chairing this body at the centre of government, rather than through a single line department, is what allows the Compact to align sovereign compute, model assurance, workforce and R&D investment across departments that would otherwise pursue AI policy in isolation from one another. Binding authority matters for the same reason: a coordination body that can only advise leaves each department free to disregard its guidance when convenient, which is precisely the fragmented, siloed approach this submission argues has failed so far. The Compact should set sovereign capability standards for defence, essential services and critical democratic systems, and govern procurement conditions so that firms deploying AI in critical infrastructure must demonstrate compliance with sovereign capability and labour standards as a condition of Commonwealth eligibility. This gives the Compact real leverage over industry behaviour, since access to Commonwealth contracts, not just regulatory approval, becomes conditional on meeting these standards. Union participation in the Compact should extend beyond workforce representation to include co-design of deployment conditions and a genuine early warning function, since workers in telecommunications, media and critical infrastructure have frontline visibility into how systems behave under real conditions that regulators and technical audits alone cannot fully replicate. Shared standards, joint assurance frameworks, interoperable infrastructure and cross-agency expertise allow Australia to respond to AI enabled risks with speed and coherence.

The central insight is that strategic resilience is institutional. It is built through regulators with authority, an AI Safety Institute with technical depth, and a security community capable of governing AI systems that shape national stability. When these institutions are strong, coordinated and sovereign, Australia can anticipate risks, respond effectively and maintain control over the technologies that underpin its future. Without them, AI becomes a source of systemic vulnerability rather than national strength.

Conclusion

Artificial intelligence is reshaping Australia's economy, workplaces and strategic environment at a pace that demands deliberate national governance. Left to diffuse through markets alone, the evidence shows AI does not automatically deliver shared prosperity. Instead, it risks deepening existing inequalities between firms, sectors and workers, amplifying returns to capital while wages and job quality stagnate, and

in the workplaces this submission has documented, converting into measurable psychosocial harm rather than remaining an abstract economic cost. That harm does not stay contained within the firms generating it. As workers' compensation schemes tighten eligibility for psychological injury, the cost shifts to superannuation insurance and ultimately the public purse and ultimately the public purse, a genuine externality that any honest assessment of AI's productivity benefits must account for. But governed deliberately, AI offers the potential to lift productivity, rebuild industrial capability, strengthen national security and expand opportunity across suburbs and regions. Without sovereign infrastructure, enforceable workplace protections grounded in both industrial relations and work health and safety law, transparent public sector deployment and institutions capable of independent assurance, AI will entrench dependency, concentrate power and expose Australians to avoidable harm. The choices made now will determine whether AI becomes a driver of shared prosperity and national resilience, or a source of insecurity and inequality.

The evidence is clear: sovereignty and fairness are inseparable. Australia cannot be sovereign if it relies on foreign governed systems to run its defence, energy grid, hospitals or democratic information environment. Workers cannot be empowered if the algorithms that shape their employment are opaque, unaccountable or beyond contestation, and cannot be safe if those same algorithms determine their workload, their rostering and their exposure to disciplinary action without a genuine opportunity for human review. The case studies in this submission show these are not separate failures. An algorithm that sets an unsafe workload and an algorithm that disciplines a worker against an undisclosed standard are the same failure of governance, applied to different points in a worker's employment. National capability is built on industrial foundations, via a skilled workforce, sovereign compute and cloud, trusted model assurance systems, strong regulators, and a security community equipped to govern AI in high-stakes environments. Those foundations cannot be built without the active participation of the workforce that operates them, and that participation cannot be limited to industrial bargaining alone. It must extend to the work health and safety consultation this submission has shown is too often missing entirely, consultation that begins before a decision is made and continues as a system's impact becomes apparent, not consultation that begins only once a change has already been decided. That earlier, ongoing engagement is not a cost imposed on business for the benefit of workers alone. Systems designed with genuine input from the people who operate them are more likely to fit how work is actually done, and regulatory certainty introduced early gives industry a stable basis for long-term investment rather than the deferred risk of disruptive correction after harm has already accumulated. Good governance and genuine productivity are not in tension here; they are the same project.

Australia has the opportunity to lead, but leadership requires action. Recognising unsafe staffing, work intensification and algorithmic management as work health and safety obligations, not matters of managerial discretion; embedding enforceable

governance in workplaces through reform of the Fair Work system; building sovereign AI capability across compute, cloud and data governance; strengthening regulators and the AI Safety Institute so they have real technical authority rather than nominal oversight; closing the gap that leaves platform and public sector AI accountable in name but not in practice; and supporting small and family businesses through accessible skills and infrastructure are the essential steps toward a secure and inclusive technological future. Realising these opportunities — ensuring that both the economic value AI creates and the wellbeing of the workers who operate it remain protected in Australia, turning sovereign capability into export strength, and defending Australia’s information environment against synthetic manipulation — depends on treating this as a single, coherent national project rather than a series of disconnected policy responses.

Appendix A: Principles underpinning SDA policy positions

SDA policy is driven by providing value to our members whose work is regulated by an industrial system that, despite recent reforms, failed them for decades.

Australians need to be supported by an economic system that has working people at its centre. Our predecessors built an industrial system which provided the foundations for shared prosperity. It is now our responsibility to utilise the reformed industrial framework for the current and future generations. Decades of concerted attacks on our industrial relations system have seen inequality grow and economic and political power become further concentrated in the hands of a few.

The world of work has changed and will keep changing. There is an unprecedented intersection between work and care. Income and gender inequality have combined to increase disadvantage. Predictable, secure hours of work that provide a living wage are at the centre of decent work. But there has been growth in insecure work, digitalisation is now a matter of course, safety concerns have persisted, and automated, digital and generative technologies must be shaped to enhance, not undermine, decent work.

We believe that fundamental not incremental change is needed. In contributing to policy, we seek to drive a new system that acknowledges the change that has occurred and will be fit for purpose in the emerging world of work.

The SDA engages in topics that help drive this agenda and we are guided by ten principles that we believe will create value for our members.

Those principles are:

1. **Address Inequality & Enshrine Fairness**
Minimum expectations must be set and adhered to.
2. **Equity & Empowerment**
All workers must be supported to progress so that no-one is left behind.
3. **Mobility & Security**
A socially successful economy must provide opportunity for all, regardless of their background. Systems must be built in ways that support success and adaptation in a rapidly changing world of work.
4. **Delivering Prosperity & Growth For All**
A foundation for prosperity and economic growth must be achieved.
5. **Protection in Work & Beyond**
Workplaces and the community must be healthy and safe for all workers and their families during and beyond their working lives.
6. **Workers' Capital & Superannuation**
Workers' capital and superannuation must be an industrial right for all workers and treated as deferred earnings designed for dignity and justice in retirement.
7. **A Strong Independent Umpire**
A strong, independent, cost effective and accessible industrial umpire and regulator must be central to the future system of work in Australia.
8. **Protection & Support for Our Future**
Protecting and supporting our future requires a strong and vibrant retail industry and supply chain providing decent work and jobs with fair and just remuneration and contributing to the

economy including through skilled workers.

9. **Work & Community**

Work is a fundamental human activity that provides for personal, social and economic development. Work as it operates in community must build and protect a balance between life at work and life outside work so that workers can contribute to society through the wider community.

10. **Institutional Support for Collective Agents**

Institutional support must provide for collective agents (registered organisations) in all industries so that they are recognised, enshrined and explicitly supported as central to the effective functioning of the system.

Details of specific policy positions can be discussed by contacting the SDA National Office.

