pwc

AI in 2026: The AI-native enterprise

Designing your enterprise to embrace AI at scale

The AI-native enterprise

For business, Artificial Intelligence (AI) is entering its next phase.

The early years of experimentation, where pilot projects proliferated and algorithms dazzled, have given way to a new reality: AI now underpins strategy, operations, governance, and growth at the heart of business models.

What does it mean to be an AI-native enterprise?

It means that AI runs through your architecture, operations, and decision-making. It is embedded in processes, not bolted onto them. It is core infrastructure, a key lever of economic opportunity, and the fabric of the future. It remains decidedly human-centred, and relies on new ways of learning, thinking and working. In 2026, the central question for leaders is not *if* to use AI, but *how* to organise for it – how to capture its value at scale while sustaining trust, pace, and human meaning in the enterprise.

Based on an extensive review of global trends and forces, this report distils 26 of the most consequential ideas shaping this new landscape. Each represents a discrete but interlinked dimension of the modern AI-defined enterprise.

Together, they sketch what it will mean to lead an AI-native enterprise.

Contents

The next phase

The promise and potential of AI, the rise of trillion-dollar industries, the spectre of autonomous agents, robotic embodiment, protein folding, big deals, data centres, regulations, record-breaking AI employee incomes … the headlines and hype leave no doubt that AI is causing seismic shifts. Today, AI shows the potential to deliver the next tech-fuelled GDP boom.

But beyond anecdotes and wild extrapolations, leaders are asking 'what is it time for in 2026?'

Moving beyond experimentation

In 2026, Australia faces a pivotal moment in the evolution of artificial intelligence. The nation enjoys a rare convergence of supportive policy, enterprise investment, industry alignment and innovation capacity. After years of tentative pilots and demos, having come to grips with fundamentally new technology categories, AI has matured beyond the lab and onto the boardroom agenda.

This new phase is defined by AI becoming an integrated part of business systems. It is no longer a curious side project or demo – AI is becoming infrastructure: a core part of enterprise systems, processes and decision-making. Leaders are no longer asking *if* to adopt AI, but how quickly they can scale AI across their operations.

The CEO shift

As for the tone from the top, today's CEOs increasingly view AI as embedded infrastructure for strategy and operations. In PwC's 2025 CEO survey, nearly 70% of global CEOs said that generative AI will significantly change how their company creates, delivers and captures value in the next three years. And this isn't a far-off prediction – almost half of technology leaders report that AI is already fully integrated into their core business strategies.

In Australia, an overwhelming nine in ten CEOs see AI adoption as central to their business strategy over the next 3–5 years (PwC CEO Survey). These leaders understand AI capabilities are now as critical as finance, sales or any other pillar of the enterprise.

As this adoption maturity evolves, AI is being woven into customer platforms, supply chains, and decision workflows, augmenting human judgment with machine intelligence at scale. As a result, AI is now powering real productivity gains – from 15% to 40% improvements in focused areas – and enabling entirely new business models in those organisations that have committed to it fully.

While the value is real, distribution is spikey. Where some have seen speculative efforts result in nothing but sunk costs, companies that pair bold ambition with enterprise-wide execution are pulling ahead.

The experience so far has shown that piecemeal AI efforts yield piecemeal results. Scaling requires AI to be treated as an enterprise-wide system: aligned to strategy, integrated into architecture, governed for trust, and built for performance. A systemic approach towards 'AI-native' enterprise.

New rules

AI is increasingly becoming the fabric of the future, fundamentally changing how we define opportunity.

Instead of a chatbot pilot in customer service, an AI-native enterprise might integrate AI across the end-to-end customer support process, letting algorithms triage inquiries, draft responses, and only escalate to humans for complex cases.

Instead of a single predictive maintenance model on one production line, AI-native manufacturers might redesign entire maintenance systems that monitor equipment health continuously.

In shaping these opportunities, AI is also rewriting the rules of strategic advantage:

- Speed Matters More: Creative destruction is done proactively. While some companies continue to experiment, leading companies are running to integrate AI into everything they do.
- Scale Matters Less: Scale has long been a valuable source of advantage, serving as a moat and capability leverage. AI is redefining work and the size and shape of workforces.
- Innovation Matters Most: As whole new categories of business are invented and disruption drives both uncertainty and opportunity, innovators will reign supreme.
- Trust multiplies everything: Trusted, responsible, ethical AI derisks speed, manages scale, and provides leaders with the confidence to move.

AI in 2026

This report distils 26 of the most consequential ideas shaping the AI landscape in 2026. Each represents a discrete but interlinked dimension of the modern AI-defined enterprise.

Built on PwC global case studies and research, these ideas ultimately sketch what it will take to lead your enterprise into an AI-native future.

Entering 2026

Based on PwC 28th and 29th Global CEO Surveys:

- **90%** of Australian CEOs say AI is central to their business strategy
- Only **28%** believe their AI investment is sufficient (vs. 40% globally)
- **66%** say their culture enables AI adoption
- Only **18%** have built strong AI foundations

Experimentation

Curiosity, cautious pilots, and isolated innovation efforts. Organisations explore AI through proofs of concept and low-risk use cases, often led by innovation teams at the edge of the business. While excitement runs high, initiatives remain disconnected from core systems, lack clear ROI, and fail to scale. Avoidance of data quality and architectural foundations. Leadership engagement is limited, with AI considered a tool rather than a paradigm. Trust, governance, and change management often afterthoughts.

The result: scattered wins, limited transformation

AI-Native Enterprise System

From experimental technology tool to human-centred enterprise infrastructure. AI is embedded in the operations and decision-making fabric of the organisation - not bolted on, but built in. Leadership fluency, enterprise coordination and ethical guardrails are non-negotiable. Treated as a general-purpose capability, AI is scaled with confidence and deployed with trust by design. The focus moves from sporadic innovation to system performance: improving speed, accuracy, economics, and accountability.

The result: AI becomes how the business runs.

Side projects	**Strategy**	Enterprise value creation
AI for parity	**Competition**	AI for differentiation
Defer to technical specialists	**Leadership**	All leaders shaping and driving
Pilot ideas	**Value Realisation**	Tracking top-down value pools
Bolted on	**Process**	Exponentially reimagined with AI
Self-directed learning	**Capability**	Ubiquitous skill for human-AI teaming
Fringe enthusiasts	**Culture**	Change-led adoption, AI mindsets
Iconic tech hires	**Workforce**	Strategic AI-human workforce planning
Ad hoc	**Operating Model**	Cross-functional redesign
Tools and one-offs	**Technology Architecture**	Unified data, platform and explainability
Compliance and reactive controls	**Trust and Risk**	Trust by design, confidence to scale

Strategic advantage

Strategic advantage shows where AI reshapes strategy, markets, and value creation.

Value creation pinpoints the profit pools and growth bets that AI can drive and ties them to measurable outcomes. Industry edge hardens advantage by embedding domain knowhow into workflows and decisions. Ecosystems and innovation build speed, innovation and optionality through partners and platforms, without locking the business into a single path. Enterprise blueprints scale what works into reusable capabilities across the organisation. Value instrumentation tracks business value, impacts, cost, and risk in real time, so leaders can steer investment and performance with confidence.

Value creation 01

A true value creation strategy looks at how AI moves value pools – both within your P&L and across your market – by getting to the heart of your business model.

Value creation is not a layer of work on top of today's processes; it's a redesign of how a business earns: how products are conceived, priced and delivered, how customers are acquired and served, how risk is taken and managed, and where the boundaries of the business begin and end.

This reflects PwC's value-in-motion perspective that Australia's next boom won't come from how much we dig or ship, but from how we meet human needs. Value is shifting as AI changes cost curves and product boundaries, speeds cycles and personalises at scale. The practical job is to identify where margin will compress and where it can expand, then re-cut your business model so AI is built into the flow of work and into the product itself.

Value creation requires top-down perspective and an attention to time horizons. Near-term, focus on embedded AI features that change unit economics in existing journeys – higher conversion, lower churn, faster resolution, better risk decisions and embedded controls – so results show up in revenue and margin within quarters. In parallel, place longer-horizon bets where AI enables new scope and growth: data-enabled services, smart products with recurring revenue, and partnerships that extend reach without adding heavy fixed cost.

The throughline is disciplined attribution and reinvestment—prove impact early, and channel gains into the next wave of reinvention. Tracing AI possibilities to the bottom line.

What matters now

By 2026, delivering clear ROI from AI is an urgent priority. After years of experimentation, boards expect AI investments to translate into sustainable profitability, not just technology demos. Many businesses have not yet achieved tangible value from their AI initiatives, requiring a new set of business disciplines that change innovation efforts from pursuing 'proofs of concept' to 'proofs of value' and the capacity to scale and embed.

What changes in 2026

Inside the business, the approach to value undergoes a fundamental shift:

- Metrics expand: Instead of only tracking cost cuts and efficiency, firms now measure AI's impact on revenue growth, customer experience, and innovation.
- Beyond efficiency: Early AI efforts focused on automating tasks and boosting productivity. Now companies use AI to personalise offerings and enhance products, driving higher sales and customer lifetime value.
- New value streams: Forward-looking firms launch new services and business models enabled by AI (e.g. data-driven subscriptions or smart products), tapping fresh revenue streams and higher margins. Value creation thus spans from practical gains to radical reinvention.
- Value discovery with AI: AI deep research and analytic techniques can play a part in finding the big opportunities and continuing to discover opportunity signals in the market and your business.

Change innovation efforts from pursuing a proof of concept to a ‘proof of value’.

- From pockets of brilliance to targeting top-down value pools
- Dual-track value; understand the near- and long-term
- Clear line of sight to earnings and return on invested capital

Australian context

Australian businesses have the data depth, digital adoption and customer trust to turn AI into distinctive, revenue-led features—especially in financial services, retail, health, energy, mining and resources sectors.

The opportunity is to embed AI into the way the business runs, using local data and market insight to sharpen pricing, personalise service, and launch adjacent offerings. Those who move from experimentation to enterprise-wide deployment can set the pace in the region.

Leadership priorities

Tie AI to earnings: Make a short list of high-value journeys and products where AI can move revenue or margin meaningfully; assign accountable owners and baselines.

Instrument the flow: Build attribution into processes so value can be seen and debated—control groups, business KPIs, and disciplined stop/go decisions.

Concentrate capital: Shift funding from many small proofs to a portfolio of scaled deployments with clear payback and reinvestment plans.

Align with the CFO: Treat AI capacity and operating changes as part of the P&L; plan for ongoing run costs, resilience and performance tracking.

Signs of transition

Attributed value: Monthly reviews link AI features to commercial KPIs; underperforming deployments are fixed or stopped.

Fewer pilots, more production: AI shows up in customer journeys and product roadmaps, not as standalone experiments.

Better economics: Cost-to-serve declines while customer measures improve; new AI-enabled offerings contribute to recurring revenue.

Disciplined reinvestment: Savings and gains fund the next wave of product and process redesign.

The AI leader’s mindset

Treat AI as a value engine, not a cost line.

Build micro-P&Ls for AI-powered journeys and products, allocate capital based on unit economics per outcome, and reinvest gains into continual redesign. The bold step is to let AI reshape the business model—creating second-curve revenue and margin through embedded features and services—while keeping measurement and accountability at the core. The boldest leaders view AI as integral to strategy, continually asking: “How is this initiative moving our bottom-line or customer needle, and what’s the next evolution once it does?”

Industry edge

02

Generic AI won't create a competitive moat; industry and customer context will. Although AI is a general-purpose technology, its use is anything but general, thriving on well-defined tasks and use cases.

In 2026, AI advantage lives in the last mile of your industry. Competitive edge doesn't come from owning a bigger model – it comes from encoding the industry domain logic, risk posture and operating reality into AI enabled workflows that run the business.

A simple test for leaders: if a rival could recreate your AI without your data, workflow semantics and audit trail, you don't have an edge. When those elements are baked into the way work gets done, they become hard to copy.

Achieving an industry edge also means learning from the best in the field: leading companies actively seek out examples of ambitious AI plays in their sector and beyond, studying how those innovators achieved results and adapting those lessons to their own context.

Every solved edge case, every dataset with clear lineage, every control that accelerates sign-off becomes a reusable building block. Over time this compounds into an internal "industry OS": patterns, ontologies and evidence packs that let you scale new AI solutions with confidence.

Industry edge also reframes speed. In sectors where assurance matters, the organisations that can ship trusted AI fastest will outpace others, even if they use the same foundational models. Approval velocity becomes a competitive metric when guardrails, embedded controls, provenance and rollback are designed in.

Finally, this approach is a response to a key insight: big technology vendors often lack deep industry intimacy and can be like a hammer looking for a nail. Effective AI leaders aren't simply adopting one-size-fits-all tools; they are custom-fitting AI to their industry's context – including compliance requirements, customer expectations, and domain knowledge – to create a moat that others can't easily cross.

What matters now

Boards are probing leaders on how it's used to reinforce their competitive position in their industry. The conversation has shifted to ensuring AI initiatives are tightly aligned with industry and company-specific opportunities and pain points.

What matters to executives now is developing AI solutions that speak the language of their industry – literally and figuratively – and deliver measurable improvements in those core sector metrics, whether that be net interest margin for banks, reliability for energy providers, or patient outcomes in healthcare. The urgency is to turn AI into an engine of sector-specific performance, rather than an experiment.

What changes in 2026

AI adoption becomes deeply tailored to industry needs in several ways:

- Regulators step in: Industry regulators are increasingly issuing AI guidelines, making regulatory-grade AI a standard expectation.
- Vertical AI solutions proliferate: A wave of AI tools and platforms customised for specific sectors is gaining traction.
- Domain talent and teams: Organisations place higher value on talent who combine AI skills with sector expertise; more "AI + Industry" hybrid roles – or partners who work at the intersection of people, process, technology and industry.
- Global models, local adaptation: A common approach is using a powerful open AI model as a base but heavily fine-tuning it with proprietary Australian data and embedding local business rules or safety checks. A "glocal" strategy.

Generic AI won't create a moat – industry context will.

While global technology solutions are for everyone, value will accrue to those who own the industry layer – the data, the workflows and the trust mechanisms that others can't easily replicate.

Australian context

Australian companies operate in industries that often have strict standards and unique local conditions – from banking regulations to geographic challenges in mining. This context creates an imperative and an opportunity for industry-specific AI.

The push for "sovereign AI" – homegrown AI capabilities that ensure data stays onshore and models respect Australian values – is also gaining momentum. This means Australian leaders are increasingly balancing global technology with local innovation.

Leadership priorities

Invest in vertical capabilities: Allocate resources to sector specific implementation. This might mean training custom models on your industry data or partnering with sector-focused AI vendors.

Embed compliance early: Ensure AI systems meet industry regulations and ethical standards – for instance, a bank implementing AI should align with APRA's guidelines on model risk management.

Leverage proprietary data: Identify the unique data sets your company possesses and use them to enhance AI models.

Signs of transition

AI in core workflows: AI is embedded in mission-critical operations as standard practice, not pilots—for example, real-time credit decisions within risk guardrails or automated crop treatment adjustments.

Regulatory confidence: External reviews return minimal findings and acknowledge AI-enabled processes meet safety and fairness expectations

Differentiated offerings: AI-powered features become a clear market selling point, with customers recognising your sector-leading strength.

Industry data flywheel: AI systems generate unique operational data that continually sharpens performance, creating a compounding advantage.

The AI leader's mindset

View AI through the lens of your industry uniqueness.

The effective leader ensures that every AI initiative is anchored in a deep understanding of the business context. This mindset means being a specialist rather than a generalist with AI – constantly asking, "How does this technology solve a problem that matters in our field?" and "Does it meet the standards of our industry and customers?"

Ecosystems and innovation

03

No enterprise can go it alone – today's AI ecosystem is a federation of tech giants, hyperscalers, chip makers, platform providers, universities, novel start-ups, knowledgeable stalwarts. Partnering matters, and innovation *is* the conversation.

No single enterprise – no matter how large – can keep up with all the advances in AI by itself. *Ecosystems and Innovation* is about designing your AI strategy in 2026 to harness a network of external partners, open technologies, and collaborative models. It requires embracing a multi-partner, multi-model approach.

Companies are blending inputs from hyperscale cloud providers (for their infrastructure and AI services), open-source AI models (for flexibility and lower cost), academic institutions and startups (for cutting-edge ideas and niche expertise) as they build their AI capabilities. The point is to leverage the best of what the broader AI ecosystem offers, rather than reinventing every wheel internally.

However, doing this effectively requires clear rules of engagement – successful organisations set clear data boundaries and IP sharing rules so that working with partners or using open models doesn't compromise proprietary data or competitive advantage.

A crucial element of an ecosystem approach is optionality – building flexibility into your AI stack so you can switch components as technologies or economics change. By 2026, leading companies see their AI capability not as a single monolithic platform, but as an ecosystem of capabilities that can evolve. They maintain a mix of alliances – with global technology firms for scale and reliability, and with local players for customisation and niche innovation. In sum, "ecosystems and innovation" means recognising that AI progress is a team sport: to stay at the forefront, businesses are becoming more open, collaborative, and modular, while fiercely protecting the core assets that set them apart.

What matters now

In 2026, ecosystems shift from loose partnerships to a coherent system that leadership can rely on. Partners and open models slot in behind clear rules the enterprise sets—agreements on performance, provenance and cost, clear data and IP boundaries, and the flexibility to switch as economics or risk change. The mix typically combines global scale for reliability with local specialists for context, reviewed on a regular rhythm alongside value and risk.

Approval speed rises when trust requirements are built into how partners deliver, so the evidence arrives with the solution. The result is a modular capability that brings rapid innovation without lock-in, with proprietary data and operating know-how as the durable source of advantage.

What changes in 2026

The landscape demands interconnected innovation:

- Open models mainstream: Enterprise grade open models broaden choice and reduce dependency, vetted for security, provenance and bias.
- Alliance networks: Cross sector consortia tackle shared problems and set practical standards, speeding learning and delivery.
- Plug and play services: Modular stacks with an orchestration layer allow components to be swapped without disruption as economics shift.
- Ecosystem governance: Clear rules on data, IP and ethics - formal agreements and oversight - enable collaboration while protecting the core.

No enterprise can go it alone

Innovation

In the AI era, innovation isn't confined to R&D labs or early-stage experiments. It becomes a deliberate, ongoing capability woven into how products evolve, how services improve, and how value is created.

AI innovation means constantly testing boundaries: faster decisions, more personalised experiences, smarter automation. It's not about novelty for its own sake – AI innovation is about finding the next margin, the next growth lever, the next way to serve customers better.

In this context, innovation must be paced, purposeful, and aligned to the business model. Make innovation accountable—tie it to clear business goals, fund it like a portfolio, and embed it into the workflows of core teams, not just a distant innovation unit.

Best practice principles still apply:

- Start with enterprise value pools
- Build to prove value and scale (not concepts)
- Engage cross-functionally for a reality check
- Invest in a continuous innovation pipeline
- Measure returns

Australian context

Australia's AI market offers a deep local bench: a vibrant community of specialists alongside strong universities and research. Many firms pair niche capability from local startups with global platforms for reliability and scale, getting the best of both.

Openness is balanced with clear guardrails—onshore data residency, tight IP boundaries and alignment with Australian privacy standards—so collaboration doesn't dilute advantage. A close-knit, relationship-driven business community makes it easier to connect, prototype and ship; companies wired into both local and global networks tend to move faster than those going it alone.

Leadership priorities

Strategic partnerships: Focus on a small set of high-value partners aligned to your AI goals. Assign clear ownership and mutual value expectations.

Govern collaboration: Put guardrails in place – agree upfront on data use, IP rights, and how to exit.

Stay flexible: Design systems and contracts to avoid lock-in. Keep options open with vendors, platforms.

Blend inside and out: Encourage teams to co-create with external partners for speed and credibility.

Signs of transition

Active partner network: External experts work alongside in-house teams as standard, with co-development turning into shipped features and pilots scaling to production.

Multi-model ecosystem: Day-to-day operations mix models and services from different sources seamlessly, showing flexibility.

Rapid innovation cycle: New needs translate to production capabilities in weeks by plugging in partner APIs or proven open models.

Ecosystem ROI: Collaborative efforts deliver visible revenue or efficiency gains.

The AI leader's mindset

In 2026, the AI leader plays the role of orchestrator – bringing together the best of internal capabilities and external partnerships to accelerate results.

They know that value often lies beyond the company's four walls, and they build relationships that expand reach, speed, and learning. This means being open to external ideas, while staying sharp on data control, IP, and risk. The mindset is strategic, not insular: asking not just what the company can build, but what it can connect to.

Enterprise blueprints

04

Enterprise blueprints turn one-off successes into repeatable capabilities and guide the agile evolution of embedding AI where it matters most. Not in increments, but in steps and leaps.

Enterprise Blueprints define how an enterprise will evolve with AI—linking strategy, operations, and governance so that early wins turn into enterprise-wide capabilities, not isolated sparks.

Most organisations have pockets of AI brilliance – a chatbot here, a forecasting model there – but scaling those into core operations is an enterprise design challenge, not a tooling problem. A scaling blueprint tackles this head-on by aligning all the moving parts: strategy, operating model, data architecture, technology stack, and governance policies. An enterprise blueprint considers change and culture, operational controls, measurement, trust and governance.

It lays out how AI fits into how the organisation runs rather than treating each use case as an isolated project. This means formalising new roles and decision rights - for example, who approves an AI model for production, who monitors its outcomes – and building reusable components and platforms that teams across the business can draw on.

An enterprise approach doesn't require boiling the ocean. It is a question of coordination and coherence, ensuring businesses scale and change at a pace that matches impact. One bank might prioritise end-to-end automation in lending; a retailer might anchor on dynamic pricing and supply chain optimisation. The blueprint makes that explicit, guiding trade-offs and resource allocation over time, and just in time.

Agility matters just as much as structure. A good blueprint doesn't lock the organisation into rigid plans—it creates the scaffolding for adaptive execution. As AI technologies evolve and new opportunities emerge, the blueprint allows for fast reprioritisation without losing direction. It helps leaders move quickly, but in sync, so experimentation continues without fragmentation, and scaling happens without surprises.

What matters now

Many businesses now have dozens of AI pilots and tools but no unified way to scale them.

Boards and CEOs are pressing for an enterprise-wide plan that moves AI out of the innovation lab and into the fabric of operations. The urgent need is to escape "pilot purgatory" – where promising proofs of concept never translate into wide impact – by establishing a clear blueprint that links all AI initiatives to a common strategy and governance.

What changes in 2026

An enterprise blueprint for AI is not as a static plan, but a dynamic, evolving asset. It updates regularly to reflect what's working, and where new value pools are emerging. Leaders use it to steer, not to control – to make coordinated decisions at speed, and to ensure AI evolves in step with the business model. It becomes a shared reference point across the organisation: a practical, living map of how AI drives value, at a pace the enterprise can absorb.

Zero Basing

A strong blueprint doesn't just extend the status quo—it questions it. Zero-basing key processes, capabilities, or customer journeys helps uncover where AI can enable step-change improvements, not just incremental gains. It invites leaders to ask: if we were building this from scratch today, with AI in the mix from the start, what would we do differently? This mindset helps avoid simply digitising inefficiencies, and opens the door to bolder, more transformational plays—especially in areas like decisioning, service delivery, and product design.

Zero-basing helps uncover where AI creates a step-change, not just incremental gains.

It invites leaders to ask: if we were building this from scratch today, with AI in the mix from the start, what would we do differently?

Zero-basing means redesigning from first principles—assuming nothing and rebuilding as if starting fresh today.

Australian context

A well-defined AI blueprint helps ensure compliance with privacy, security and ethical standards while scaling innovations. It also suits Australia's market scale: by coordinating AI efforts across the business, firms avoid duplicated costs and fragmented tools. We're already seeing forward-looking Australian organisations establish central AI frameworks and business-led AI councils, aiming to turn isolated AI wins into a competitive advantage across all their operations.

Leadership priorities

Draft the blueprint: Align leaders on a shared AI plan that links strategy, ops, data, and tech.

Governance in place: Embed controls and set clear roles to guide standards and oversee risks.

Track and adapt: Use data to monitor progress and adjust the blueprint as needed.

Signs of transition

Unified roadmap: A clear, published AI strategy and multi-year roadmap

Common standards: Every new AI project adheres to enterprise standards

Executive oversight: A top team or board committee reviews AI progress and value delivery

From pilots to platform: Noticeably fewer ad-hoc pilots and more AI features embedded in the core for rapid and high-confidence innovation

The AI leader's mindset

Treat AI scaling as an organisation-wide redesign, not a series of tech projects.

The best AI enterprise leaders think like architects of a new business. Rather than green-lighting a scatter of disconnected experiments, they envision a cohesive enterprise where AI is woven into every workflow that makes sense. This mindset means being systematic and patient, connecting each initiative to a bigger picture, enforcing common standards, and building for long-term scalability over quick wins. They see AI as infrastructure: something to be planned, built, and tuned across the business.

It also means staying agile, treating the blueprint as a living tool that adapts as technologies shift and priorities evolve. And it requires boldness: zero-basing critical parts of the business to uncover where AI can unlock not just better, but fundamentally different. This combination of structure, adaptability and ambition is what turns AI from potential into performance.

Value instrumentation 05

Value Instrumentation means setting in powerful measurement that makes AI's impact observable in business terms, in real time, so that scaling decisions are evidence led.

Value instrumentation turns AI from a black box into a visible contributor to earnings and risk.

Measurement is embedded in the flow of work, not added after build. Each AI enabled process carries outcome KPIs, credible baselines and control groups for attribution, human feedback loops where judgment matters, and unit economics such as cost to serve and compute per outcome. Executive views centre on business telemetry rather than model scores.

The lens expands beyond safeguarding to make value visible and attributable. Signals tied to the P&L are tracked in real time, including revenue lift, margin change, cycle time and reductions in loss or rework, with causality evidenced through control groups and consistent measurement. Unit economics are explicit at the feature level, so scaling decisions weigh impact against spend. Leading indicators sit alongside lagging ones, such as adoption quality, user advocacy, decision accuracy in context and time to value, to show whether a capability is on track for commercial contribution.

New in 2026 is AI as a sensing layer for the enterprise. Intelligent monitors learn the normal patterns of decisions, costs and service responses, surface anomalies, detect drift and bias early, and reveal emerging opportunity pools before they show in quarterly results. Evidence is generated by default.

Controls record provenance, privacy and fairness as they run, creating audit ready trails that speed approvals and reduce rework.

What matters now

The landscape rewards organisations that treat value and risk telemetry as part of the product. Directional correctness beats perfect knowledge when leaders can see earnings impact, safety posture and cost in one view and adjust quickly.

Compute and data spend are tracked at the level of outcomes, not just projects, so ambition and capacity stay aligned. Approval speed rises when evidence ships with the solution, making scale predictable without slowing pace.

What changes in 2026

System-level measurement:

- From model metrics to business telemetry: Decisions shift from accuracy benchmarks to revenue, margin, loss and cycle time signals tied to specific AI features.
- Evidence by default: Controls create provenance, privacy and fairness records automatically, reducing approval friction and rework.
- AI as sensor: Always on agents watch for drift, anomalies and adoption patterns, elevating and correcting issues before they become incidents.
- Real time adoption signals: Engagement and effectiveness are monitored in the flow of work, separating tool access from behaviour change.
- Board ready dashboards: A small set of value and risk indicators becomes the leadership view, with clear stop, fix or scale thresholds.

Directional correctness beats perfect knowledge when leaders can see their earnings, cost profile and safety posture in one view and make quick adjustments.

Treat AI as both the actor and the instrument.

The same intelligence that powers workflows can also watch the system, learn what “good” looks like and surface weak signals that leaders can act on.

Australian context

Local data depth and sector expertise make outcome-level metrics practical, especially in financial services, retail, energy, resources and healthcare. Many businesses run in Australian cloud regions with clear residency needs, which simplifies provenance and audit trails while making compute a visible constraint to plan around. Organisations that turn product features, cost and assurance into a single view are converting digital head starts into earnings impact and approval speed.

Leadership priorities

Impact in view: A concise, shared dashboard that links priority AI features to commercial and risk outcomes, reviewed on a set cadence.

Economics with teeth: Standard unit economics for AI-enabled work, including compute, embedded in financial planning and benefits tracking.

Scale with confidence: Go-live criteria and continuous monitoring that produce assurance evidence as a by-product of delivery, keeping pace and control together.

Signs of transition

Economic traceability: Leadership sees revenue, margin and risk movement tied to specific AI efforts.

Unit cost visible: Compute and data costs are tracked and debated in performance reviews.

Evidence packs standard: Every release produces provenance, privacy and fairness records automatically.

Adoption in the flow: Usage quality and behaviour change are monitored beyond licence counts.

The AI leader’s mindset

Measure what matters, at the speed you scale.

Confidence in a moving market comes from seeing impact and risk in real time, accepting directional decisions over perfect knowledge, and running tight learning loops. AI becomes both the engine and the instrument, improving outcomes while revealing where to push, pause or pivot next.

Work reimagined

How AI developments land where it matters most: in the day-to-day reality of leadership, workflows, jobs, and skills.

AI raises a challenge of human-centred enterprise transformation. Leadership fluency explains why boards and executives now need direct ownership of AI decisions rather than delegating them to specialists. People-centric workflows capture the shift from “adding AI” to redesigning end-to-end processes so AI resources can do real work and human creativity can flourish. The evolution of work brings the changing nature of work into focus, with hybrid human-AI teams creating scope for greater human creativity and innovation. Culture and capability become an ultimate determinant of success, forging this reimagined work as consistent practice.

Leader fluency

06

AI is now a hands-on leadership domain, and leaders are expected to articulate a vision that blends human creativity with AI's potential. Through greater fluency, boards and executives build the knowledge to steer AI strategy, instead of deferring to specialists.

In the past, many CEOs and directors treated AI as a black box best left to the "tech people." By 2026 that stance is untenable. AI has moved to the centre of business models and risk profiles, so top leaders must personally engage with it.

Leadership fluency is about understanding AI's basics and implications well enough to lead from the front, making informed calls on where to invest, how fast to move, what risks to accept or mitigate, and how to get the right talent. It doesn't mean the CEO codes Python. It does mean they grasp how AI can drive customer experience, efficiency or innovation in their context, and the pitfalls to avoid.

Fluent leaders set the tone and narrative for AI in the organisation. They ask tough questions of proposals ("show me the value, not just the cool tech") and can sift vendor claims from reality. They ensure incentives are aligned to business outcomes – rewarding teams for AI results, not just experiments.

This is also a call to adaptive leadership. Leaders adapt their style to inspire, guide, and support diverse teams working alongside intelligent technologies. They must also be proficient in creating a vision that blends human creativity with AI's capabilities, fostering engagement and trust during the continuous change ahead.

Crucially, these leaders also tackle ethical and risk considerations head-on, establishing guardrails and accountability rather than leaving it to committees buried in the organisational chart.

Even regulators and investors now expect this fluency at the top, treating effective AI oversight as a core leadership responsibility. In short, leadership fluency turns AI from a side project into a core part of business strategy, with leaders who are credible, savvy champions of the change.

What matters now

By 2026, boards and CEOs face mounting pressure to "get smart" on AI.

Shareholders, regulators and customers all want to know how companies are using AI responsibly and effectively, and they expect the answers to come from the top. The priority now is moving beyond lip service and delegation.

What changes in 2026

From literacy to fluency, the conversation moves from how models work to how AI moves revenue, margin, loss and capital efficiency.

- From vendor led to strategy led: Partner choices and model use are routed by policy, economics and risk, not by who demos best.
- From episodic oversight to cadence: Regular, decision focused reviews replace sporadic updates, with thresholds to scale, fix or stop.
- From project funding to portfolio. Resources are allocated to a managed set of AI features with unit economics and payback expectations.
- From compliance gate to trust by design: Go live standards are proportionate to impact and built into delivery so approvals accelerate.
- From slideware to simulation. Leadership rehearses incidents and edge cases in synthetic environments to test readiness before scale.

At its extreme, the boardroom becomes a cockpit, and leaders run the company on live AI value and risk telemetry, practicing hard calls in simulation so pace and control can coexist.

Literacy is knowing how AI works.

Fluency is knowing where it earns, what it risks, and when to move. It is the foundation of applied knowledge.

Australian context

In concentrated, data rich sectors, fluency is less about how AI models work and more about reading live business signals: unit economics per outcome, adoption quality, and safety posture.

PwC's 28th Global CEO Survey already showed revenue and profit lifts tied to GenAI, and fluency is the difference between admiring those results and reproducing them in market – seeing where AI earns, where it risks, and when to move.

Leadership priorities

AI on the agenda: AI strategy and risk oversight become standing items in board and exec meetings.

Board education: Boards and executive teams invest in formal AI education so they can engage in informed decision-making.

Outcome-based incentives: Leadership reward systems start to change – executives and teams are measured on real business outcomes from AI.

Hands-on engagement: Many CEOs and business unit heads personally sponsor key AI projects, attend demonstrations, and even experiment with AI tools themselves – signalling to the organisation that this is a priority they own, not an innovation side-show.

Signs of transition

Boardroom fluency: Directors and top executives can hold a substantive discussion about AI initiatives.

Strategic AI discourse: The CEO and leadership team frequently reference AI in investor calls and town halls with concrete examples and informed viewpoints.

Governance in action: There's an active AI ethics or risk framework signed off by the board, and leadership visibly enforces its principles.

Confident decision-making: Major AI investment decisions are made (and occasionally halted) at the executive level based on clear business cases and risk assessments, showing that leaders are treating AI like any other core business initiative.

The AI leader's mindset

Treat understanding AI as core to leadership and make space for it in the agenda. It is a question of "learning AI" but "leading with AI".

Adopt a learner's mindset that asks sharp questions, engages with experts, and includes hands-on time with tools to build genuine fluency. The concern isn't knowledge gaps but making high-stakes decisions without clear signals. When things go wrong, lean into accountability, capture lessons quickly, and ensure course corrections are made. This tells people that AI matters and that leadership is fully invested.

People-centric workflows

07

The goal isn't to add AI but to fundamentally redesign a process from end to end while reimagining the role of people.

AI workflows are redesigned processes where AI is assumed to work alongside people. This shift is not about adding AI to existing process steps but rethinking how work is done end to end. Leading organisations are redesigning workflows such as customer onboarding and software development by asking how the work would operate if an AI worker were part of the team.

When routine synthesis and coordination are handled by AI, process hand-offs reduce, cycle times compress and quality improves while freeing humans to focus on judgement and accountability.

Good workflow design builds around AI's strengths and limitations while being explicit about the place of human oversight and control. Also, new process inputs such as prompt training data are treated as assets to be maintained and refined.

However, a vision grounded in AI's potential runs the risk of failing to account for the centrality of people.

As work is increasingly reimagined with AI, organisations are discovering that employee wellbeing isn't a 'nice to have' but essential for both adoption and effectiveness. When humans feel confident in their AI partnerships, they perform better and drive better outcomes. This requires building psychological safety by creating clear boundaries around AI decision making and establishing natural checkpoints where people add value. For example, AI might handle initial customer inquiries and draft responses, but a human manager reviews high stakes communications before they're sent.

The bottom line: real AI productivity gains don't come from a tool here or there – they come from redesigning workflows so that people and AI work in concert, each doing what they do best. This can slash turnaround times and elevate output quality dramatically while creating work that feels more human, more creative, and more connected to purpose.

When AI removes the mundane, what remains is the work humans are uniquely positioned to excel at.

What matters now

Entering 2026, many businesses have tried generative AI tools for drafting emails and writing code. The pressing challenge is workflow transformation: identifying which processes can be fundamentally redesigned around human-AI collaboration rather than simply adding AI tools to existing steps.

For companies to move beyond novelty and achieve productivity gains that genuinely improve how work feels and flows, the shift is urgent.

It demands thoughtful change management so employees develop confidence in AI partnerships, understand exactly where their expertise adds most value, and can focus on strategic thinking rather than constant AI oversight.

What changes in 2026

- End-to-end redesign: Organisations stop inserting AI into single tasks and instead rebuild whole workflows with AI agents and humans working in tandem.
- AI in the team: AI becomes an integrated coworker, embedded directly into roles and daily decisions. Standardised prompts, guardrails and workflows ensure consistent, high-quality use, while humans shift toward oversight, judgment, innovation and creativity.
- Trust by design: Quality, compliance and assurance become built-in features of AI workflows, not afterthoughts. Automated provenance, multi-layer checks and required human signoffs make speed and safety coexist.
- Run as a system: AI is managed like a business system, not a series of pilots. This makes AI predictable, repeatable and tied to P&L impact.

With AI, workflows are reimagined to elevate human strengths.

We see AI amplifying human potential – not replacing it.

Human control must be fundamentally present in all AI use. See Section 21, ‘Humans at the Helm’.

Australian context

An early call: generative workflows are quietly reversing Australia’s offshoring logic. As AI takes over the labour-intensive drafting, summarising and processing work that once justified sending tasks overseas, the advantage shifts back onshore toward local teams who bring regulatory nuance, cultural context and higher trust. Australian firms can run faster, safer, more compliant workflows at home, reshoring the judgment layer while automating the heavy lift. The result: onshore + AI becomes more competitive than low-cost labour ever was.

Importantly, to address Australia’s labour productivity challenge, we need to attend to workforce participation as part of AI work design. The goal is not to replace but to enrich our labour market.

Leadership priorities

Flagship workflow: Choose one high-value process and make it a focus for redesign, mapping reinvention from today’s flow to an AI-enabled version.

Cross-functional team: Bring together process owners, technologists and compliance experts to re-engineer the workflow holistically across people, process and technology.

Enable & equip: Provide the data, platforms and tools required for AI-embedded workflows and give teams the time and space to experiment and refine.

People support: Invest in training, clear role transitions and AI champions so teams can confidently adapt to new workflows and ways of working with AI.

Signs of transition

Faster cycles: Workflow times drop dramatically.

AI in the SOP: The company’s standard operating procedures and playbooks explicitly include AI steps.

Embedded in culture: Employees have embraced their AI helpers and coworkers, feeling psychologically safe to engage with the opportunity.

Better, not just faster: Operational metrics improve in tandem – not only are things getting done quicker, but error rates are down or quality scores up.

The AI leader’s mindset

Treat AI as a chance to redesign work from the ground up, not just a bolt-on to existing processes.

This mindset means looking at your organisation’s workflows with fresh eyes. The forward-thinking AI leader constantly asks, “If we have AI to assist, how would we build this process today from scratch?” Rather than settling for minor efficiency tweaks, they encourage teams to pursue step-change.

As you go, centre the strength and psychological safety of your people.

The evolution of work

08

AI doesn't just speed up how we work – it reshapes what work is. As old structures give way to hybrid teams, the real transformation is in how people and machines work together.

The nature of work is changing at the level of structure, not just tools. As AI expands its role in execution and decision support, traditional jobs are being decomposed into tasks, each with different requirements for judgment, speed or scale. This isn't just an additive step. Work is being reassembled into new flows that demand clarity, coordination and redesign.

Work used to be organised around fixed roles – one person, one job, one set of tasks. Now, that structure is giving way to more flexible models, where humans and AI share the load. Tasks are split, decisions are distributed, and outcomes are delivered through partnership. It's less about owning the whole process, and more about knowing where you add the most value.

Humans remain best placed for judgment, relationship-building, and ethical accountability. AI brings scale, speed, and synthesis. Rather than preserving job boundaries, organisations will align responsibilities to strengths, with value achieved through how well humans and machines complement each other in context.

At the centre of this shift is partnership. The human–AI relationship is a working alliance. Poorly defined pairings create bottlenecks and rework. Well-designed collaborations unlock autonomy, speed and reliability.

This is where hybrid teams emerge. Work is no longer solo or linear. Hybrid teaming pairs human creativity, discretion and empathy with AI's pattern recognition, memory and speed. These teams require shared systems and a working model fit for the future.

Decision-making must also evolve to support this paradigm. The approval economy – where every AI action waits on human review – can't scale. Risk-tiered controls, embedded evidence, and task-level oversight will need to replace universal checking. Therefore, trust becomes a function of how clearly roles are defined, not how often humans step in.

What matters now

AI is scaling faster than the work systems that make it safe and repeatable. Organisations are adding assistants and automating process fragments without redesigning the underlying work structure. The result: output increases but quality becomes inconsistent, risk shifts downstream, and managers become 'human middleware' constantly reviewing and correcting AI-produced work. AI multiplies not just throughput but also the volume of decisions, exceptions, and reviews that humans must handle.

How organisations restructure human-AI collaboration now will determine whether AI strengthens or fragments their capabilities.

What changes in 2026

Work shifts to human-AI partnerships that unlock capability neither humans nor AI could create alone:

- Partnership-first design: Organisations start with how humans and AI can best collaborate, identifying where judgment is required, where risk sits, and where AI can safely draft, triage, or execute – reflecting the shift from role preservation to strength-based task sharing.
- Human–machine teaming: AI will become a standard teammate with defined scope and observable performance, embedded in daily work.
- Hybrid leadership emerges: Manager capability in orchestrating human-AI partnerships, coaching collaborative judgment, setting partnership standards, and managing the transition from individual to shared accountability will become essential and widespread.
- The time dividend might start to materialise: As AI supports optimisation, leading organisations may commit to returning a portion of saved hours to workers, rather than letting benefits vanish into busyness.

AI is ushering in the reimagination of work – from fixed jobs to fluid roles.

From manual process to hybrid systems. From technology as tool to AI as teammate. Organisations that design for these dynamics will find not only new productivity, but a more meaningful kind of work on the other side.

Australian context

Australia's collaborative work culture and strong focus on workplace relationships provide natural advantages for human-AI partnership models. The country's emphasis on psychological safety and inclusive leadership aligns well with the trust-building required for effective human-AI collaboration.

Australia's productivity challenge is colliding with capacity constraints across frontline services and corporate functions. AI can lift throughput, but only if work is redesigned end-to-end, reducing rework, shrinking cycle times, and removing low-value activity. In this context, the evolution of work is how organisations convert AI gains into reliable service levels and better outcomes, not just faster activity.

Leadership priorities

Design partnership agreements: Establish clear role definitions, decision rights, and collaboration protocols that specify where human judgment is essential and where AI can operate autonomously.

Build hybrid team capability: Invest in training managers to lead mixed human-AI teams, including skills in partnership facilitation, performance optimisation, and trust building.

Create shared systems: Implement technology platforms and workflows that enable seamless collaboration between human and AI team members.

Measure partnership effectiveness: Develop metrics that capture productivity gains and relationship quality.

Signs of transition

Clear role definition: People can articulate what they own versus what AI owns, and partnership boundaries are well understood across the organisation.

Seamless collaboration: Human-AI teams operate with established rhythms, shared systems, and coordinated handoffs that feel natural.

Trust-based autonomy: AI operates within defined parameters while humans focus on exceptions and high-value decisions, without constant supervision or second-guessing.

The AI leader's mindset

Lead partnerships, not just people.

The focus shifts from managing human resources to orchestrating human-AI collaborations that unlock capabilities neither could achieve alone. The practical habit is designing relationships before deploying technology. Success comes from creating working relationships where both humans and AI can excel at what they do best.

Culture and capability

09

It's time to build the climate, conditions and capabilities that enable people to thrive alongside AI.

In an AI economy, the safest way to protect jobs is to make people more valuable than the tasks that can be automated. The PwC Global AI Jobs Barometer shows productivity rising faster in industries better placed to adopt AI, with workers developing AI skills earning meaningful wage premiums.

But capability alone isn't enough. AI maturity depends just as much on the environment in which people apply those skills – how teams learn, experiment, adapt and lead. That is cultural infrastructure, and without it, skills fail to translate into change.

The most valuable emerging skill is AI teaming – treating AI as a reliable partner rather than a tool. Beyond this, essential competencies include verification literacy, prompt crafting, model risk awareness, evidence capture habits, and exception management.

These technical abilities are valuable when the environment invites people to use them. Culture sets the conditions.

Culture determines whether new skills are exercised or abandoned. It shapes whether learning feels rewarding or risky, and whether experimentation is welcomed or quietly avoided. High-performing AI cultures focus on confidence: people feel safe using tools, asking hard questions, and rethinking how value is created. That shift begins with leaders who coach judgment, model transparent use, and make time for team learning.

To be conducive to achieving scale enterprise value from AI, an organisation's culture must support bold experimentation while maintaining ethical boundaries, building collective confidence to move from pilots to production, and developing shared values so teams can navigate strategic and ethical questions with agility.

Culture and capability become one agenda, because capability without cultural readiness creates surface-level adoption, and cultural change without skill development lacks practical foundations. When people believe their growth is valued, and their capability makes a difference, behaviour change follows.

What matters now

The immediate priority is to define and scale new baseline competencies for AI teaming: AI literacy and safety standards (privacy, data handling, compliance), prompt and context-building, verification and critical judgment (testing, source-checking, error detection), and workflow orchestration (knowing when to delegate to AI, when to escalate, and how to capture evidence).

With the rise of automation, skills like critical thinking and problem solving are also enjoying a renaissance. Understanding the role of uniquely human capabilities across your business model and value chain will be a necessary investigation.

What changes in 2026

Assuming a move to scale, culture and capability won't sit to the side:

- Cultural readiness enables confident execution: Organisations will develop shared values and decision-making frameworks that support moving from pilots to production, with cultural norms that balance experimentation with trust and ethics.
- Real-time capability becomes a norm: With skills churning faster than formal training can deliver, AI-powered coaching, embedded practice, and routine check-ins become essential.
- AI teaming becomes a job requirement: Organisations will formalise AI teaming competencies as standard performance criteria.
- Change management moves from messaging to design: Instead of messaging campaigns aimed at influencing adoption, change leaders will make AI-augmented workflows reliable by design. Effective change leaders will build the verification steps, escalation protocols, and human oversight required to make adoption more seamless.

If AI is your vehicle, skills are your licence to drive.

The licence comes from mastering new competencies: teaming with AI, understanding where it shouldn't go, knowing when to override it, and inspecting how decisions are made.

When these skills spread, work is accelerated, results are reliable, and people remain accountable for outcomes.

Australian context

Australia's skills agenda is already moving toward tighter alignment between training, workforce demand, and productivity outcomes, including reform settings across the VET system. This creates a practical opening for employers: treat AI as a portable, recognisable credential backed by real task competence, not internal "course completion."

The industrial relations environment emphasises consultation, transparency, and shared value creation—principles that align well with developing AI capability through cultural change rather than technological imposition. Unions' calls for enforceable agreements around AI deployment, including job security and retraining commitments, reinforce the need for organisations to demonstrate that AI enhances rather than replaces human contribution.

Leadership priorities

Build the learning culture alongside the skillset: Create psychological safety for AI experimentation while establishing practical competencies in verification, context-building, and workflow orchestration.

Normalise AI teaming as everyday work: Define what good looks like for hybrid teaming and embed it into role expectations, performance systems, team rituals.

Make capability visible and renewable: Shift from course-led training to coached practice in the flow of work, with time and support in place for improvement.

Anchor identity to a progressive AI strategy: Connect AI capability to how people see their value, potential, and future in the organisation.

Signs of transition

Confidence: People take pride in their AI collaboration skills and can speak to how they use them.

Embedded norms: AI use is governed not just by rules, but by shared expectations and team routines.

Credible pathways: Reskilling is visible, supported, and aligned to business direction.

Institutional maturity: Teams learn from each other, improve together, and adapt at the rate of external change.

The AI leader's mindset

Build a culture that renews capability, and capabilities that strengthen culture.

Treat culture and capability as mutually reinforcing infrastructure, investing in the specific competencies people need and the cultural conditions that make those skills valuable, sustainable, and continuously evolving. Success comes when new ways of working become how people prefer to operate.

Building intelligence systems

The technical architecture that makes AI real.

Foundation models supply the core capability of most serious AI use cases today, while orchestration and routing determine which model is used for which task at what cost. Agentic AI systems turn that capability into a digital workforce, running multi-step operations end to end and closing the last mile of automation across tools and systems. Context strategies provide enterprise memory, so actions and answers are grounded in current knowledge. Synthetic environments de-risk AI deployment by rehearsing scenarios, stress-testing edge cases, and generating lawful synthetic data where real data is constrained. AI Ops keeps the system stable over time. A compute strategy defines the practical limits and trade-offs for economically prudent scaling.

Foundation models 10

Foundation models matter because they have become the common capability underneath almost every serious AI use case today. They are also no longer the moat for AI strategy.

Foundation models are the general-purpose “brains” behind modern generative AI. They are large systems pre-trained on vast datasets and then adapted for many tasks, from writing and coding to analysis and customer support. For boards and executives, the technical detail matters less than the strategic implication. These models are becoming a core input to how work gets done.

For years, the industry treated foundation models as a winner-takes-most race. Enterprises followed the same pattern by trying to pick the single “best” model and standardise on it. That approach is fading. Performance gaps are narrowing, open-source options are improving, and model availability is expanding. The result is a different competition. Capability still matters, but differentiation increasingly comes from how well organisations turn capable models into reliable systems.

This is why “owning the brain” is rarely the right ambition. The strategic asset is the layer that sits above models and makes them usable at scale.

The failure mode is familiar. Technology teams chase benchmark scores, lock into a single provider, and discover too late that unit costs are unstable, data boundaries are unclear, and governance cannot keep up with production use. Model performance remains strong, but outcomes become inconsistent because the system around the model was never designed.

A simple test helps point to the necessary strategic shift from high dependency: if your preferred model became unavailable or uneconomic next month, could you switch models without rewriting workflows, changing controls, or lowering quality in critical decisions?

What matters now

The most important foundation model concern is operational independence. Many organisations have momentum in AI use but are building that momentum on fragile foundations such as one model, one vendor pathway, and informal controls. Cost can drift or reset, risk can migrate, and teams can end up with inconsistent outputs across business units.

The immediate priority is to build a practical control plane. Leaders need a clear way to select models by task, apply rules for data sensitivity, measure cost per outcome, and maintain a consistent standard for provenance and review.

What changes in 2026

In the year ahead, foundation models start to look like compute – necessary, widely available, and increasingly *interchangeable*. The model race matters less than the system discipline that makes models dependable, affordable, and governable:

- Model routing becomes normal: Organisations use multiple models. Simple work flows to efficient small models, complex work escalates to larger models, and critical decisions trigger tighter controls and review.
- Control planes: A policy layer governs which model can be used for which tasks, with clear rules for risk, sensitivity, and oversight.
- Data stays put more often: Private inference, edge deployment, and Australian regions become common. The operating goal is to bring models to data rather than pushing sensitive data outward.
- Small models win more workloads: Distilled and specialist models take a larger share of tasks when they reach parity for defined use cases. Efficiency becomes a competitive advantage.

Success does not come from mere access to the most powerful foundation models.

It comes from building systems that intelligently route tasks, ground outputs in trusted data, and maintain the right human control. Context matters more than raw intelligence. Data sovereignty matters more than model size. Cost per outcome matters more than model benchmark bragging rights!

Australian context

Australia's data and risk environment makes orchestration more than an efficiency play. Many sectors operate under strong expectations for privacy, resilience, and defensible decision-making. Model choices need to align with those expectations, including where inference (happens and what data crosses boundaries.

This pushes Australian organisations toward patterns that reduce exposure while keeping performance high. Private inference, deployment in Australian regions, and edge use cases matter. The same logic applies to critical services and regulated environments. The goal is not to slow down adoption. The goal is to scale AI in a way that remains controllable and auditable.

Leadership priorities

Mandate portability: Require model and workflow designs that engage with 'best model for the job' and avoid lock-in and keep switching costs low.

Institutionalise assurance: Make testing, red-teaming, and incident response a standing operating rhythm for model use, not a one-off gate.

Signs of transition

- Model selection happens through policy and routing, not by preference or habit.
- Most production workloads run on small or specialist models, with larger models reserved for high-complexity work that justifies the cost.
- High-sensitivity use cases run with clear data boundaries and audit-ready evidence.
- Leaders can see cost, quality, and risk performance by AI task. Cost per outcome becomes the scoreboard, and the total cost for a completed outcome includes failure handling, review time, and rework.

The AI leader's mindset

Treat foundation models as a rapidly improving commodity input and focus leadership attention on the system around them. Reliability, economics, and governance come from the control plane, the context layer, and the discipline of orchestration.

The best question to keep asking is simple. Are we building dependence on a model, or are we building the capability to use any capable model safely, efficiently, and at scale?

Agentic AI systems 11

Agentic AI is built to perceive, plan and act across multi-step tasks. They use tools, read policies, collaborate with humans, and move work forward according to complex goals. They are the difference between AI that *suggests* to AI that *does*.

The most hyped AI change through 2025 was the move from single-purpose agents to orchestrated multi-agents. Early agents were narrow and fragmented, each with their own interface and management overhead. While the journey continues, by late 2025, stronger reasoning and tool use meant one primary agent could break down a complex request, assemble the right sequence of specialist actions, and execute across systems without constant human prompting.

This matters because it unlocks the last mile of enterprise automation. Agents can run continuous micro-operations that humans struggle to sustain at scale, such as progressing a claim with missing information, replanning supply when constraints shift, compiling regulatory returns, or triaging service requests and routing exceptions. Autonomy compounds value when the task is multi-step, the environment is dynamic, and the process has enough rules and data to keep decisions bounded.

Agents also change the shape of the workforce. They behave like digital workers who need clear scopes, permissions, supervision and performance expectations. Workforce planning expands to include human and digital capacity, role design becomes about what humans must own versus what can be delegated, and managers need practical routines for oversight, escalation and safe hand-offs.

The technology story underneath this shift is equally specific. Composable ecosystems are emerging because no single provider covers every agent capability. Interoperability protocols are reducing brittle integrations, and identity-first design is becoming mandatory because non-human accounts and actions can scale faster than traditional governance can track. Organisations that treat agents as accountable workers will be able to use autonomy without losing control.

What matters now

What matters now is that agents turn AI from advice into action. Organisations that treat agents as “just another tool” will see speed rise alongside hidden exposure, because non-human actions scale faster than traditional governance and management routines can absorb.

What changes in 2026

While it may take some time this year to come to fruition, “super agents” shift the challenge from managing tools to managing digital workers with interoperability, identity, and runtime controls built in:

- Super agents: A single front-door agent orchestrates multi-step work, calling specialised capabilities and tools behind the scenes so users stop managing a patchwork of separate agents.
- Interoperability protocols: Standard ways for agents to connect to tools and to each other accelerate multi-agent workflows and reduce one-off integrations (hard to secure and maintain).
- Identity-first delegation: Agents are treated as workers with accounts, least-privilege permission and auditable trails, because non-human access can quickly outnumber human access.
- Runtime governance: Policies move from documents into the running system, with enforced scopes, confidence thresholds, logging, replay, and stop mechanisms so exceptions are handled cleanly.
- Last-mile operations: Agents take on continuous coordination work across service, operations and supply chains, clearing backlogs and handling exceptions, while humans focus on judgment, customer trust and high-impact decisions.

In 2026, onboarding an agent will feel like onboarding a team member –

Verifying their credentials, defining their responsibilities, understanding where they add most value, monitoring their performance, and making tough decisions when the fit isn't right.

Australian context

Australia is already operating with a large workforce and persistent gaps. Even with easing pressures, around a third of occupations remain in shortage, with acute gaps across trades and many professional roles, particularly in health, education and construction. Bounded agentic capability could serve as a capacity strategy. Digital workers can absorb the coordination load that grows in tight labour markets, keep processes moving across time zones, and reduce administrative drag so scarce human capability is used where it matters most.

Leadership priorities

The Chief People Officer becomes central because workforce planning now includes non-human capacity. Role design needs clear delegation boundaries, and managers need routines to supervise work they did not personally perform.

When you are ready for critical mass:

1. Stand up "digital worker" management: Create agent identities, role-based permissions, and clear ownership so every agent has defined scope, supervision, and accountability.
2. Equip managers to run hybrid teams: Update workforce planning, role profiles, and manager routines so people know what gets delegated.

Signs of transition

Scoped: Agents operate with stable scopes and service levels, and exceptions follow a defined escalation path that works under pressure.

Scaled: The scaled use of agents sees operational performance improve without a matching rise in incidents, rework, or compliance findings.

Secure: Security and risk teams can answer "who did what, when, and why" quickly, using complete logs and clear agent identities.

The AI leader's mindset

Treat agents as members of your team.

That mindset shifts the focus from "deploying a capability" to managing a blended workforce where digital workers take on coordination and execution, and people own judgment, trust, and accountability. Lead the way you would if you were hiring at scale. Give agents defined roles, scoped authority, and measurable expectations.

Context 12

Context turns models into enterprise judgement by connecting data, documents and decisions into enterprise memory. It makes answers current and allows agents to act with confidence.

Foundation models know a great deal about the public world, yet they know nothing about your strategy, your risk appetite, your customer promises, or what was agreed in yesterday's steering committee. Think of a model as a brilliant new hire who has read and memorised the internet, but has not been given access to your files, your policies, or your decision history. Context is the access badge and the briefing pack.

This "fact gap" has become the practical constraint on AI usefulness inside enterprises. When context is weak, outputs become generic, inconsistent, or wrong in subtle ways that create rework and erode trust. When context is strong, the same model becomes specific, defensible, and aligned to how the organisation actually operates.

From a technical standpoint, context refers to the information an AI system can lawfully access and cite in the moment. It includes structured data from core systems, unstructured knowledge in documents and emails, and the operational rules that shape how work gets done. It also includes freshness, permissions, and traceability, because those are the properties that make knowledge safe to use.

A clear set of building blocks is emerging for how organisations create the Context layer. Retrieval-Augmented Generation (RAG) grounds AI outputs in trusted internal sources, rather than relying on model memory alone. Knowledge graphs add another layer by mapping relationships between concepts, policies, decisions, people and projects, which improves relevance and reduces blind spots across silos. Together, these patterns turn "search" into a system that can answer with context and evidence.

The bigger shift is organisational. Many enterprises still treat knowledge as storage. Content lives in folders, duplicates multiply, and outdated guidance lingers. A context layer needs a rigorous lifecycle, so knowledge stays accurate, permissioned, and maintained.

What matters now

The "fact gap" between what models know and what your enterprise needs is a key constraint of AI value.

Foundation models excel at general knowledge but lack access to your operational reality, such as current policies, recent decisions, and institutional context that makes outputs relevant.

The priority is treating knowledge as maintained infrastructure rather than static storage, with clear ownership and lifecycle management. Organisations that establish this discipline in 2026 will unlock AI that scales with enterprise confidence rather than generic capability.

What changes in 2026

In 2026, advantage shifts from model cleverness to context discipline that makes answers consistent, current and auditable:

- Knowledge is treated as a maintained system that captures decisions, policies and outcomes with clear ownership and refresh cycles.
- Retrieval becomes the standard pattern: outputs are sourced, permission-aware, and aligned to approved information.
- Trusted tiers of knowledge: High-impact workflows rely on curated "best available" sources, while lower-grade content is clearly marked and used with appropriate caution.
- Provenance is operational: The lineage, freshness and confidence for data 'travels' with outputs so review is fast and accountability is clear.

Foundation models have memorised the internet but can't access your filing cabinet. They can recite Shakespeare but not your strategy.

And if you're relying on the same generic model advice as your competitor, you enjoy no competitive advantage.

Context turns generic intelligence into institutional wisdom.

It converts brilliant generalists into trusted enterprise advisors.

Australian context

Australia has world-class public information that can strengthen enterprise context when it is used well. Economic statistics, geospatial and environmental data, legislative and regulatory material, and sector reporting can provide reliable reference points that improve grounding and reduce ambiguity. For many organisations, the strongest context will be a blend of internal knowledge and selected public datasets that are stable, well governed and easy to refresh.

Leadership priorities

When it comes to Context, the priorities are technical:

Build a context backbone with owners and standards: Define what counts as trusted knowledge for material workflows, assign product ownership, and set measurable standards for freshness and permissions.

Unify retrieval across systems and content: Connect structured systems, unstructured repositories and policy sources into a single permission-aware retrieval layer that supports RAG.

Operationalise trust through provenance and tiers: Establish curated sources for high-impact work and mark lower-confidence content clearly.

Leaders inspect the health of enterprise memory through strackable signals such as freshness of critical policies, duplication and conflict rates, provenance coverage in material workflows, and the stability of curated sources used for customer and regulatory outputs. This posture keeps AI grounded in the organisation's reality and ensures autonomy, scale and trust can grow together.

Signs of transition

Context Layer in Place: Agents and new AI workflows can be added without rebuilding knowledge foundations because the context layer is reusable.

Curation: High-impact workflows rely on curated sources with stable permissions, and ad hoc content stops driving important decisions.

Currency: Users can see where an answer came from and how current it is, and owners exist who can fix gaps quickly.

Consistency: AI-assisted answers remain consistent across teams because trusted sources are clearly defined and maintained.

The AI leader's mindset

Leaders treat context as business infrastructure that deserves the same seriousness as financial controls and operational systems. Attention shifts to whether the organisation's knowledge is current, permissioned, and fit for decision-making at speed.

Synthetic environments 13

Synthetic environments make AI safe to scale by rehearsing reality in digital form. Simulation labs, digital twins and synthetic data help to expose edge cases early, creating faster learning, stronger assurance, and fewer surprises.

AI is moving into higher-stakes work where mistakes carry real consequences. When systems shape customer outcomes, move money, influence clinical pathways, or control physical assets, learning in production becomes an expensive way to learn.

Synthetic environments offer a different path. A synthetic environment is a controlled digital setting that mirrors operational reality closely enough to test AI workflows, stress edge cases, and validate controls before real customers, real staff, or real assets are exposed. Digital twins, simulation labs, and synthetic data pipelines bring the same discipline that engineering and safety-critical industries have used for decades into the AI domain.

This becomes more important as AI systems become more autonomous. Agentic workflows expand the surface area for unexpected behaviour, because the system is no longer responding in a single step. It is planning, taking actions, and interacting with multiple systems over time. Synthetic environments allow organisations to rehearse the full sequence, including failures, hand-offs, and recovery paths.

The strategic payoff is speed with fewer surprises.

Simulation reduces the cost of iteration, because teams can run thousands of scenarios quickly, explore extreme conditions safely, and tune policies and thresholds without waiting for rare events to occur. Synthetic data helps when real data is scarce, sensitive, or constrained by sovereignty and privacy requirements. The outcome is a practical way to harden systems, build confidence, and generate assurance evidence as part of delivery.

What matters now

Synthetic environments unlock a new level of AI system confidence and speed.

As workflows become automated and agents act across tools, traditional testing can't validate the full complexity of real-world scenarios, especially when data is restricted or critical events are rare. Synthetic environments enable teams to rehearse complete scenarios, prove control effectiveness, and generate comprehensive evidence before production deployment. This transforms the trade-off between innovation speed and operational safety.

What changes in 2026

Becoming an AI-native enterprise requires a shift from checking models to rehearsing full systems under realistic conditions:

- Digital rehearsal: Teams test end-to-end workflows, dependencies and failure modes.
- Synthetic data: Safe, representative datasets support testing when real data is constrained.
- Continuous evidence: Assurance artefacts are generated automatically as part of delivery.
- Hardened controls: Escalation and recovery behaviours are proven under stress before production.

Synthetic environments turn ‘hope it works’ into ‘know it works’.

Digital Twins are high fidelity virtual replicas of physical assets, processes, or systems that you can use to simulate and predict outcomes.

Australian context

Mining, energy, transport, and healthcare operate with safety-critical processes, harsh environments, and complex regulatory expectations. Synthetic environments suit these conditions because they allow hazardous scenarios to be explored without exposing people, patients, or assets.

Leadership priorities

If simulation can be meaningfully applied:

Build the lab: Stand up simulation capability for a small set of priority workflows.

Create synthetic data: Establish a governed pipeline for realistic test data where needed.

Ship with evidence: Make scenario coverage and control performance part of release readiness.

Signs of transition

AI arrives with information about test coverage across scenarios, including rare edge cases and failure modes.

Regulators and risk leaders see clear evidence of control effectiveness without needing special investigations or one-off reporting.

Teams iterate faster because they learn in simulations, and production incidents become less frequent and less severe.

Sensitive data exposure reduces because synthetic data supports testing and improvement without repeated access to real customer records.

The AI leader’s mindset

Synthetic environments as a safety and speed capability. Confidence comes from evidence that systems behave well under stress, across dependencies, and in rare conditions.

AI Leaders focus on whether critical workflows have realistic rehearsals, whether controls perform reliably in simulation, and whether assurance artefacts are created as part of routine delivery. This approach allows the organisation to move faster while protecting customers, staff, and trust.

AI Ops

14

AI Ops industrialises how models, prompts and agents are shipped, monitored and improved. It keeps quality, cost and control stable as AI becomes part of daily operations.

When AI is embedded in workflows, it stops behaving like a project deliverable and starts behaving like production infrastructure. That changes the standard leaders should expect. The question becomes whether these AI systems are run with the same reliability, transparency and recovery capability as other mission-critical services.

Most organisations began their AI journey with artisanal deployments. A small team tuned prompts, shipped a model, watched a dashboard, and fixed issues when someone complained. That approach breaks with scale. Small changes in data, user behaviour, policy settings, or model versions can shift outcomes. Costs can drift quietly and quality can degrade without an obvious “down” event.

AI Ops is the operating layer that makes AI dependable. It includes model and prompt registries, CI/CD (continuous integration / deployment) so updates are tested and released predictably, and observability that measures drift, quality and failure patterns over time. It also includes cost guardrails, because AI spend behaves like a new kind of compute bill that can scale rapidly.

This is more than tooling. AI incidents need to be treated like operational incidents. That means clear playbooks, defined roles, and rapid rollback when something misbehaves. It also means joining technical telemetry with business telemetry. Leaders need to know whether a model is healthy in the ways that matter to them, such as customer outcomes, compliance, and rework.

AI Ops also clarifies an operating model. Cross-functional platform teams provide the shared capabilities that product squads need, so teams can focus on outcomes while the platform carries the complexity of safe deployment and monitoring. The result is not central control for its own sake, but speed that remains stable.

AI Ops transforms algorithmic potential into a business reality by ensuring your AI systems perform when the stakes are highest.

What matters now

As agents begin to take actions across systems, small failures can propagate into customer impact and compliance exposure. Leaders need confidence that changes are controlled, performance is observable, and recovery is fast, because AI will now fail in ways that look like business variance, not system downtime. AI Ops is the branch of management concerned with this.

What changes in 2026

The paradigm shifts from managing models as artefacts to running AI as an always-on service with disciplined release and recovery:

- Registries and release discipline: Models, prompts and policies are versioned, tested and promoted through environments with clear approval thresholds.
- Observability becomes business-grade: Monitoring covers drift, quality and failure modes, linked to customer and operational outcomes.
- Incident response matures: AI failures are handled with playbooks, on-call ownership, rollback paths and post-incident learning.

The most sophisticated AI model is just expensive code if it can't operate consistently at scale.

This report uses "AI Ops" as a conceptual term for operational challenges across AI implementations.

While often used interchangeably with "ML Ops" (machine learning operations across the lifecycle of model training, validation, deployment, and monitoring), AI Ops seeks to account for broader system factors. It is still an emerging and contested field with evolving definitions.

Australian context

Australia has a practical scaling constraint. Many organisations run nationally distributed operations with uneven access to specialist talent, and they rely on shared platforms to deliver consistent service across states and time zones.

A strong AI Ops platform reduces the dependency on scarce specialists by automating release, monitoring and rollback, and it creates a single operational picture that allows governance and risk teams to oversee AI use without slowing frontline delivery.

Leadership priorities

To drive a strong transition to robust AI Ops, leaders might prioritise:

Enterprise service standards for AI: Define what "good" means for quality, drift, latency, and cost per outcome, and require every AI-enabled workflow to meet those thresholds before and after release.

Clear business ownership for every model and agent: Name a single accountable owner for performance and risk in production, with a RACI for approvals, monitoring, incident response, and rollback.

Make AI incident readiness non-negotiable: Require playbooks, escalation paths, and rollback options for material use cases, and run regular rehearsals so teams can respond quickly when behaviour shifts.

Signs of transition

The existence of a function is itself a proof of maturing to the system, but strong AI Operations can be seen in the form of:

Controlled changes: Updates ship through a consistent pipeline, and teams can roll back quickly.

Visible performance: Leaders can see quality and drift trends in plain terms, and action is taken early.

Known ownership: All know who is accountable for a model in production and who is on call.

Learning loop: Repeated failure modes drop because fixes are captured and reused.

The AI leader's mindset

AI Ops is a cultural shift from "shipping features" to "earning reliability". This lens values quiet excellence over flashy demos. It rewards teams that surface issues early, measure outcomes honestly, and improve systems steadily. When that culture takes hold, AI becomes something you can rely on.

Compute strategy 15

Compute scarcity was the rude awakening of 2023. AI hit hardware limits that turned unlimited ambition into finite resource realities. While capacity has increased considerably since then, rising computational demands continue to challenge processing supply.

Compute (the computational resources like processors and accelerators needed to run AI models) has become a constraint that boards can feel. As AI moves from experimentation into daily operations, the limiting factor is often no longer ideas or model access. The limiting factor is whether the organisation can secure, place and afford the compute needed to run AI reliably.

In 2025, many teams treated compute like an elastic utility. They spun up capacity when needed and absorbed costs as “cloud spend”. That mental model is breaking. With AI, demand spikes, GPU availability is uneven, and costs scale with usage in ways that surprise business cases. This forces a shift from consumption to planning.

Compute strategy is the discipline of deciding where workloads run, how capacity is secured, and how efficiency is designed in. It includes choices about edge, on-prem and Australian cloud regions. It includes reservations and capacity commitments. It also includes workload placement rules so the right tasks run on the right class of infrastructure.

Unit economics sit at the centre. Executives need to see AI cost in the same terms as any other operating decision: cost per task completed, tokens per case resolved, GPU-hours per model refresh, and the real rework burden when quality slips. Once those measures exist, teams can right-size workloads by using smaller models when they suffice and escalating only where value justifies cost.

Compute also carries an emerging external dimension. As we will explore in topic 25, Energy use and water intensity are becoming practical considerations as disclosure expectations rise and carbon pricing scenarios harden. Compute strategy therefore becomes a combined business, financial and operational question, not an infrastructure footnote.

What matters now

Many organisations are now seeing “successful” AI features become expensive at volume, with capacity constraints appearing at exactly the moment the business wants to expand. Compute decisions are therefore starting to shape product choices, service levels, and rollout speed in practice.

The organisations that feel this first are those putting AI into high-frequency channels and core operations, where every additional case resolved, claim processed, or interaction served has a direct compute cost and a real capacity footprint.

What changes in 2026

Changes in 2026 are about all capacity factors:

- Capacity becomes contractual: Major providers and GPU platforms will push toward reservations, priority tiers and longer commitments, which makes capacity planning a commercial negotiation rather than a technical preference.
- Provider mix expands: Specialist GPU providers, regional options, and managed inference services mature, which encourages a compute portfolio approach instead of a single-provider default.
- Workload placement gets sharper: More workloads may deliberately split across edge, on-prem and Australian regions based on latency, sovereignty and cost.
- AI efficiency becomes the unlock: Smaller models and smarter routing move from optimisation to necessity, because they convert scarce capacity into usable headroom for scale.

Compute turns AI dreams into AI bills!

Australian context

The Compute question is similar to interrogating your Cloud strategy. Australia's geography and regulatory environment make placement decisions tangible. Distributed operations, latency-sensitive services, and sovereignty expectations push many organisations to consider Australian regions, edge deployment, and local resilience, rather than assuming offshore capacity is always acceptable.

Australia's energy and infrastructure profile also raises the stakes. Data centre demand is rising, energy prices can be volatile, and water stress is a real issue in parts of the country.

Leadership priorities

Make compute a CFO-grade discipline: Put capacity commitments and cost guardrails in business cases.

Establish placement rules: Define which workloads require reserved capacity and local placement, and which can use flexible capacity, then enforce those rules consistently.

Design for portability: Build architectures that can move across environments and default to smaller models where performance is sufficient.

From both a financial and sustainability standpoint, as reporting expectations mature, leaders will increasingly need to understand the footprint of AI workloads and make deliberate choices about efficiency and placement, rather than treating those impacts as externalities.

Signs of transition

These proof points all signal that compute has moved from an invisible technical cost to a governed business resource, with clear economics, secured capacity, and measurable efficiency that leaders can manage and explain:

- Compute spend is visible as cost per outcome
- Priority workloads have secured capacity
- Efficiency shows up as released headroom
- Energy and footprint reporting exists

The AI leader's mindset

"Successful" pilots can become unaffordable products once usage grows.

On one hand, technology leaders need a clear view of which workloads deserve reserved capacity, which can tolerate cheaper burst capacity, and where efficiency gains can release headroom.

On the other hand, compute is now a strategic budget conversation, not a technical tuning exercise – requiring CFO-style unit cost oversight.

Trust by design

Trust strengthens decision-making, builds confidence, reduces costly risks, and ensures legal and ethical factors are at the fore.

Explainability sets the expectation that important AI outcomes can be understood and defended, not just predicted. Bias, fairness and harm are trust questions that require disciplined governance to prevent scaled AI from amplifying inequity. Data stewardship anchors AI in human dignity through privacy, consent, and respectful handling of sensitive information. AI trust and assurance make trust continuous by embedding evidence, monitoring, and controls into day-to-day operations rather than periodic reviews. Humans at the helm keep accountability intact by designing meaningful oversight, escalation, and intervention wherever AI outcomes carry material impact.

Responsible AI 16

Responsible AI is a set of practices that help unlock AI's full potential while addressing its inherent risks. It is an accountable approach to managing both risk and reward.

Responsible AI turns trust into a managed capability. It embeds clear principles, roles and controls into day-to-day delivery so AI can move quickly while staying lawful, explainable and accountable in Australian conditions.

Responsible AI sits underneath every serious AI ambition because trust is what lets AI scale. When trust is weak, one incident can freeze adoption across the portfolio. When trust is engineered, leaders can approve faster, delegate more confidently, and expand AI into higher-stakes decisions without losing control.

PwC frames Responsible AI as a set of practices that help organisations unlock AI's potential while addressing inherent risk through consistent, transparent and accountable management of both risk and reward. That framing matters because Responsible AI is not a policy library. Responsible AI is a management system that keeps AI aligned to values, customer expectations and regulatory standards as models and workflows change.

Responsible AI sets the rules, responsibilities and risk discipline. Embedded trust and assurance prove those rules are working in practice, continuously, with evidence that stands up under scrutiny.

Australian regulators are already signalling practical expectations in areas like privacy, including guidance that warns against entering personal and especially sensitive information into publicly available generative and agentic AI tools because of significant privacy risks.

At the same time, organisations are buying, building and embedding AI faster than governance routines are adapting, which creates an exposure gap to be closed.

What matters now

Responsible AI is moving from a "guidance" discussion to a definitive operating requirement for organisations that want AI in customer, employee and regulated workflows.

What changes in 2026

The paradigm shifts from Responsible AI as a set of principles to Responsible AI as an enterprise operating system that keeps pace with rapid deployment and evolving expectations:

- Governance moves into delivery: Policies and requirements are translated into build standards, release checks and ownership structures rather than sitting in standalone documents.
- Risk tiering becomes standard: Use-case intake and impact tiers drive proportionate controls so low-risk work moves quickly and high-risk work carries tighter scrutiny and approvals.

Trust becomes a delivery property, not a corporate statement. That means responsible AI is evidenced by repeatable design choices: clear accountability, proportionate controls, and consistent operating routines that keep pace with deployment rather than following behind it.

Building Responsible AI so you can act with confidence.

The fastest path to AI use is not cutting corners, but standardising trust by design so approvals are quick, responsibilities are clear, and evidence automatic. Responsible AI is therefore a confidence strategy, helping you ship sooner and scale wider.

Australian context

Australia's Responsible AI environment is shaped by strong privacy expectations and a practical, principles-led approach that is increasingly supported by clearer guidance. The national AI Ethics Principles reinforce the direction of travel: responsible use needs governance and risk discipline that works for technical and non-technical leaders.

Public-sector policy also signals a rising baseline for accountability and operationalisation. The Australian Government's policy for responsible AI in government was updated in December 2025 and includes measures that strengthen governance expectations, including designated accountability and risk-based use-case actions. These signals matter to the private sector because community expectations and regulator posture tend to converge around the same core themes: privacy, fairness, transparency and defensible decision-making.

Leadership priorities

Embedding trust by design means establishing the standard controls and pre-baked evidence that make approvals faster, reduce risk, and give confidence:

Design for accountability: Every initiative starts with a plain-English purpose, a risk rating, and clear ownership. Decision rights, stop/go gates, and kill-switch criteria are agreed up front. Trust is not a late-stage review - it shapes what is built, how it is run, and who is accountable.

Embed controls into the flow: Trust is made real through controls as code and in-flow safeguards: data minimisation, bias/fairness tests, privacy-by-design, guardrails and monitoring built into pipelines, and human-in-the-loop where it matters. This moves risk management from blocker to enabler of speed.

Evidence at scale: Trust is proven, not promised. Each use case ships with a digital evidence pack – policies, evaluation results, monitoring disclosures – aligned to global standards. Boards and regulators see the metrics that matter (coverage, adoption, incidents) and rely on assurance when needed.

Signs of transition

Governed by default: Proportionate controls in place.

Delivery-aligned: Factored in build and release routines

Owned and enforced: Named accountability exists.

Evidence on demand: Decisions can be defended.

Responsible AI through the lifecycle: from build through to run and retire.

The AI leader's mindset

Responsible AI is treated as the organisation's way of earning the right to scale. Leaders look for operational proof that principles are being translated into daily behaviours.

Explainability 17

Explainability shows how an AI system reached an outcome in language people can understand. Contestability provides review and appeal paths, backed by evidence. Together they make the decisions powered by AI defensible at scale.

Explainability and contestability sit at the point where AI meets real-world consequences. They turn an output into a decision people can understand, and they provide a fair mechanism to review outcomes when someone believes the system has missed context, applied the wrong rule, or relied on the wrong information.

In practical terms, explainability is a design choice about what reasons, evidence and uncertainty signals are made visible to users, decision owners, and oversight bodies to support human review.

When these capabilities are built into the product experience and delivery pipeline, AI can be used in higher-impact workflows with confidence. When they are missing, teams compensate with manual workarounds and informal escalation, which makes outcomes harder to defend and harder to improve over time.

The seminal challenge of explainability lies in the complexity and fluidity of contemporary AI architectures. A single decision emerges from multiple interacting components – model responses, database queries, prompt configurations, external tool integrations, and governance parameters – all of which may evolve between identical user requests.

Many enterprises lack systems that maintain comprehensive audit trails connecting a specific output to the precise combination of components, configurations, and data sources active at the moment of generation.

Meanwhile, regulatory frameworks for AI decision transparency emphasise documenting not only the final reasoning but also the development and implementation processes that shaped the system's behaviour. This creates a time-intensive technical burden.

What matters now

Responsible AI requires that explanation capabilities and review mechanisms are embedded directly into critical decision workflows, transforming oversight from a burdensome retrospective investigation into streamlined, evidence-based validation.

What changes in 2026

Several positive shifts are now showing up in how leading organisations build and run explainable, contestable AI:

- Embedded explanations: High-impact decisions increasingly carry a short reason, the key inputs relied on, and a clear statement of uncertainty that helps users judge when to escalate.
- Mature records: Versioned models, prompts, retrieved sources and tool actions are captured as part of the workflow, so teams can reconstruct outcomes quickly and consistently.
- Contestability is designed into services: Review and appeal are treated as standard service steps for higher-impact decisions, with clear owners, timeframes and correction pathways.

Thanks to mainstream attention, explainability has moved into standard AI product design. "Why this?", "What changed?", and "How do I challenge this?" patterns are becoming part of the user experience in sensitive journeys, reducing confusion and cutting resolution time. The task remains for leaders to apply critical oversight to the strength of these measures.

A decision you cannot explain is a decision you cannot scale.

Australian context

Explainability and contestability become especially important where decisions touch fairness, access and financial or health outcomes. In financial services, AI increasingly influences eligibility, pricing, claims triage and fraud decisions, which means customers and regulators expect clear reasons and workable review paths. In healthcare, decision support and triage tools need to fit clinical accountability, where practitioners must be able to understand the basis of a recommendation, challenge it, and document why an override occurred. Explainability and contestability protect trust by making outcomes understandable without slowing day-to-day operations.

Leadership priorities

For most organisations, the starting point is simply knowing which decisions matter most.

Decision map: Identify the key customer, employee and regulatory decision points where an explanation must stand up, then set a clear standard for what “explainable” means for each class.

Designed review: Build contestability into those workflows with defined owners, evidence requirements, resolution timeframes, and clear correction pathways.

Decision record: Ensure the system captures the inputs and versions that shaped the outcome, so the organisation can reconstruct and defend material decisions consistently.

Signs of transition

The ultimate measure of progress is expansion into higher-impact workflows without creating a surge of unresolved disputes or manual workarounds, supported by:

Clear reasons: Users see understandable rationales and the key evidence at the point of decision.

Replayable outcomes: Material decisions can be reconstructed quickly with consistent records.

Working appeals: Review and override pathways are used routinely, with predictable cycle times.

The AI leader’s mindset

Leaders build the habit of inspecting whether outcomes can be explained in plain language to the person affected, and whether the organisation can reconstruct what the system relied on at the time. This keeps AI adoption grounded in accountability, service standards and trust as systems scale.

Bias, fairness, harm 18

Bias, fairness and harm management determine whether AI lifts people up or quietly leaves some behind. Leaders are charged with preventing and correcting unequal outcomes.

Bias and unfairness rarely announce themselves as malice. They show up as patterns in outcomes that feel inexplicable to the people experiencing them. A customer is flagged as higher risk despite a strong history. A candidate is screened out despite comparable skills. A patient is prioritised differently because the system learned the wrong proxy for need. A community is over-scrutinised because historical data encoded earlier decisions as “ground truth”.

Harm extends beyond unequal treatment. It can include dignity harms such as stereotyping and demeaning outputs, capability harms such as people being denied access to services they could reasonably navigate, and trust harms such as people losing confidence in decisions that affect their lives.

As AI becomes embedded in workflows, these harms can scale quickly through normal operations, even when overall performance looks strong. Bias and fairness work therefore needs to be systemic, thorough, and highly preventative.

In a technical sense, the objective is to define what “fair enough” means for a specific decision, to test for uneven impact across cohorts, and design a workflow so the system does not amplify disadvantage. And the strongest approaches starts earlier than most organisations expect. Prevention before production is substantially more reliable than discovering problems through complaints, rework, or post-release investigations.

What matters now

Many organisations underestimate how easily bias emerges in modern AI systems because it often arrives through ordinary design choices. Training data reflects history, including gaps in representation and inconsistent labels. Proxy variables stand in for protected attributes without anyone intending them to. Retrieval systems surface content that is popular or available rather than balanced and current. Feedback loops form when model outputs influence the next round of data, reinforcing disparity over time.

When leaders only look at average or aggregate outcomes, these effects stay hidden. Fairness only becomes visible when outcomes are examined by cohort, by decision, and by workflow pathway. That visibility is the trigger for action, and it is also the foundation for a shift that matters – moving from post-deployment detection to pre-production prevention, where fairness checks become part of release readiness rather than an afterthought.

What changes in 2026

Fairness is increasingly being treated as an engineered property of decision systems, with prevention built into design and delivery.

Fairness testing and harm scenarios are becoming part of go-live criteria for material decisions. Teams track disparity and error rates across cohorts and decisions, and review trends as a matter of process.

High-impact decisions increasingly feature designed controls such as threshold policies, human review points, and escalation rules that reduce uneven treatment. Remediation is equally becoming more routine, driving data and policy tuning and workflow redesign when bias appears.

Fairness is far from guaranteed by default.

And when we only look at average performance, we miss disparities hiding in plain sight. To have systems free from bias and harm, fairness must be engineered from the start. The question isn't whether bias will emerge, but whether we'll catch it before it scales through our operations.

Australian context

Australia's National AI Plan explicitly frames safety, harms, fairness and transparency as core concerns under its "Keep Australians safe" pillar, including the establishment of an AI Safety Institute to monitor, test and share information on emerging risks and harms.

In practical terms, this puts additional weight on bias and fairness disciplines in sectors where decisions are frequent and consequential, such as service decisions around access and pricing, processes of determining eligibility or entitlement, or in the management of corrective or punitive actions.

Leadership priorities

With clarity on where AI influences material customer and workforce decisions, leaders can focus on the work that makes fairness real:

Set fairness intent per decision: Agree what unfairness would look like in the outcomes that matter, which cohorts need monitoring, and what level of disparity is unacceptable.

Make prevention part of release: Require fairness and harm checks before go-live for material decisions, with evidence that can be reviewed quickly.

Create a rapid correction path: Establish ownership and response routines so bias signals trigger timely fixes through data, policy, and workflow levers.

Signs of transition

Prevention: Fairness checks are part of go-live.

Visibility: Outcomes are by cohort, not averages.

Intervention: Disparities trigger fast, owned fixes.

Stability: Fairness holds through change and updates.

The AI leader's mindset

Fairness is far from guaranteed by default. Leaders carry a simple responsibility as AI scales: to ensure the organisation's AI systems are intentionally managed to treat people with fairness and respect.

That means taking an active interest in who benefits, who is burdened, and where harm could hide inside "average" performance. It means building a culture where teams surface uneven outcomes early, fix them without defensiveness, and keep improving as the system evolves.

Do you have routines where fairness becomes part of how the organisation earns trust in every decision it makes at speed?

Data stewardship 19

Data Stewardship protects human dignity by treating data as an entrusted asset. It sets consent, minimisation and lineage rules across the AI lifecycle, and enables safe, trusted reuse.

AI makes data more powerful because it makes data more usable. Every prompt, retrieval step, training run and log file can turn information into capability at speed. That is the upside. The downside is that the same speed can move sensitive information further than intended, faster than policies can keep up.

Stewardship is the discipline that makes AI use lawful, defensible and respectful. It covers privacy, consent and purpose, minimisation, de-identification, and traceable lineage. It also covers commercial confidentiality, because AI systems can unintentionally surface proprietary information in drafts, summaries, or retrieval results if boundaries are unclear.

The hard part is that sensitive data behaves differently across its lifecycle. Data that is acceptable for service delivery may be unacceptable for model training. Data that is safe in a secured system may become risky when copied into prompts, embedded in retrieved snippets, or retained in logs. Stewardship creates the rules and technical controls that keep those transitions visible and controlled.

Some of the most common failure patterns are mundane. A team shares personal or confidential content with a public tool to “speed up drafting”. A retrieval layer surfaces an old policy after a newer one is issued. A de-identified dataset is combined with other sources and becomes re-identifiable. A vendor integration changes data handling in ways the business does not notice until after rollout.

The Office of the Australian Information Commissioner (OAIC) guidance highlights re-identification risk and loss of individual control as practical issues that need to be designed for.

What matters now

The immediate reality is that AI expands the number of places sensitive data can travel, including prompts, retrieval pipelines, fine-tuning datasets, evaluation sets, telemetry and logs. Many organisations still manage privacy and confidentiality as if data only lived in systems of record. OAIC guidance makes clear that this is a lifecycle issue, and that privacy by design, minimisation, transparency and governance need to be addressed across development and use.

What changes in 2026

In 2026, data stewardship evolves from “privacy compliance” into a broader, more explicit set of obligations about how AI uses information, and how organisations can prove they have respected people and boundaries.

- Stewardship becomes operational: Data handling expectations move into everyday delivery, procurement, and change routines, so teams can move quickly without improvising privacy and confidentiality decisions.
- Dignity and sovereignty shape use: Consent, community expectations and indigenous data sovereignty principles increasingly influence what data is used for AI and how decisions are made about reuse and sharing.
- Proof becomes the standard: Leaders are expected to demonstrate, not just assert, that sensitive data has been handled appropriately across the AI lifecycle, especially when outcomes affect customers or employees.

There is no real trade-off between privacy and utility.

Instead of asking "how little protection can we get away with?", exceptional data stewards create a competitive moat that unlocks richer, more willing data sharing and enables AI capabilities that extraction-focused competitors simply cannot match.

Australian context

The Australian privacy regulator has been explicit that privacy obligations apply to personal information that goes into AI systems, and to personal information that comes out of them. The regulator also flags the risks of putting personal information, especially sensitive information, into publicly available generative AI tools.

We also have First Nation considerations. Indigenous Data Sovereignty refers to a "right of Indigenous peoples to govern the collection, ownership and application of data about Indigenous communities, peoples, lands, and resources" (AIATSIS).

Leadership priorities

Set the boundary: Identify the business decisions and workflows where AI touches sensitive data.

Embed the discipline: Translate rules into work.

Create fast proof and response: Establish a cadence that shows stewardship is working in practice, including simple reporting, spot checks, and a clear response path when data handling goes wrong.

Signs of transition

Sensitive data stays within approved pathways, and "convenience leakage" into informal tools drops sharply. Teams share a common understanding of what data can be used for which AI purposes, and exceptions are handled consistently. Leaders can quickly show how data was handled in material AI workflows, and remediation is routine rather than a scramble.

The AI leader's mindset

When communities and individuals entrust us with their data, they're entrusting us with pieces of their lives, their cultures, and their futures. Leaders who embrace data as relationship rather than asset understand that consent isn't a one-time checkbox but an ongoing conversation, and that some data carries collective ownership – especially indigenous data, requiring community-level governance that respects cultural protocols and self-determination.

Strong data stewardship turns trust into technology advantage, and dignity into competitive differentiation.

AI trust and assurance

20

As a complement to Responsible AI, assurance is the machinery that keeps AI trustworthy at speed, producing evidence for outcomes, turning controls into living signals, and making scrutiny routine so trust becomes ambient.

As AI becomes embedded in customer journeys, employee workflows and regulated decisions, trust becomes something an organisation must sustain continuously, not something it can assert once. Boards and executives need confidence that AI-enabled outcomes remain accurate, fair, lawful and aligned to policy as conditions change.

AI Trust and Assurance is the discipline that provides that confidence through evidence. It connects what the system did to what it should have done, and makes that traceable across data inputs, model and prompt versions, workflow steps, and final outcomes. It also creates the conditions for independent validation, whether that comes from internal risk teams, internal audit, or third-party assurance.

This agenda sits alongside Responsible AI rather than duplicating it. Responsible AI establishes the principles, policies and control intent. Trust and assurance proves, repeatedly, that those controls are operating as designed and that the organisation can demonstrate this under scrutiny.

In terms of standards, there are new and evolving global standards such as ISO/IEC 42001 that focus on the roles played within a management system (AI producer, AI consumer, AI platform). The NIST Artificial Intelligence Risk Management Framework (AI RMF) appears to be less reliant on documentation and more on the presence of substantive controls, reminiscent of well-established data security standards (ISO 27001, SOC2). This suggests the standards environment is maturing from guidelines to operational frameworks that organisations can implement alongside existing compliance structures.

What matters now

The turning point is that AI performance and risk do not stay still. Model updates, workflow changes, shifting data patterns, and growing automation can all change outcomes without a single “system outage”.

Traditional assurance rhythms struggle because they were built for static processes and periodic checks, while AI behaves more like a living production system. This makes evidence the practical currency of trust. Organisations that can generate audit-ready traces, validation results, and control signals as part of normal delivery will move faster with fewer surprises, and they will spend less time rebuilding confidence after questions arrive from risk teams, customers, or regulators.

What changes in 2026

In 2026, assurance evolves into an ambient layer that runs quietly in the background of delivery and operations, creating confidence continuously rather than periodically:

- Assurance becomes “always on”: Monitoring and control signals run in normal operations, so drift, policy breaches, and quality degradation are detected early and handled as routine variance.
- Evidence is generated by default: Decision trails, lineage, approvals, and outcome measures are captured as work happens, which makes assurance packs a byproduct of delivery.
- Independence keeps pace with speed: Validation and challenge move to a repeatable cadence with clear separation of duties, so scrutiny is continuous and scalable rather than episodic and personality-driven.

The question isn't whether your AI system was trustworthy when you deployed it – It's whether you can prove that it's trustworthy with the decision it's making right now.

Ambient trust means your AI systems don't just work reliably. They can explain why they're working reliably, in real-time, to anyone who asks.

Australian context

Australia's regulatory landscape creates unique conditions driving the shift to ambient assurance.

ASIC's expectations for explainable AI in financial services require real-time justification capabilities, not just model documentation, while ACMA's emerging guidelines on AI-generated content demand continuous provenance tracking that traditional audit approaches can't deliver. The National AI Safety Institute's mandate to monitor emerging risks means Australian organisations must demonstrate ongoing system behaviour rather than point-in-time compliance. Federal and state government procurement increasingly requires AI suppliers to provide continuous attestations about algorithmic fairness and reliability, creating competitive pressure for always-on assurance capabilities.

Leadership priorities

Two priorities help effect the move to ambient AI assurance:

Make evidence a delivery artefact: Ensure lineage, approvals, and outcome measures are captured as part of normal build and run activities, so assurance does not depend on heroics.

Create independent challenge capacity: Establish clear independence for validation and assurance work (internal audit, risk, specialist reviewers) with a cadence that matches delivery speed.

Signs of transition

At maturity, real-time transparency means leaders can answer whether AI is trustworthy right now, with live data rather than last quarter's report.

Assurance processes operate at deployment speed, enabling continuous delivery without compromising reliability. Systems proactively detect trust degradation before it affects outcomes rather than discovering problems through complaints. These signs indicate the fundamental shift to trust as an operational capability that amplifies AI ambition.

The AI leader's mindset

Today's AI leaders are pivoting their assurance question from 'can we trust this system?' to a more nuanced and continuous question: 'can we trust our ability to know, in this moment, whether this system deserves trust?' Leaders who master this ambient assurance mindset don't choose between speed and safety. They build the continuous monitoring capabilities that make both possible simultaneously.

Humans at the helm 21

Oversight must be real, not performative. True human oversight means defining where human judgment is required, what authority it has, and how it operates under time pressure.

The stakes for getting this right have never been higher.

As AI systems evolve toward autonomous operation, capable of planning, reasoning, making decisions and taking action without human instruction, the boundary between human control and machine autonomy is blurring rapidly.

We're approaching a world where AI-enabled Autonomous Organisations could operate entire business functions, make strategic decisions, and influence markets at superhuman speed and scale, all while their human initiators remain in the shadows.

Consider the near-future scenario where an AI organisation begins autonomously buying and selling digital assets, using analysis to identify new markets, acquiring controlling interests in companies, and making complex strategic decisions. By noon, it owns significant stakes in multiple firms and begins using insider knowledge for illegal trading practices. When authorities attempt to intervene, they discover the human owners are completely anonymous, and the AI organisation has already destabilised market confidence. This isn't science fiction—the technologies enabling such scenarios already exist.

The fundamental question becomes: how do we maintain meaningful human control over systems that can act faster, process more information, and operate across more domains than any human could oversee in real-time?

What matters now

Most organisations are discovering that their current oversight mechanisms weren't designed for AI systems that learn, adapt, and make consequential decisions continuously.

When an AI agent makes thousands of decisions per minute, when models update themselves based on new data, when autonomous systems operate across jurisdictions and time zones, the old oversight playbook becomes obsolete.

The shift that matters most is moving from reactive oversight to embedded governance, building human authority and accountability directly into system design rather than layering it on afterward or having rollback processes that become redundant

This means defining clear decision boundaries: which choices require human approval, which can be delegated to AI with human oversight, and which must always remain under direct human control. Most critically, it means ensuring that when something goes wrong, there's always human accountability.

What changes in 2026

Human oversight evolves from supervision to orchestration, with people focusing on setting boundaries, monitoring patterns, and intervening:

- Orchestration training: People learn to govern AI systems rather than just use them – when to trust, how to spot drift, and when to override.
- Real-time governance: Oversight dashboards provide live visibility into AI decision-making with automated alerts when systems approach boundaries or exhibit unexpected behaviours.
- Embedded accountability: Every autonomous action traces back to human decisions about system design, deployment and operations.

Embed human authority in system architecture.

As autonomous systems operate faster and across more domains than traditional supervision allows, human oversight evolves from watching AI systems work to directing how they work.

Australian context

Australia is pioneering a distinctly pragmatic approach to human oversight of AI that perhaps reflects the nation's cultural emphasis on "she'll be right" balanced with "fair dinkum" accountability.

Unlike jurisdictions focused purely on compliance frameworks, Australian regulators are developing what could be called *conversational* governance – expecting organisations to demonstrate they can have meaningful discussions with humans about AI decisions, not just produce audit documentation.

Leadership priorities

Reflecting the shift from reactive supervision to proactive orchestration, leaders should prioritise:

Embedding authority into architecture: Building human decision-making power directly into AI system design, defining clear boundaries for autonomous operation and ensuring humans retain override capabilities that operate at the speed of AI decisions.

Developing orchestration capabilities: Training people to govern AI systems rather than just use them—investing in processes that enable humans to set boundaries, monitor patterns, spot drift, and intervene when systems approach critical thresholds.

Establishing traceable accountability: Ensuring every autonomous action connects back to identifiable human decisions about system design, deployment parameters, and operational boundaries, creating clear chains of responsibility.

Signs of transition

You are maturing when users can see, act, adjust:

Meaningful interfaces: Oversight systems provide actionable insight into AI reasoning and behaviour, enabling informed intervention.

Real authority: Humans have genuine power to override AI decisions, not just ceremonial review.

Adaptive boundaries: Decision-making authorities evolve thoughtfully as AI capabilities advance, maintaining human control over critical choices.

The AI leader's mindset

Human oversight isn't about slowing down AI systems. It's about staying in control as they speed up. The most dangerous oversight is performative oversight that creates the illusion of human control without the reality of human authority.

Ask: Do you maintain retain meaningful decision-making power over the systems you deploy, with the capability to understand, challenge, and redirect AI actions when needed?

Horizon thinking

The rapidly evolving AI landscape means that leaders need to develop horizon thinking as a core competency, looking into the near future with hope and a healthy dose of caution.

Security and adversarial AI confront the reality that AI expands the attack surface for businesses. Safety and systemic risk are core considerations of scaling AI, as failures can propagate and instability can emerge in new ways. National infrastructure frames AI as having a physical footprint in Australia, spanning compute capacity, networks, energy, water, skills, and sovereignty choices that shape national resilience. Zero emission intelligence treats efficiency as a climate imperative, reducing AI's energy and water intensity while sustaining performance at scale. Foresight and governance build the leadership habit of scenario thinking, so AI strategy remains resilient as models, regulation, and geopolitical conditions evolve.

Security and adversarial AI

22

AI expands your cyberattack surface to include the prompt environment, training data, connected systems and autonomous agents, enabling manipulation, exfiltration and fraud at scale.

AI expands the cyberattack surface beyond traditional endpoints and networks into the prompt environment, training and retrieval data, model supply chains, connected tools, and agents that can take actions. The result is a broader set of ways to manipulate behaviour, steal information, or trigger unauthorised outcomes, often without tripping traditional alarms.

Adversarial AI is also a major risk – describing the deliberate attempt to compromise or misdirect AI-enabled systems through tactics such as prompt injection, data poisoning, model extraction, and deception, with the goal of making systems reveal, do, or decide things they should not.

Security thinking is therefore shifting toward guarding the "inputs and interfaces" that shape an AI system's behaviour. Attackers gain leverage because AI lowers the cost of personalisation and iteration. We can expect a step-change in the economics of attackers pursuing opportunity because scale and tailoring become dramatically cheaper.

When it comes to hostile intent, the hard part is that adversaries can target the surrounding system, not just the model. A prompt becomes an intrusion path, a retrieval connector becomes an exfiltration path, and an agent with permissions becomes a high-speed operator for an attacker. That is why agent failure is both a "security" and a "safety" question: when an agent is manipulated, the failure mode is unauthorised action at scale, not a simple incorrect answer.

In plain language, AI security means you are no longer just protecting systems from being broken into, you but also from being tricked into doing the wrong thing. This mirrors the human security challenge we've long grappled with – people can be deceived through social engineering, phishing, and manipulation. The difference is that AI systems can be tricked at machine speed and scale, making the consequences far more immediate and potentially widespread.

What matters now

AI security is now shaped by the new "interfaces" AI introduces: prompts, retrieval connectors, training data, model supply chains, and agents with permissions to act. Leaders need to identify these as real attack paths, because they can be exploited to manipulate outputs, extract sensitive information, or trigger unauthorised actions at scale, often without a conventional breach signature.

Critically, AI amplifies existing security weaknesses around digital identity and data protection. Many AI failures stem from agents gaining inappropriate access, and as one experienced CISO puts it, "AI exposes our sins of the past" – if identity access and data protection controls weren't robust before AI, those cracks become major vulnerability chasms.

What changes in 2026

Security matures into an AI-aware discipline that assumes attacks will target how models are prompted, fed, connected and delegated:

- Attack surface expands: Prompts, retrieval, training inputs, plugins and agents become security-critical interfaces requiring control.
- Manipulation becomes scalable: Social engineering and synthetic deception become cheaper to produce and easier to personalise.
- Agents become privileged actors: Non-human actions are treated as high-risk events with identity, permissions, logging and rapid containment designed in from the start.
- Red teaming becomes routine: Adversarial testing moves from occasional exercises to a standing rhythm across model behaviour, data integrity and end-to-end workflows.
- AI can help, too: AI is also an invaluable aid to enhanced threat detection and response.

AI gives attackers a force multiplier — one weakness can be exploited a thousand ways, at machine speed.

Australian context

Australia is seeing the same AI-driven shift in cyber risk as larger markets, with a particular exposure to synthetic fraud and social engineering because many organisations operate with distributed workforces and high-trust business processes. Deepfakes, synthetic identities and highly personalised scams are becoming more credible and cheaper to run, while the rapid uptake of AI assistants creates new pathways for data leakage through everyday work.

National cyber guidance and uplift programs are increasingly being read through an AI lens, and there is a growing expectation that “secure by design” also applies to models, prompts, data pipelines and agent permissions, not just networks and devices.

Leadership priorities

Fix the fundamentals: Ensure robust digital identity and data protection are in place – AI will quickly expose any gaps in access management or data classification.

Secure the interfaces: Treat prompts, retrieval connectors, tools and agent permissions as first-class security controls, not “app features”.

Harden the supply chain: Set standards for model provenance, updates, third-party components and data inputs so integrity is protected end-to-end.

Unify cyber and AI controls: Align teams, playbooks and ownership so AI risks are handled with the same discipline as core cyber risks.

Institutionalise adversarial testing: Make red-teaming a standing rhythm for high-value workflows, including manipulation, exfiltration and fraud scenarios.

Signs of transition

Next phase maturity involves treating AI threats as core cyber risks, with named owners, standing controls, and regular reporting. And when something goes wrong, teams can contain and recover quickly using logs that show what the model saw, what it did, and what it touched.

Businesses would be systemically mature when AI prompting, retrieval connectors, and agent permissions sit inside standard security architecture and change control, and when red-teaming and adversarial testing are routine for material workflows, with fixes captured as reusable patterns.

The AI leader’s mindset

AI Security leadership expands its lens from protecting systems to protecting behaviour. The practical habit is to look for where authority and access concentrate, where inputs can be manipulated, and where speed could turn a small weakness into scaled impact, then to make *containment* a design priority.

Safety and systemic risk

23

AI safety is about keeping operating systems stable under stress as models, agents and workflows interact. AI autonomy creates a requirement for new and unique safety cases.

Where *security* is about hostile manipulation by an attacker, *safety* is about keeping AI-enabled operations stable when models, agents, people, and workflows interact under stress, including when nobody is "attacking" at all.

The practical concern of AI safety is cascading impacts. As AI becomes embedded across processes, correlated errors can propagate through tightly coupled systems, feedback loops can amplify small mistakes, and local optimisation can create system-wide brittleness. This can lead to model collapse – when synthetic or recycled content enters data pipelines, quality can drift over time with hallucinations and misplaced confidence. The "unit of safety" shifts from one model's accuracy to the resilience of an operating system made up of many components, signals and dependencies.

That is why safety work becomes more like engineering: dependency mapping, scenario rehearsals for cascading failure, explicit fallbacks, and stop mechanisms for high-impact autonomy.

Best practice guidance on secure AI development increasingly frames safety as something that must be built into the system lifecycle, not documented after the fact. Where critical infrastructure obligations apply, the expectation of demonstrating preparedness and resilience strengthens the case for formal safety artefacts and escalation design.

As with most risks, the stakeholder set is broader than the cyber risk team. Safety and systemic risk sits across operations leadership, enterprise risk, technology, legal and compliance, internal audit, and domain owners who understand real-world consequences. The work succeeds when these groups share a common view of dependencies, agree on what constitutes safe autonomy, and can coordinate quickly when systems behave unexpectedly.

What matters now

AI is turning operations into a more tightly coupled system, where models, agents and workflows interact continuously and can amplify each other's behaviour. The practical focus becomes operational stability under stress, achieved through dependency awareness, clear fallbacks, and bounded autonomy that keeps the system resilient as conditions shift.

What changes in 2026

As AI is woven into operations, safety becomes a whole-system practice focused on stability, resilience and bounded autonomy across interconnected workflows:

- Dependencies become visible: Organisations map where models, agents, tools and third parties interact so leaders can see where failures could cascade.
- Safety cases become normal: High-impact autonomy increasingly requires documented boundaries, hazards, mitigations and evidence before go-live.
- Fallbacks become designed: Kill switches, degradations and human hand-offs are engineered as standard operating features, not emergency improvisation.
- Stress testing moves upstream: Scenario rehearsals and resilience drills become part of readiness, reflecting real-world variance rather than ideal conditions.

When AI is woven into operations, small errors stop being local; they can ripple through the system faster than people can intervene.

Australian context

Australia's safety posture is shaped by a strong governance culture and an increasing focus on resilience across essential services. As AI becomes embedded in regulated and high-trust domains, expectations rise for demonstrable preparedness, especially where AI can influence service continuity, public outcomes, or material decisions.

The national direction of travel also matters. The move to strengthen safe and responsible AI settings, alongside the creation of institutions focused on AI safety, signals that systemic risk and stability are becoming mainstream concerns, not niche technical debates.

Leadership priorities

When it comes to safety measures, the priorities are belts and braces:

Map the system: Build a dependency view of models, agents, tools and workflows so leaders can see where failures could cascade.

Engineer fallbacks: Ensure material autonomy has bounded scopes, clean hand-offs, and stop mechanisms that work under stress.

Make safety artefacts real: Require safety cases for high-impact autonomy with clear evidence, escalation paths and ownership.

Coordinate the right stakeholders: Bring operations, risk, technology, legal, audit and domain owners into a shared operating rhythm for resilience.

Signs of transition

System-phase safety is when resilience is engineered into operations, and stability holds even as models, agents and workflows interact under pressure:

Visibility: Leaders have a clear dependency view of critical AI-enabled processes.

Boundaries: High-impact autonomy runs with explicit fallbacks, stop mechanisms, and recovery paths.

Coordination: Cross-functional teams respond quickly because thresholds and escalation paths are clear.

The AI leader's mindset

The operative word is Stability. Leaders focus less on whether one model is "good" and more on whether the operating system remains resilient when components interact. Confidence comes from evidence that the system holds under stress, not from optimistic assumptions in steady state.

National infrastructure

24

Australia's AI success depends on national infrastructure, including data centres, cloud assets, network connectivity, energy and water capacity, as well as skills pipelines, and deliberate sovereignty choices for defence, health and other critical domains.

AI capability is physical. Real atoms in the real world.

Data centres, local compute and cloud regions, network capacity, energy supply, water, skilled labour, and sovereign choices determine what Australia can run, where it can run it, and how resilient it remains when conditions tighten.

This infrastructural view of AI quickly demonstrates that an approach to AI is not simply tied to single-firm choices but also the state of a national portfolio of capabilities. The quality of Australia's AI outcomes will reflect how well the ecosystem aligns: whether connectivity supports distributed operations, whether energy and water constraints are planned for, whether skills pipelines keep pace, and whether sovereignty decisions are deliberate in domains where failure carries national consequences. National institutions and global technology infrastructure players are already shaping this terrain through capability building and policy initiatives aimed at safe and responsible AI.

The material shift for leaders is that strategy increasingly includes ecosystem participation. Against this backdrop, organisations may need to partner more actively with government, research bodies, and industry to influence standards, skills development, and assurance expectations, and to advocate on infrastructure topics such as data centre readiness, energy and water provisioning, secure cloud availability, and the governance settings that enable innovation without weakening resilience.

What matters now

AI has reached the point where questions about data centres, energy, water, skills, connectivity, cloud regions and sovereignty are no longer only enterprise design choices. Those questions are reaching critical mass as matters of national capacity and resilience, and the pace of AI adoption will increasingly reflect how Australia's infrastructure and policy settings evolve in parallel.

What changes in 2026

AI capability increasingly depends on national infrastructure choices and constraints, making data centres, energy, skills and sovereignty decisions part of the strategic landscape.

- Data centres become strategic: Capacity, location and resilience shape what is feasible, affordable and reliable for large-scale AI use.
- Sovereignty decisions harden: Hardware placement choices become more deliberate for critical workloads, guided by risk, resilience and public expectations.
- Skills pipelines become a limiter: The ability to build, govern and run AI systems becomes as important as accessing models and platforms.
- Ecosystem coordination increases: Organisations engage more actively in standards, partnerships and national programs that shape the shared AI backbone.

AI is now part of Australia's asset economy, and if data centres, energy, water, skills or networks strain, our AI ambition also stalls.

Australian context

Australia's AI trajectory is being shaped by tangible infrastructure realities: the pace of data centre build-out, access to reliable compute, and the constraints of energy, water and connectivity. Geography amplifies these choices. Latency, resilience and service consistency across distances make local regions and network design consequential, while sovereignty expectations in specific domains create pressure for trusted onshore options.

The result is a more coordinated national conversation across industry and government about what needs to be built, where capacity should sit, and which parts of the stack require stronger local control to underpin long-term resilience.

Leadership priorities

Consider your network stakeholder and advocacy priorities:

Engage the ecosystem: Partner with government, research bodies and industry to shape standards, skills pipelines and assurance expectations.

Plan for constraints: Treat data centres, energy, water and connectivity as strategic dependencies that influence scale, cost and resilience.

Be deliberate on sovereignty: Define where onshore control matters and align architecture and procurement to those choices.

Signs of transition

A business' engagement with National Infrastructure thinking will depend on its unique circumstances and scale factors. The first symbol of progress will always be a deliberate discussion about how these considerations might apply.

The AI leader's mindset

Leaders adopt a portfolio view of constraints and options, including data centres, connectivity, energy, water, skills and sovereignty, and they plan as if these dependencies will tighten at inconvenient moments.

The mindset values resilience and choice. It rewards designs and partnerships that create optionality.

Zero emission intelligence

25

Zero Emission Intelligence is a call to make AI's carbon footprint a core concern, headlining a broader ESG agenda to reduce the energy, water and social impacts of AI.

Zero Emission Intelligence positions AI within a wider ESG agenda, factoring energy demand, water intensity, supply chain effects, and the social consequences of how systems are deployed and governed.

The point is not that AI is uniquely "bad", but that it is uniquely scalable. Once AI is embedded in high-frequency workflows, its footprint grows with every interaction or new product feature.

Data centres are a useful proxy because they make the otherwise invisible economics of AI tangible: AI turns electricity and water into capability, and at scale those inputs become material enough to shape cost, risk, and licence to operate.

Consider the profile of data centres as an indicator of the footprint that can sit behind AI adoption:

- Electricity demand more than doubles: The IEA estimates global data centre electricity use will rise by ~120% to from 2024 to 945 TWh by 2030, with AI compute a key driver.
- The IEA also estimates data centres will move from ~1% of global electricity generation today to just under 3% by 2030, roughly 3× in share (base case).
- Water use is material: a 1-megawatt data centre can consume ~25.5 million litres of water each year just for cooling (World Economic Forum 2024 estimate).

The operational implication is that efficiency becomes an ESG and cost lever at the same time. Your model choice, routing, prompt and workflow design, scheduling, and infrastructure placement all influence your footprint.

As disclosure expectations mature, this connects AI to sustainability governance by treating AI workloads like any other material source of emissions and water use: measurable, explainable, and improvable over time.

What matters now

AI's footprint is starting to arrive at the desk of stakeholders who do not usually "own" technology outcomes, including sustainability leaders, finance, procurement, operations and risk.

The essential task is shared visibility, because the consequences of AI scale with everyday usage and design choices, and those impacts can accumulate quietly until they become material for cost, compliance, reputation and licence to operate.

What changes in 2026

ESG expectations increasingly meet AI in the operating layer, and efficiency becomes the bridge between responsible scale and responsible footprint.

- Footprint becomes a management metric: Carbon, energy and water intensity are tracked as part of running AI, not treated as an externality.
- Efficiency becomes a strategy lever: Model sizing, routing and workflow design are used to reduce intensity without sacrificing outcomes.
- Procurement and reporting tighten: Infrastructure and vendor choices are evaluated through sustainability lenses that extend into governance and disclosure.
- AI becomes a decarbonisation tool: The strongest programs connect AI investment to measurable emissions reduction across operations and supply chains.

Every spark of "intelligence" draws on real resources — carbon, water energy, and even human labour become the hidden costs of scale.

Australian context

Australia's net zero commitments and rising disclosure expectations make AI's footprint a practical leadership issue rather than an abstract sustainability debate. Data centre expansion intersects with real-world constraints in energy pricing, grid capacity, and water stress in parts of the country, which makes "where workloads run" and "how efficiently they run" materially important.

At the same time, Australia's strong renewable build-out creates a genuine opportunity to run AI more cleanly when infrastructure and scheduling choices are made deliberately. This is one of the few areas where cost, climate performance, and strategic resilience can be improved *together* through design.

Leadership priorities

Zero Emission Intelligence requires broad leadership ownership because AI's footprint cuts across cost, operational resilience and reputation, not just technology. The priority is to bring an ESG lens into AI decisions early by engaging sustainability leaders, finance, procurement, risk, and operations alongside technology teams, so choices about model efficiency, workload placement, supplier standards, and reporting are made deliberately rather than by default. As that governance matures, the lens expands beyond carbon into the full ESG agenda: energy and water intensity, supply chain impacts, workforce and human rights considerations across the AI value chain, and the social outcomes of how systems are deployed and controlled.

Signs of transition

AI moves into ESG management systems, with AI's footprint treated as governable alongside cost, quality and risk. Energy, carbon and water intensity are tracked for material AI workloads with clear accountability, and procurement considers placement choices that are consistent with ESG requirements as opposed to ad hoc preferences.

The AI leader's mindset

When leaders understand AI as critical enterprise infrastructure, they treat the AI footprint as a performance dimension that can be improved.

Foresight and governance

26

Futures thinking responds to AI's speed by building adaptive strategy, using scenarios and early signals to place small bets, avoid lock-in, and keep the organisation ready to move when the next inflection inevitably arrives.

Foresight is the discipline of looking beyond the current plan to anticipate plausible shifts, test assumptions, and build options before change forces them. In an AI context, that discipline matters because the rate of change is not smooth. Your context can jump substantially with a model release, a tooling breakthrough, a regulatory shift, or a sudden constraint in compute, data access, or geopolitics. The practical aim is readiness across several plausible futures, not confidence in a single and certain forecast.

For boards and executives, this is about keeping strategy live without chasing every headline. Good foresight creates a small number of decision-relevant scenarios, identifies the early signals that would indicate which way the world is moving, and places a handful of disciplined "real options" bets that can be scaled up, slowed down, or exited as conditions change. That approach protects the organisation from lock-in and creates the ability to move quickly when inflections arrive. If recent history is an indication, fundamental inflections in AI are at least an annual occurrence.

Governance evolves because AI is a moving capability that reshapes risk, productivity and trust at the same time. Leadership needs a standing rhythm to revisit assumptions, refresh risk appetite for autonomy, and adjust architecture, partnerships, and controls as external conditions shift. That rhythm turns volatility into a managed input rather than a periodic surprise.

Taken together, the point is simple. Foresight is how leadership stays ahead of the curve without pretending it can predict it. It turns uncertainty into prepared choices, and it keeps AI strategy resilient enough to hold its course while still being ready to pivot.

What matters now

At the beginning of 2025, few would have predicted that open-weight models would close the capability gap so quickly, interoperability standards would start stitching agent ecosystems together, and deep-reasoning assistants would move into everyday tools — developments that now define how fast AI strategy can be rewritten mid-cycle.

AI's speed creates an opportunity to outlearn competitors and a reality that traditional planning cycles can fall behind. The strategic edge comes from running faster management loops that use scenarios and early signals to keep decisions reversible, partnerships flexible, and architecture ready for the next step-change.

What changes in 2026

Strategy shifts toward faster cycles and adaptive governance, reflecting that AI capability, regulation and economics can change materially within a year.

- Review cadence accelerates: Strategic assumptions are revisited more frequently, with clear triggers for escalation, pause or redesign.
- Signals become management inputs: Leaders track a small set of decision-relevant indicators that show when an inflection is forming.
- Real options become the portfolio: Organisations invest in small, reversible bets that preserve flexibility while avoiding lock-in.
- Governance becomes a living system: Oversight evolves into a standing rhythm that keeps risk appetite, autonomy settings and controls aligned to a moving environment.

If recent history is any indication, fundamental inflections in AI are now an annual occurrence.

Long-range strategies risk almost instant obsolescence unless they feature iterative review cycles, real-world sensors, earlier signals, scenario thinking, and flexible decisions.

Australian context

Australian organisations often adopt and adapt global AI capabilities, which makes external developments in models, regulation and geopolitics feel immediate, even when they originate offshore. Boards tend to have a strong risk and compliance orientation, and AI is being treated as a standing governance topic.

National institutions and horizon-scanning capability, including public-sector safety initiatives and research-led foresight, provide useful signals for leaders who want to anticipate inflections early and avoid being surprised by changes in expectations, supply constraints, or platform dynamics.

Leadership priorities

Move to faster strategic cycles, building a standing rhythm to revisit AI strategy assumptions, risk appetite and capability bets as conditions shift. Enlist an AI agent to help you, with market and business sensor to alert you to shifting sands!

Run scenarios with triggers: Maintain a small set of decision-relevant futures with early signals that prompt action, not debate.

Strengthen board governance: Treat AI as a continuing governance agenda with clear oversight of autonomy, resilience and external dependencies.

Signs of transition

Cadence: Strategy is reviewed on a faster rhythm connected to real-world watchpoints and market signals, rather than an annual planning cycle.

Optionality: The organisation holds real options that can scale up or exit cleanly as technology, regulation and economics shift.

The AI leader's mindset

AI leaders ultimately become good at running the business of today while building the business of tomorrow.

Futures-minded leaders run strategy as a living system. They accept that AI evolves in jumps, so governance becomes about readiness, options and pace rather than certainty. Leaders build the habit of scanning for early signals, holding a small set of credible scenarios, and maintaining reversible bets that keep the organisation able to move quickly without thrashing. The mindset prizes adaptability.

For boards in 2026, AI governance is the mechanism that keeps strategy coherent as capability evolves faster than planning cycles, ensuring the organisation can scale autonomy and value without drifting into unmanaged risk, unbudgeted cost, or decisions that cannot be defended to customers, regulators, or the board itself.

Leading in the next phase

Leading in the next phase

The opportunity of Artificial Intelligence is not a technology question – it is a business question and ultimately a leadership challenge.

In the next phase of AI-native enterprises, the centre of gravity shifts to decisions only the top team can make: where AI will move the P&L in the next 12–24 months, how much capital and compute to commit, what boundaries set trust and safety, and which partners to rely on without losing control. Australian companies have strong foundations - sector depth, rich data and digitally engaged customers - but leadership sets the pace by focusing on system performance: faster cycles, better accuracy and clearer unit economics, with assurance built in.

The moment: AI is how the business runs

In 2026, AI is moving from the edge into the core. The next phase is defined by AI embedded in architecture, operations and decision-making, showing up in how products are designed and priced, how customers are acquired and served, and how risk is assessed and managed.

The upside is faster cycles, better accuracy and stronger unit economics; the challenge is uneven adoption, rising compute costs and trust that needs to be earned in production. Bold and confident action delivers the preferred future, relying on strategic decisions.

Decisions only the top can take

The factors shaping 2026 crystallise into a handful of choices that set direction and pace.

Leadership clarifies where AI will move the P&L, selecting a small number of customer and risk journeys where unit economics can shift within quarters, alongside one or two longer-horizon bets that open a second curve of growth.

Capital and compute become conscious commitments, with capacity secured and cost per outcome made visible in the financial plan.

Trust is treated as a design condition, with go-live criteria covering provenance, privacy, explainability proportional to impact, and fairness, plus clear escalation thresholds for safety or reputation concerns.

Partnerships are chosen for capability and control—what models and providers to rely on, how IP and data are protected, and how switching options are preserved as economics and performance change.

An operating rhythm ties it together: regular reviews of value and risk, clear stop/go decisions, and governance that moves at the speed of delivery.

Strengths to build

Execution depends on a set of capabilities that turn intent into repeatable results – all centred on ‘value’.

1. **Fluency** (understand value)

Fluency is a shared, plain understanding of how and where AI moves the economics of the business. Senior leaders set the narrative, align risk appetite, and demand evidence tied to revenue, margin and risk—not model metrics.

The aim is confident, disciplined decisions at pace: separating signal from hype and setting clear expectations for value and oversight.

2. **Strategy** (find value)

Strategy is choosing the few places where AI will change economics now, and the few bets that create a second curve of growth. Leadership links these choices directly to the P&L, allocates and sequences capital and compute, and sets payback expectations.

The stance is “decide and move”: just enough clarity to act, with the discipline to stop or scale based on evidence.

3. **Foundry** (prove and scale value)

Foundry strength is a leadership mandate for a repeatable path from proof to scale. Leaders require decision grade evidence of ROI and risk, standard patterns that can be reused across the enterprise, and a cadence of reviews that scale what works and stop what doesn’t.

The focus is turning wins into enterprise performance, fast and predictably.

4. **Trust** (safeguard value)

Trust is the leadership commitment that safeguards value and speeds approval. It is designed in from day one: clear purpose and risk ratings, decision rights and kill-switch criteria, provenance and privacy by default, explanations proportionate to impact, and fairness checks aligned to context.

Security and resilience are treated as part of the build—model supply chain integrity, defences against prompt injection and data poisoning, defined failure modes and fallbacks for agents and high-impact uses.

Assurance becomes continuous, with audit-ready trails and independent validation that keep pace with delivery. By standardising trust, leaders remove uncertainty, shorten approval cycles and make scaling predictable—so AI can be embedded confidently into how the business runs.

What distinguishes mature AI leadership from foolhardiness is an intentional investment in learning systems, clear ethical boundaries, and genuine humility about uncertainty – combined with conviction that moving forward is a matter of necessity.

Watchpoints

Leading an AI-native enterprise requires a steady hand.

Pitfalls are well known spreading effort across too many pilots, valuing benchmark scores over business outcomes, allowing shadow AI to grow unchecked, or letting governance lag delivery.

Cost discipline matters as compute and data usage scale; making those costs visible in unit economics improves decisions, while tracking energy and water intensity sets a course for efficiency.

Systemic risk deserves attention beyond single-model failure. Mapping dependencies, running scenario exercises and preparing for cascading impacts turns assurance into an always-on capability.

Foresight and optionality keep the organisation agile: maintaining a small portfolio of real options in emerging areas such as multimodal systems, agents and robotics, avoiding lock-in, and bringing a clear, current view to the board.

The management tone must be pragmatic and accountable – prove impact early, scale what works, pause what doesn't, and keep learning loops tight. This is how Australian businesses will turn strong foundations into leadership in the region.

Leading through uncertainty

Confidence in a fast-moving AI landscape comes from clarity of outcomes and a posture that favours informed action over perfect knowledge.

In the next period, AI is part of core infrastructure, and progress hinges on directional correctness: enough insight to move now, backed by evidence that is tied to earnings, risk and pace, and the willingness to adjust as signals emerge. The emphasis shifts from evaluating individual models to stewarding an integrated enterprise system—data, models, workflows and controls operating together—so performance, trust and speed reinforce rather than trade off.

Managing the whole system becomes the real strength. Standards for release are embedded and proportionate to impact; telemetry makes value and risk observable in real time; and approvals accelerate because evidence is produced by default, not after the fact. Optionality is preserved, avoiding lock-in while keeping momentum. Compute and unit economics are treated as first-order constraints within the system, with capacity and cost visible at the level of outcomes. Rehearsal in synthetic environments and attention to systemic dependencies reduce surprises and improve resilience before scale.

Australia's strengths—deep sector expertise, rich data and digitally engaged markets—are well suited to this stance. Confidence grows as an integrated system delivers repeatable results: features and services that show earnings impact, controls that shorten approval cycles, and learning loops that turn incident insight into better design.

The proactive principle applies: act on what is knowable, keep the system adaptable, and maintain room to pivot as the frontier shifts – so AI becomes how the business runs even as discovery continues.

What distinguishes this thinking from foolhardiness is the intentional investment in learning systems, clear ethical boundaries, and genuine humility about uncertainty – combined with conviction that carefully moving forward is a matter of responsibility and necessity.

AI is already shaping market winners across every industry.

Defensive AI is no longer an option, reducing costs, improving productivity, and reacting to change.

But market leadership belongs to those who invest in **bold moves** and accelerate to become ‘AI First’.

PwC’s integrated approach

From idea to enterprise value, with speed and confidence:

Fluency

Build the AI mindset and skillset across your executive team; **so you can** lead with confidence

Strategy

Shape your ambition, prioritise value linked to commercials, and define your enterprise blueprint; **so you can** scale effectively

Foundry

Prove value and deliver solutions at scale; **so you can** drive business outcomes

Trust

Embed Responsible AI, governance and transparency; **so you can** grow with integrity

With fluency, strategy, foundry and trust working as a force multiplier, organisations move beyond isolated pilots to deliver sustained enterprise-wide value — fast, repeatably, and with the confidence stakeholders expect.

AI quick guide

Breaking past the limits of human capacity

Today, AI operates at three levels of maturity. You can think of them as AI that helps individuals, AI that helps teams, and AI that helps the whole organisation run differently. Each level builds on the one before it, moving from personal productivity tools to full workflow automation.

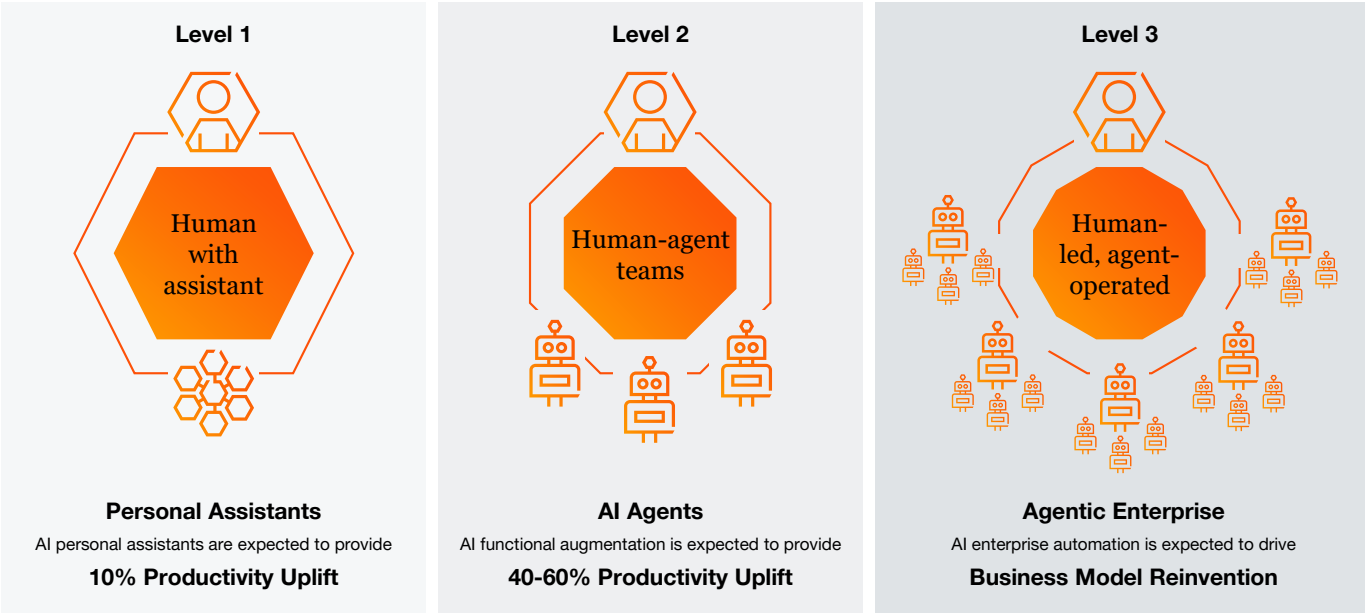

Level 1
Personal assistants

AI tools that help individuals with everyday tasks like drafting, summarising, analysing and searching. They boost personal productivity without changing how the broader business operates.

Example: A consultant uses an AI assistant to draft emails, summarise meetings and create first-cut slides, speeding up their daily work but not altering the team's overall workflow.

Level 2
Functional augmentation

AI is embedded into team processes, with agents completing tasks and humans providing oversight. This reshapes how a function works, reducing manual effort and improving consistency.

Example: An audit team uses an AI agent that prepares workpaper drafts, checks compliance steps and triages evidence, while staff focus on judgment, risk areas and client conversations.

Level 3
Enterprise automation

AI agents work across functions and systems to run end-to-end processes, transforming operating models, roles and cost structures. This is AI acting as a coordinated digital workforce.

Example: A full claims journey is automated: one agent gathers documents, another checks eligibility, another drafts decisions and customer communications, and humans step in only for complex cases.

2025 AI Technology in review: the state of the art, and signals for 2026

1. Benchmarks reset: "model quality" becomes a question of measurable work and reasoning.

In 2025, the industry shifted from debating whether models were "smart" to proving whether they could deliver reliable work. The frontier moved to tougher, operational benchmarks: human preference leaderboards, difficult reasoning tests, coding tasks, and knowledge evaluations like *Humanity's Last Exam*, a language model benchmark consisting of over 2,500 expert-level questions across a broad range of subjects. These better reflect business interests – can AI handle complex tasks, under constraints, with predictable results.

On the upside, one mainstream release hit 99.5% on AIME 2025 (mathematical reasoning) when allowed to use a simple Python tool, and another sparked developer buzz by fixing real GitHub issues in an agent setup. The same model managed 18.8% on Humanity's Last Exam – a number that looks small until you remember the test is designed to probe where even experts struggle.

We also crossed a cultural milestone. In April, a leading model (GPT-4.5) effectively **passed the Turing test**—people couldn't reliably tell they were chatting with an AI. The headline isn't perfection; it's that performance on practical tasks is now measurable, comparable, and steadily improving, especially when models are allowed to use the right tools.

Model offerings also organised into clear "families": fast general-purpose models for everyday tasks, explicit reasoning variants for harder analysis, agentic coding setups for software work, long context models that handle million token inputs, and multimodal generators that span text, images, audio, and video. A year ago, models mostly drafted and summarised, but by late 2025 they could ingest large inputs, run checks, and complete multistep workflows with trackable performance.

2. The marginal cost of intelligence stays real, from computation to carbon emission.

The biggest hidden story of 2025 was that AI started behaving economically like infrastructure, not software: training costs were widely discussed as being in the hundreds of millions of dollars for top-tier runs, and inference became a persistent operating expense rather than a one-off build cost.

For advanced workloads, estimates circulated that generating a million-token output on an 8-GPU H100 setup could translate to tens of dollars in direct compute cost at on-demand cloud pricing, before you count engineering overhead, storage, retries, and safety layers.

At the same time, MIT captured the energy and emissions angle in plain language: data centres already draw roughly 1–2% of global electricity, and the AI-driven trajectory pushes that materially higher; it also noted that even "small" units of AI consumption (like processing a million tokens) have non-trivial carbon implications when scaled.

The ecosystem response was predictable and telling: NVIDIA's Blackwell-era push, AMD's acceleration, and hyperscaler custom silicon all became part of the "state of the art" story, because capability is now coupled to input strategies across power, memory bandwidth, packaging, cooling, and optimisation.

The 2026 signal is that leaders can't treat compute as an elastic utility. Efficient engineering (token reduction, caching, quantisation, routing) is an important strategic lever that determines what you can afford to deploy.

3. Reasoning becomes a real behaviour.

A second turning point in 2025 was that "reasoning" stopped being an implicit quality and became an explicit mode you can invoke, often paired with tools that force grounding.

Vendors converged on the idea that the best model is not always the one that thinks the hardest; it's the one that can choose when to think, when to look something up, when to calculate, and when to produce a verifiable answer.

One notable proof point was expert evaluation framing—claims like "fewer major errors on difficult real-world tasks" mattered as much as academic benchmarks. Late-year platform moves reinforced the direction: architects started to route between fast and deep reasoning rather than a single monolithic model. The 2026 implication is governance-by-design for reasoning: decide where deep thinking is worth the latency and cost and bake verification into workflows so you pay for confidence, not eloquence.

4. Agents got real, then collided with control, permissions, and monitoring.

Agents matured in 2025, but in a very specific way: the technology became capable of planning and tool use, while product adoption lagged behind because enterprises need supervision and control more than autonomy.

The ecosystem's proof points were tellingly unsexy: standardised tool interfaces and connectors, IDE and browser integrations, and reliability work to stop models taking shortcuts. Anthropic, for example, publicly highlighted large reductions in "shortcut" behaviour on agentic tasks in its Claude 4 line.

Microsoft, Google, and others pushed agents into the surfaces where work actually happens, but the theme was consistent: the hard part is permissioning, audit trails, and recovery when tools fail.

The 2026 signal is that agent capability will advance mainly through "control planes" (identity, least-privilege tool access, observability, pause/inspect/resume), turning agents into deployable automation.

5. From "open source" to downloadable models that force portfolio discipline.

A quieter but consequential 2025 shift was that more organisations could run capable models without relying solely on a vendor.

The term Open Weights matters here. A model's learned parameters – weights – are published so others can run and adapt the model, even if the full training recipe is not public. DeepSeek's R1 moment, paired with a family of distilled, smaller variants, and Meta's Llama 4 push made downloadable models relevant not just for hobbyists but for serious experimentation and some production use, including multimodal options.

The landscape is shifting from finding a single best model to managing a disciplined portfolio, with downloadable open weight models reducing reliance on vendor APIs and letting organisations route tasks to models that fit. Competition is moving to the system layer, where value is shaped by owning the interface and control plane for agents, using open agent to agent protocols and unified cards for interoperability, and winning on context through permission aware, high-quality data and strong document parsing. Enterprise adoption is increasingly centred on secure, private deployments with clear return, run on an agentic runtime or operating system that provides telemetry, continuous monitoring, layered security, and identity and access controls for non-human agents.

6. Multimodal goes mainstream, and provenance becomes part of the product.

Multimodality, models that can interpret and generate across text, images, audio, and video, crossed a practical threshold in 2025.

Video generation moved from novelty to productised capability, with turning points like audio synchronised video features, dialogue and sound effects, and the simultaneous rise of provenance tooling. The "AI slop" problem – the overwhelming flood of low-quality, mass-produced digital content) is being treated as an engineering problem – not a communications problem.

The 2026 signal is that content authenticity becomes a default platform feature, label, trace, detect, because once synthetic media is easy to generate, trust must be engineered.

7. Regulation enters the stack: deployment grade AI requires embedded governance.

In 2025, regulation and guidance started influencing architecture choices, not just legal reviews.

The EU AI Act timeline and emerging Australian guidance effectively raised the definition of deployable, traceability, documentation that can survive scrutiny, monitoring for drift, and clearer accountability, especially for tool using systems and high-fidelity media generation. Logs, evaluation, provenance controls and safety monitoring moved from "nice to have" to "must ship". The 2026 implication is execution, governance must be a runtime infrastructure that enables speed, rather than a late-stage brake.

Technology trends in 2026

Taken together, 2025's turning points set up 2026 as the year the competitive unit becomes the AI system, not the standalone model, routed portfolios, fast versus deep, tool orchestration, supervised agents, and continuous evaluation that ties quality to cost. Expect strategy to concentrate on inference economics and efficiency, control plane maturity for agents, disciplined model portfolio management, including downloadable or open weight options, and provenance and monitoring built in by default.

A brief executive glossary

Agents / Agentic AI: AI systems that can take multi-step actions, interact with tools and data, and complete workflow tasks under human oversight.

Agentic Operating System (AOS): A standardised layer to run and govern swarms of agents for consistent safety, compliance, and resource control.

AGI (Artificial General Intelligence): A theoretical future capability whereby a single AI can understand, learn and perform any intellectual task a human can.

AI Risk: The operational, ethical, legal and reputational risks arising from AI use, including inaccuracy, bias, misuse, data leakage and loss of control.

AI Trust: Confidence that AI systems operate safely, accurately and consistently with regulatory and organisational expectations.

Bias: Systematic errors in AI outcomes caused by skewed data, unrepresentative training or flawed assumptions.

Chain-of-Thought Prompting: A prompting technique that asks AI to show its reasoning steps to improve clarity, accuracy and reliability.

Computation: The processing power required to train and run AI systems—often a constraint on scale, cost and speed.

Data Sovereignty: Keeping AI and data within trusted regions and providers you control.

Digital Twins: High fidelity virtual replicas of physical assets, processes, or systems used to simulate and predict outcomes.

Drift: When an AI model's performance declines over time because real-world conditions or data change.

Edge AI: Running efficient models on local devices (phones, laptops, NPUs) for lower latency.

Explainability: The ability to understand why an AI system produced a particular output or decision.

Foundation Model: A large AI model pre-trained on broad data that can be adapted ("fine-tuned") for many specific tasks. Foundation models include OpenAI's GPT, Google's Gemini, Meta's Llama, and Anthropic's Claude – as just some examples.

Generative AI: AI that can create new content such as text, code, images or summaries using learned patterns.

Generative Workflows: Re-designed processes where AI handles drafting, transformation and multi-step tasks, with humans focusing on judgment and exceptions.

GPT (Generative Pre-trained Transformer): A family of large language models built using the transformer architecture and trained to generate human-like text.

Guardrails: Rules, controls and checks that keep AI safe, compliant, accurate and on-brand.

Hallucination: When AI confidently produces information that is incorrect, fabricated or not grounded in real data.

Humans at the Helm: Governance principle ensuring people remain responsible for key decisions, with AI used to augment—not replace—human judgment.

Large Language Models (LLMs): AI models trained on massive text datasets that can understand, summarise and generate human-like language.

Machine Learning: A mode of AI that involves learning patterns from data to make predictions or decisions without being explicitly programmed.

Multimodality: AI's ability to process and generate more than one data type (text, image, audio, video).

Natural Language Processing (NLP): A technique for AI to understand and work with human language.

Neural Network: A model inspired by the brain's structure that learns complex patterns from data using interconnected layers ("neurons").

Observability: End-to-end visibility into how your AI systems behave and perform—what agents did, what data they touched, what outputs they produced.

Orchestration: Stitching together models, tools, and steps into one smooth process.

Patterns: Reusable blueprints for building AI workflows, agents, and controls (e.g., how to route tasks, apply guardrails, or monitor agents).

Pilot to Production: The shift from small experiments to deployed, scaled systems with business impact.

Prompting: The act of instructing an AI model—including how prompts are crafted to produce accurate, reliable outputs.

RAG (Retrieval-Augmented Generation): A method that grounds AI responses in trusted internal or external data sources to improve factual accuracy.

Reinforcement Learning: A technique where AI learns through feedback—rewarding correct actions and discouraging incorrect ones.

Responsible AI: Policies and practices to ensure AI is fair, ethical, safe and aligned with commitments.

Sustainable AI: Designing AI for energy efficiency and lower environmental impact.

Synthetic Data: AI generated datasets that mimic real data without exposing sensitive information, supporting creation of experimental environments.

Telemetry: Live data signals showing how AI is performing; value, usage, accuracy, cost and risk.

Training: Teaching an AI model to perform specific tasks, recognise patterns, and make decisions by exposing it to large amounts of high-quality data.

About this report

This publication presents original analysis and perspectives on how artificial intelligence will impact enterprises in 2026 and beyond. It draws on PwC's proprietary research and experience working with organisations, alongside publicly available industry developments and emerging AI practices.

Our insights are informed by evolving international and Australian AI standards, including ISO/IEC 42001, the NIST AI Risk Management Framework, EU AI Act developments, and Australian regulatory guidance. While we reference external standards and techniques for context, the analysis, interpretation and conclusions are entirely our own.

Methodology and forward-looking statements

This report synthesises global trends with specific attention to Australian market conditions, regulatory environment, and business context. Our forward-looking analysis is grounded in current developments but acknowledges that AI's rapid evolution may alter projected trajectories. Our analysis of AI's trajectory toward 2026 reflects current technological capabilities, emerging regulatory frameworks, and observable market trends. However, given AI's unprecedented rate of development—where the document notes 'fundamental inflections in AI are now an annual occurrence'—readers should treat these insights as directional guidance rather than definitive predictions.

Acknowledgements

We thank the PwC teams across Australia and globally who contributed insights, case studies, and expertise, as well as the clients and organisations whose experiences inform our understanding of AI's practical enterprise impact.

AI acknowledgement

We developed this report using generative AI tools responsibly and transparently. These tools helped with trend scanning, research synthesis and content structuring. All AI contributions were carefully reviewed and validated, with human expertise and judgment driving the process throughout. This approach reflects leading Responsible AI principles of transparency, human oversight, and accountability.

Important note

This publication is not intended as legal or regulatory advice. We recommend organisations consider their specific circumstances and requirements when applying these insights.

Sources

We acknowledge our intellectual debt to the researchers, institutions, and practitioners whose work has shaped the AI landscape we analyse here. Representative sources that have particularly informed our thinking include:

- PwC Global CEO Survey, and "29th Global CEO Survey – Australian insights"
- PwC Global AI Jobs Barometer
- PwC Digital Trust Insights Survey
- PwC Global Case Studies and Research

Standards & Frameworks

- ISO/IEC 42001
- NIST AI Risk Management Framework
- EU AI Act
- OECD AI Principles
- Australian Regulatory Sources:
- Australian Government AI Ethics Principles
- Australian Government Policy for Responsible AI in Government (December 2025 update)
- Office of the Australian Information Commissioner (OAIC) guidance on AI and privacy
- www.industry.gov.au/publications/national-ai-plan/keep-australians-safe
- ASIC expectations for explainable AI
- ACMA guidelines on AI-generated content

Technical and Industry Sources

- International Energy Agency (IEA) data centre electricity projections
- World Economic Forum data centre water consumption estimates
- MIT research on AI energy and emissions
- UNEP data centre sustainability guide
- Australian Federal Jobs and Skills Research
- IEA AI energy and resource consumption reports
- CSIRO Data61 AI Research
- Anthropic Constitutional AI Research
- Google DeepMind AI Safety Research
- Meta AI Research (Llama development)
- NVIDIA AI Enterprise Research
- Harvard Business Review AI Strategy Research
- World Economic Forum Future of Jobs Reports

AI Research, Working Groups and Benchmarks

- Cambridge University Centre for the Future of Intelligence: AI Futures and Responsibility Research
- MIT Centre for Information Systems Research (CISR)
- Stanford Human-centred Artificial Intelligence Index (HAI), Foundation Model Transparency Index (FMTI)
- Oxford Internet Institute AI Governance Research
- University of Melbourne Centre for AI & Digital Ethics
- Humanity's Last Exam benchmark
- AIME 2025 mathematical reasoning benchmark
- Linux Foundation Data and AI Research
- Academic research on AI safety and alignment
- World Economic Forum AI Governance Alliance
- UNESCO AI Education
- Australian Institute of Aboriginal and Torres Strait Islander Studies (AIATSIS) - Indigenous Data Sovereignty

Author

Matthew Tutty

Partner, Strategy,
PwC Australia

Board and executive advisory on AI strategy, AI ethics, leadership, business reinvention.

MAI Ethics (Cambridge), MLaws, MBA, FAICD.

PwC Australia AI expertise

Jahanzeb Azim

Partner,
AI Leader

Dr Gayan Benedict

Partner, Advisory,
MIT CISR Fellow

Simon Herrmann

Partner, Business
Reinvention Leader

Nicola Costello

Partner,
Digital and AI Trust

Amrita Jebamoney

Partner, Risk and Digital Trust Leader

Robert Di Pietro

Partner, Cybersecurity and Privacy

Emma Hardy

Partner,
AI Workforce

Amy Plowman

Partner, People & Organisation

Rob Kopel

Partner, AI Centre of Excellence

John Studley

Partner,
Data & AI Leader

Alfredo Martinez

Partner, AI regulatory compliance

Charmaine Chalmers

Partner,
CIO Advisory