

AI REPRESENTATION RISK IN AUSTRALIA

The state of AI business representation in Australia

What 320 AI representation audits reveal about how AI systems find, describe and compare Australian businesses, and the commercial cost of getting it wrong.

Inside this report

A synthesis of recurring patterns from 320 AI representation audits, situated within the wider Australian and global evidence base.

01 Executive summary

What the audit work and the wider Australian evidence base together reveal.

02 The Australian business landscape

Scale, structure and digital readiness of the businesses exposed to AI representation risk.

03 AI is now a buyer surface, not a niche channel

How fast Australian buyers have moved to AI-mediated discovery.

04 Australian businesses are not ready for this surface

Adoption is wide, governance is shallow, the gap is structural.

05 The AI Representation Risk Framework™

The seven layers of how AI representation can fail.

06 Findings by risk layer

What the audits revealed across each of the seven layers.

07 Cross-pattern observations

Where the layers compound and what that tells us.

08 Commercial implications

What inaccurate AI representation costs in lost trust, lost revenue and brand damage.

09 The policy gap

Why representation governance now sits with the business, not the regulator.

10 What this means for your business

Strategic priorities and what to look for first.

11 About this report

Methodology, sample, limitations and contact.

12 References

Sources cited throughout the report.

Want to check your own business first?

Run the free AI Representation Snapshot for an early view of how AI describes you.

RUN THE FREE SNAPSHOT

Executive summary

Across 320 AI Representation Audits conducted by Iconic Marketing during the build and testing phase of its AI representation methodology, recurring patterns emerged in how AI systems describe, recommend and contextualise Australian businesses. Each audit tested AI platform outputs across repeated runs, with three runs per AI platform per audit. This report synthesises those observations through the AI Representation Risk Framework™ and situates them within the wider Australian and global evidence base on AI adoption, AI search behaviour, and AI output quality.

One conclusion runs across every layer of the framework. AI visibility is not simply a traffic problem. It is an accuracy problem, a trust problem, and a buyer interpretation problem. Australian businesses can lose consideration before any enquiry is made, because AI systems may omit them, describe them incorrectly, compare them unfavourably, or present them as less credible than they are.

WHAT AI BUSINESS MISREPRESENTATION MEANS

AI business misrepresentation occurs when an AI system inaccurately, incompletely, or weakly describes a business because the public signals available to that system are fragmented, outdated, unclear, or inconsistent. It is distinct from deliberate misinformation: no one needs to be intentionally spreading false claims for the business to be misrepresented. The problem is structural. It emerges from the gap between what a business knows about itself and what AI systems can find, interpret and synthesise from public sources.

AI systems do not read a business's website the way a human buyer does. They extract signals: entity information, service descriptions, location details, structured data, pricing cues and third-party references from directories, social profiles and external publications. When those signals are strong and consistent, AI can describe the business with reasonable accuracy. When they are weak or contradictory, AI fills the gap with inference.

The strategic stakes are no longer hypothetical

By 2025, 56 per cent of Australians were using AI tools, up from 45 per cent the year before, with 30 per cent already using AI assistants to answer search queries⁵. By mid-2025, two-thirds of Australians were turning to AI or wanted to, and almost four in ten were using AI assistants instead of traditional search engines⁶. Forty-eight per cent of Australians were using AI for product searches, rising to 66 per cent of those under 45⁷.

At the same time, an EBU/BBC study covered by Reuters found that 81 per cent of AI assistant responses contained some type of problem and 45 per cent contained at least one significant issue⁸. Independent reporting by The Guardian found AI search summaries producing dangerously inaccurate health information¹⁵, and the Reserve Bank of Australia has warned firms must verify AI outputs because hallucinations and injection attacks are now business risks⁹.

These two trends, rising buyer reliance on AI and a documented quality problem in AI outputs, converge into a single business risk. Australian businesses are increasingly being represented to buyers by systems that are demonstrably unreliable, and most of those businesses do not know what is being said about them.

The seven recurring risk layers

The patterns observed across the audit work fall into seven distinct risk layers, each describing a different way AI representation can fail. The layers are listed below in canonical framework order, from most heavily weighted to least heavily weighted.

-
- L1 Discovery Suppression.** AI cannot find or surface the business in response to relevant buyer queries.
 - L2 Factual Hallucination.** AI invents services, products, pricing or locations to fill signal gaps.
 - L3 Competitor Drift.** AI introduces or recommends competitors during descriptions or comparisons.
 - L4 Interpretation Drift.** AI categorises and describes the business differently than intended.
 - L5 Entity Confusion.** AI mistakes the business for a different organisation.
-

L6 Authority Gap. AI describes the business with hedging or qualified language.

L7 Source Decay. AI relies on outdated sources, describing a previous version of the business.

These layers rarely appear in isolation. The audit sample showed most businesses affected by three or more layers simultaneously, with the layers compounding rather than offsetting each other. A business with weak authority signals is also more likely to be hallucinated about, more likely to be supplanted by competitors, and more likely to be described from outdated sources.

How to read this report

The audit work is proprietary and is presented here as a directional evidence base, not a statistically representative national sample. The 56 audited businesses are not a randomised cross-section of the Australian market, and pattern frequency is reported qualitatively rather than as percentages.

The audits are interpreted alongside published Australian and global research. Where pattern observations align with independent evidence, that evidence is cited inline. Where the audit work points to behaviour not yet captured in published research, that is identified as Iconic Marketing's own observation.

THE COMMERCIAL IMPLICATION

Buyers are increasingly forming impressions of businesses through AI systems before any human contact occurs. When AI representation is inaccurate, incomplete or hedged, the business inherits a disadvantage it did not earn and may not know exists. AI representation is now a strategic surface that needs to be governed, not just monitored.

To assess how AI may be representing your business, start with the [free AI Representation Snapshot](#).

The Australian business landscape

Before examining how AI systems are misrepresenting Australian businesses, it is worth establishing the scale of the population at risk and the digital baseline from which most of those businesses operate. Both are larger and more uneven than commonly assumed.

Scale of the population at risk

The Australian Bureau of Statistics reports 2,729,648 actively trading businesses in Australia at 30 June 2025, with the number of businesses increasing 2.5 per cent in 2024–25 and a net increase of 66,650 businesses over the year¹. The Australian Small Business and Family Enterprise Ombudsman reports the same population of GST-registered entities and notes that, using the ATO definition of less than \$10 million turnover, 98 per cent of Australian businesses are small and 92 per cent have turnover below \$2 million².

2.7M+

Actively trading businesses in Australia at 30 June 2025. 98 per cent are small businesses; 92 per cent have turnover below \$2 million.

This is the population for which AI representation now matters. Each of these businesses is potentially the subject of AI-mediated buyer research: AI assistants asked to recommend providers, summarise services, describe locations, or compare options. Each is exposed to the seven risk layers in the AI Representation Risk Framework™, regardless of whether the business itself is using AI internally.

The digital readiness gap

Digital readiness across this population is highly uneven. The ABS Digital Intensity Index shows that digital sophistication increases with firm size: micro businesses with zero to four persons had only 5 per cent established digital intensity, while businesses with 200 or more employees had 53 per cent established digital intensity³. The result is a structural gap between the digital maturity of large firms and the digital maturity of the small businesses that make up the overwhelming majority of the Australian economy.

The 2025 Australian Digital Inclusion Index found that 45.6 per cent of Australians had recently used generative AI tools, while around 20 per cent of Australians remained digitally excluded, and warned of an emerging 'AI divide' across the population⁴. The risk extends beyond consumer access. Businesses with weaker digital foundations are also the businesses with weaker public signals on the surfaces AI systems use to describe them.

Why this matters for AI representation

Where digital signal density is low, AI systems do not stay silent. They infer. They fill gaps with the closest match. They borrow attributes from similar entities. They lean on outdated or secondary sources. The structural digital gap in the Australian business population is therefore not just a question of who can use AI internally. It is a question of which businesses are most exposed to being misrepresented by AI externally.

THE STRUCTURAL READING

The 2.7 million Australian businesses sit on a digital foundation that AI systems now read, interpret and present to buyers. Businesses with thinner foundations are not absent from AI outputs. They are present, but represented from inference rather than evidence. That is a different kind of risk.

AI is now a buyer surface, not a niche channel

The most consequential shift in the last twenty-four months is not how Australian businesses use AI internally. It is how Australian buyers use AI to research the businesses they might purchase from. AI has moved from a curiosity to a primary discovery surface across a meaningful share of the population, with adoption rising fastest among the demographics businesses most want to reach.

Adoption is now mainstream

auDA's Digital Lives of Australians 2025 report found that 56 per cent of Australians use AI, up from 45 per cent the year before. Thirty per cent already use AI assistants to answer search queries, and 18 per cent use AI assistants to solve problems⁵.

Adobe's mid-2025 study found that 66 per cent of Australians were turning to AI or wanted to, up from 59 per cent only three months earlier. One in three Australians regularly use AI assistants, and almost four in ten use AI assistants instead of traditional search engines. Seventy-five per cent of AI assistant users find them helpful⁶.

On product research specifically, 48 per cent of Australians now use AI for online product searches, rising to 66 per cent among under-45s. Seventy-eight per cent expect AI shopping tools to become mainstream. Sixty-one per cent would trust AI for product recommendations, even though 59 per cent have received unexpected or laughable AI search results⁷.

4 in 10

Australians now use AI assistants instead of traditional search engines (Adobe, July 2025).

The behavioural shift this creates

The implication for businesses is structural, not incremental. The buyer no longer arrives at a business's website having shortlisted it through a Google search. The buyer increasingly arrives, or fails to arrive, having already had a conversation with an AI system. By the time the buyer is on the website, the description has been formed, the alternatives have been suggested, the comparison has been made, and the trust position has been set.

Three behaviours follow from the data above.

1. **Buyers ask AI systems to recommend, summarise and compare.** Where the audited business is omitted, drift-affected, or hedged, this happens before any contact.
2. **Buyers trust AI recommendations even when they have seen AI produce errors.** The willingness to act on AI output exceeds the confidence in AI accuracy.
3. **Younger buyers are pulling adoption forward.** Two-thirds of under-45s already use AI for product research. The trajectory is one-directional.

The invisible buyer stage

TRADITIONAL SEARCH · VISIBLE TO ANALYTICS

User search

›

Website results

›

Website visit

›

Evaluation

AI-ASSISTED · FIRST THREE STEPS INVISIBLE TO ANALYTICS

User question

›

AI explanation

›

Shortlist formation

›

Website visit

›

Enquiry

Strategic implication: a business can lose consideration before there is any search impression, website session, form submission or phone enquiry to analyse.

AI output quality is documented as unreliable



81 per cent of AI assistant responses contained some type of problem and 45 per cent contained at least one significant issue.

European Broadcasting Union and BBC, reported by Reuters, October 2025

This rising buyer reliance sits alongside a documented quality problem. An EBU/BBC study covered by Reuters in October 2025 found that 45 per cent of AI assistant responses contained at least one significant issue, and 81 per cent contained some type of problem. A third of responses had serious sourcing errors, and 72 per cent of Gemini responses had sourcing issues⁸.

The Guardian, in early 2026, reported on Google's AI Overviews providing dangerous health information, with experts warning that AI-generated search summaries may produce different answers on different days¹⁵. Australia's own ACSC and ASD small business guidance explicitly warns that AI systems can make up information or hallucinate, and that prompt-injection attacks can manipulate AI outputs¹⁷. The Reserve Bank of Australia's November 2025 bulletin reinforces that firms must verify AI outputs because hallucinations and injection attacks are real risks⁹.

THE CONVERGENCE

Australian buyers are increasingly using AI to discover, describe and shortlist businesses, on a surface that independent research shows is unreliable in the majority of responses. Whether or not a business chooses to engage with AI internally, its representation on this surface is being shaped externally, every day.

Australian businesses are not ready for this surface

If Australian buyers have moved decisively to AI-mediated research, the supply side has not kept pace. Adoption of AI inside Australian businesses is rising sharply, but adoption is shallow, governance is thin, and the smaller businesses that make up the bulk of the economy lag the most. The combination produces a representation gap that is widening, not closing.

Adoption is wide but shallow

The Reserve Bank of Australia's November 2025 bulletin found that roughly two-thirds of Australian firms had adopted AI in some form, but adoption was shallow. Almost 40 per cent used AI minimally, and only around 30 per cent used AI substantively for tasks like forecasting and inventory management. Larger firms adopted AI more readily, while smaller firms cited barriers including lacking digital readiness and uncertain ROI⁹.

Indeed Hiring Lab, drawing on the Federal AI Adoption Tracker, found a clear size gradient: 78 per cent of large businesses (200 to 500 employees) report some level of AI adoption, falling to 72 per cent for medium-sized businesses, 60 per cent for small businesses, and 36 per cent for micro businesses. Broad usage remains rare, at 8 per cent in services and 1 per cent in mining¹⁰.

The Governance Institute of Australia's August 2025 survey found nearly nine in ten organisations using AI. Yet almost half of employees had received no AI training, 93 per cent could not effectively measure ROI, and 88 per cent struggled to integrate AI with existing systems¹¹.

36%

Of Australian micro businesses report any AI adoption, against 78% of large businesses (Indeed Hiring Lab, 2026).

Awareness, accuracy and data risk are real barriers

The University of Technology Sydney's Human Technology Institute, in a January 2025 study of 133 SMEs, found that generative AI often exceeds expectations but adoption is hindered by low AI awareness, concerns about accuracy and quality of AI outputs, and data-risk concerns. Nearly half of respondents worried about AI accuracy and emphasised the need for human oversight and training¹².

BizCover's 2025 Small Business AI Report, drawing on a survey of 965 small business owners, found 66 per cent already using AI and 14 per cent planning to adopt soon. Forty-eight per cent considered AI important to daily operations, but 49 per cent worried that AI could harm creativity¹³.

The early baseline

It is worth recalling how recent this shift is. The ABS's Characteristics of Australian Business 2021–22 release found just 1 per cent of businesses reported using AI, while 59 per cent used cloud computing and 41 per cent used digital platforms to reach customers¹⁴. The trajectory between 2022 and 2025 represents one of the steepest adoption curves in recent business technology history, but adoption depth and governance maturity have not moved at the same pace.

The representation gap that follows

These three observations, shallow adoption, weak governance, and a sharp size gradient, matter for AI representation in two ways.

First, businesses that are not yet using AI internally are also less likely to be testing how they appear in AI outputs externally. The blind spot compounds. The business is unaware of what AI is saying about it because it is unaware that the question is one worth asking.

Second, the businesses most exposed to representation risk, smaller, less digitally mature, with less internal AI capability, are precisely the businesses least equipped to identify and remediate that risk. The size gradient in adoption maps closely onto the size gradient in representation vulnerability.

THE ASYMMETRY

On the demand side, AI is mainstream and growing. On the supply side, AI is shallow and uneven. The businesses most exposed to misrepresentation are the businesses least likely to know it is happening. That asymmetry will not close on its own.

The AI Representation Risk Framework™

The AI Representation Risk Framework™ organises the failure modes observed across the audit work into seven distinct layers. Each layer describes a different way an AI system can represent a business inaccurately, incompletely, or with insufficient authority. Each is weighted according to how heavily it impacts buyer interpretation and commercial outcomes.



The seven layers are listed in order of importance. Discovery Suppression carries the heaviest weight because a business that AI cannot reliably find does not exist in AI-mediated conversations, regardless of how strong the rest of its digital presence is. Source Decay carries the lightest weight on its own, but materially compounds the impact of every other layer when present alongside them.

The seven layers are not a maturity model. They are independent risk categories that frequently appear together. The most common pattern in the audit sample was simultaneous exposure across three or more layers, with each weakness reinforcing the others.

The framework treats AI representation as a structural surface, not a channel performance question. It distinguishes between absence, inaccuracy, comparison, interpretation, identity, confidence and recency. Each is a different failure mode that calls for a different response.

The seven layers, in summary

RISK LAYER	FRAMEWORK WEIGHT	WHAT IT COSTS THE BUSINESS
L1 Discovery Suppression	25% Highest weight	AI cannot find or surface the business in response to relevant buyer queries. The business does not exist in AI-mediated conversations.
L2 Factual Hallucination	20%	AI fills signal gaps with inferred or fabricated information: wrong locations, services not offered, invented credentials. Buyer expectations are set against reality.

RISK LAYER	FRAMEWORK WEIGHT	WHAT IT COSTS THE BUSINESS
L3 Competitor Drift	20% 	AI compensates for weak signal by drawing on competitor signal. Competitors are introduced into descriptions, or recommended where the business should be the answer.
L4 Interpretation Drift	15% 	AI categorises and describes the business differently than intended. Strategic consultancy reads as digital agency. Premium reads as general.
L5 Entity Confusion	10% 	AI mistakes the business for a different organisation. Recognises the name, but applies someone else's details.
L6 Authority Gap	7% 	AI describes the business with hedging or qualified language. Even when accurate, low-confidence descriptions undermine credibility before contact.
L7 Source Decay	3% Lowest weight 	AI describes a past version of the business. New services, locations or positioning are invisible because peripheral sources have not caught up.

The remainder of this report examines each layer in detail. Each layer has its own definition, its own representative scenario, its own observed pattern in the audit sample, its own underlying signal cause, and its own consequence for buyer interpretation. Where independent evidence supports the observed pattern, that evidence is cited.

Read the full [AI Representation Risk Framework™](#) methodology.

Findings by risk layer

Each of the seven layers is examined below in the canonical framework order, from most heavily weighted to least heavily weighted. Each layer combines the framework definition, a representative scenario, what was observed across 320 AI Representation Audits conducted across 56 businesses, the underlying signal causes, and the commercial consequence for the business when the layer is active.

L1

RISK LAYER · 25% · HIGHEST WEIGHT

Discovery Suppression

Discovery Suppression measures whether AI systems can find, identify and surface the business in response to relevant buyer queries. It is the most heavily weighted layer in the Framework. A business that AI cannot reliably discover does not exist in AI-mediated conversations, regardless of how strong the rest of its digital presence is.

WHAT THIS RISK LOOKS LIKE

A buyer asks an AI assistant to recommend a lawyer in Penrith. Your business does not appear in the response. Not because you do not qualify, but because your digital signals are not strong enough for AI to confidently include you.

WHAT WAS OBSERVED IN THE AUDIT SAMPLE

Discovery Suppression was the most frequently observed risk layer across the audit work. It was a recurring pattern observed in some form: either complete omission from category-level recommendation queries, or omission from queries with specific qualifiers such as 'best,' 'top,' 'trusted,' or 'specialist' that buyers commonly use.

Suppression rarely correlated with poor business quality or limited operating history. It correlated more reliably with thin web presence, sparse external citation, weak entity definition, or absence from the kinds of source pages AI systems treat as authoritative for a given category.

WHAT DRIVES IT

- Limited primary-source content that defines the business clearly in its category.
- Weak or inconsistent entity signals across the public web.
- Low representation in structured sources AI systems weight: industry directories, association listings, well-structured services pages, credible third-party references.
- A website that does not state the business's category, services and locations in extractable form.

THE BUSINESS CONSEQUENCE

The loss of consideration before any contact is made. The buyer has already shortlisted competitors before the audited business has any opportunity to present itself. This is the most strategically dangerous of the seven layers because it is the most invisible: there is no enquiry to track, no rejection to learn from, no signal that the buyer ever considered the business at all.

Buyer-side validation. Adobe's 2025 research found that almost four in ten Australians now use AI assistants instead of traditional search engines, and the figure was rising month over month. For any business absent from the AI shortlist, this is the cohort it is now missing entirely, and the share of that cohort is one-directional.

L2

RISK LAYER - 20%

Factual Hallucination

AI systems generate content based on available signals. When those signals are incomplete, outdated or contradictory, AI fills the gaps with inferred or fabricated information. Factual Hallucination measures the risk that AI systems are producing inaccurate descriptions of the business: wrong locations, incorrect services, outdated pricing or invented credentials.

WHAT THIS RISK LOOKS LIKE

An AI assistant tells a potential buyer your business is based in a different city, specialises in a service you do not offer, or has been operating for far fewer years than you have. The buyer moves on before ever contacting you.

WHAT WAS OBSERVED IN THE AUDIT SAMPLE

Hallucination was less frequent than Discovery Suppression, but it appeared in a meaningful share of audits and was particularly serious where it occurred. Hallucinated content tended to cluster in two patterns: filling gaps where the business's authoritative content did not state key facts plainly, and importing facts from a similarly named or similarly positioned business.

Hallucination was disproportionately observed in audits where the business had recently changed its services, locations or trading name, but had not updated peripheral sources to match.

The pattern is consistent with independent research. The EBU/BBC study covered by Reuters found 81 per cent of AI assistant responses contained some type of problem and 45 per cent contained at least one significant issue, with a third showing serious sourcing errors⁸. The Reserve Bank of Australia's November 2025 bulletin warns that hallucinations and injection attacks are real business risks that require human verification of AI outputs⁹. The ACSC's 2025 guidance to small businesses explicitly states that AI systems can make up information¹⁷.

WHAT DRIVES IT

- Sparse primary-source content, leaving AI to infer rather than extract.
- Absence of clear, structured statements of key facts on the website.
- Conflicting information across web properties: the website says one thing, a directory says another, a press release says a third.
- Similarity to other businesses that have stronger or more confident signals.
- Older content that asserts facts no longer true.

THE BUSINESS CONSEQUENCE

Customer expectations are set incorrectly before the buyer makes contact. The business inherits friction it did not earn: enquiries about services that do not exist, callers asking for staff who do not work there, leads filtering on pricing or capability the business never advertised, and complaints rooted in expectations it never set. Hallucinated representation is structurally harder to remediate than the other layers, because the inaccurate content is often anchored to outdated or secondary sources that are slow to update.

L3

RISK LAYER · 20%

Competitor Drift

When AI systems lack sufficient signal from the business, they compensate by drawing on the signals of businesses in the same category, including competitors. Competitor Drift measures the risk that AI systems are conflating the business with competitors, introducing competitor names into descriptions of its services, or recommending alternatives when the audited business should be the answer.

WHAT THIS RISK LOOKS LIKE

A buyer asks AI which marketing agency they should contact in the Hunter region. Your business is mentioned briefly, but a competitor is positioned more prominently and described in more specific, credible detail, because their signals are stronger than yours.

WHAT WAS OBSERVED IN THE AUDIT SAMPLE

Competitor Drift was a common pattern in the audit sample, and a particularly counterintuitive one for senior leaders. The audited businesses were frequently larger, more credentialed or longer-established than the competitors AI introduced or recommended ahead of them. The competitors that outperformed in AI representation were rarely outperforming in market quality. They were outperforming in signal legibility.

The competitive set produced by AI was, in several audits, materially different from the competitive set the business itself tracked. AI-mediated competitors were often businesses the audited brand did not consider strategic peers, but which AI systems had structurally easier evidence to confidently describe and recommend.

WHAT DRIVES IT

- Tighter category claims by competitors. A competitor that says 'specialist in X for Y customers in Z region' is more legible than one that lists ten services across three states.
- More consistent entity description across the public web.
- Better-structured services and About content directly addressing buyer-language questions.
- More extractable evidence of credentials, methodology or proof points.

THE BUSINESS CONSEQUENCE

Market share erosion that is not visible in traditional competitor analysis. The audited business may track competitors using channels the buyer no longer uses to discover services. By the time the loss is visible in pipeline data, the buyer-perception gap has been compounding for months. The buyer's competitive set is shaped by AI legibility, not by AI judgement of quality.

Buyer-side validation. PayPal's research found 61 per cent of Australians would trust AI for product recommendations, even though 59 per cent had received unexpected or laughable AI search results. The willingness to act on AI's competitive recommendations exceeds buyer confidence in AI accuracy.

L4

RISK LAYER · 15%

Interpretation Drift

AI systems do not always interpret the business the way it intends to be understood. Interpretation Drift measures the gap between how the business positions itself and how AI systems categorise and describe it. A business that defines itself as a strategic consultancy may be interpreted by AI as a general marketing agency, which affects which queries surface it, which buyers it reaches, and how it is compared against competitors.

WHAT THIS RISK LOOKS LIKE

You have spent years building a reputation as a premium residential builder. AI describes you as a general contractor. Buyers looking for quality are passed to competitors. The ones who find you arrive with the wrong expectations and the wrong budget.

WHAT WAS OBSERVED IN THE AUDIT SAMPLE

Interpretation Drift appeared frequently across the audit work, particularly among businesses with mixed positioning, broad service mixes, or website copy written in generic industry language. It was the most commonly observed risk layer for businesses that had grown organically across multiple service lines without a formal repositioning exercise.

In several audits, AI systems converged on a generic description that resembled the business's least distinctive sales page rather than its strongest differentiating content. Where a business led with a clear proposition on its homepage but allowed less distinctive pages to dominate the site's structural signals, AI systems frequently reproduced the less distinctive characterisation.

WHAT DRIVES IT

- Vague website copy, particularly on services and About pages.
- Service descriptions written in generic industry vocabulary rather than the language of buyer outcomes.
- Multiple positioning claims across the site without a clear primary one.
- Heavy reliance on category nouns such as 'strategy,' 'marketing,' 'consulting' without specific, extractable definitions of what those words mean for this business.
- Third-party references that describe the business in generic terms because the business itself does not provide more specific language.

THE BUSINESS CONSEQUENCE

Drift toward generic descriptions makes the business look like everyone else. The investment made in differentiation, methodology or sector expertise is invisible to AI-mediated buyers. The business becomes harder to choose, not because it has been compared and found wanting, but because there is nothing distinctive in the description for the buyer to attach to. Drift also affects which buyers reach the business at all: those looking for the actual proposition do not find it, and those who do find it arrive with mismatched expectations.

L5

RISK LAYER · 10%

Entity Confusion

Entity Confusion occurs when AI systems mistake the business for a different business entirely. This happens when another organisation shares a similar name, operates in the same category, or has overlapping public identifiers that cause AI to blend the two together.

WHAT THIS RISK LOOKS LIKE

A buyer asks an AI tool about your business and receives information that partially or fully describes a different company. Your name is recognised, but the location, services, ownership or description that follows belongs to someone else. The buyer forms an impression of your business based on another organisation's profile.

WHAT WAS OBSERVED IN THE AUDIT SAMPLE

Entity Confusion was particularly pronounced where business names were generic, where multiple entities used similar trading names, where the business had changed name or merged, or where directory listings were inconsistent. It was also common where founder or director profiles were stronger than the business profile, leading AI to describe the business through the founder lens or vice versa.

The risk has a public dimension in Australia. The ACCC reported in July 2025 that it had received at least 360 reports about 60 online retailers misrepresenting themselves as local Australian businesses, with these so-called 'ghost stores' often using AI-generated images of owners or teams to create credibility¹⁸. Entity Confusion in AI outputs is one route by which legitimate businesses can be tangled with bad actors operating under similar trading names.

WHAT DRIVES IT

- Generic or commonly used business names.
- Multiple entities trading under similar names in the same category or geography.
- Directory listings that haven't been updated, especially after rebrands or location changes.
- Founder profiles that are stronger than the business profile, creating ambiguity about which entity is the subject of a query.
- Inconsistent legal name, trading name and brand name use across the public web.

THE BUSINESS CONSEQUENCE

Confusion at the point of buyer search. Reputation belonging to the business is not credited to it. Reputation belonging to another business is wrongly attributed. In the worst cases, association with bad actors operating under similar names can damage the legitimate business directly. Buyers form impressions that are partially accurate and partially imported from another entity, and confidence in the description erodes when buyers verify and find inconsistencies.

L6

RISK LAYER · 7%

Authority Gap

Authority Gap occurs when AI systems cannot describe the business with confidence, producing hedging language, vague responses or qualified statements that signal uncertainty to buyers. Even where the business is visible and technically accurate in AI results, a low-confidence description undermines credibility before a buyer has made any direct contact.

WHAT THIS RISK LOOKS LIKE

A buyer asks an AI assistant about your business. Instead of a clear, specific description, the response includes phrases like 'may offer' or 'is believed to provide' or 'limited information is available.' The buyer reads uncertainty into that language and moves on to a competitor described with more conviction.

WHAT WAS OBSERVED IN THE AUDIT SAMPLE

Authority Gap was commonly observed where the business had limited third-party citations, sparse structured data, weak About / Services / Methodology content, no peer-reviewed or industry-recognised credentials surfaced in extractable form, or where the business's claims lived primarily in marketing copy rather than substantive documented evidence.

Hedging language was particularly damaging in head-to-head comparison queries. Where competitors were described declaratively and the audited business was described with qualifiers, the resulting comparison read as a recommendation against the audited business, even where the AI made no such explicit recommendation.

WHAT DRIVES IT

- Limited third-party citations of the business in credible sources.
- Sparse or inconsistent structured data on the website.
- Marketing claims that are not corroborated by extractable evidence.
- Absence of methodology, evidence, case study or proof content in formats AI systems can readily ingest.
- Weak founder or organisational authority signals across professional surfaces.

THE BUSINESS CONSEQUENCE

Even when included in AI responses, the business is presented as a less confident option. Hedging language reduces consideration weight. The hedge is read as risk. Where the business has made meaningful investment in capability or credentials but those investments are not surfaced in AI-extractable form, that investment is effectively invisible at the moment of buyer evaluation.

L7

RISK LAYER · 3% · LOWEST WEIGHT

Source Decay

Source Decay measures the risk that the information sources AI systems draw on to describe the business are outdated, stale or no longer accurate, producing descriptions that reflect a past version of the business rather than its current reality. It is the lowest-weighted layer in the Framework on its own. However, when present alongside other layers, Source Decay materially compounds their impact.

WHAT THIS RISK LOOKS LIKE

Your business has evolved significantly over the past three years. New services, new locations, new positioning. But the sources AI draws on most heavily still reflect how your business looked in 2021. AI describes an old version of you to a new buyer. Or worse: AI references a founder who left 18 months ago.

WHAT WAS OBSERVED IN THE AUDIT SAMPLE

Source Decay was most common where the business had rebranded or repositioned without updating peripheral sources, where old PR or directory profiles still ranked well in third-party indexes, or where current website content was not structured in a way AI systems could confidently extract from.

Several audits found AI systems confidently citing content the business had taken down or replaced, or describing the business using language from earlier marketing campaigns the business had since moved away from. In multiple cases, Source Decay was the underlying cause of what initially presented as Factual Hallucination: the 'invented' facts were in fact accurate facts about a previous version of the business.

WHAT DRIVES IT

- Old PR coverage that continues to rank in third-party indexes after the business has changed direction.
- Directory profiles that have not been updated after rebrands, location changes or service changes.
- Older versions of the business's own content cached or quoted elsewhere.
- New positioning that has not yet earned the third-party citation depth of the older positioning.
- Website content that is less structurally extractable than older sources, even where the older sources are out of date.

THE BUSINESS CONSEQUENCE

The business's recent investment in repositioning, new services or new market focus is invisible to AI-mediated buyers. The business is being represented as a previous version of itself. In some audits, this gap was years wide. Where the buyer makes contact and the current reality differs from the AI-formed impression, the burden of explanation falls on the business.

These risk layers are assessed in the [AI Representation Audit](#), which tests how AI systems surface, describe, compare and interpret a business during buyer research.

Cross-pattern observations

The seven layers of the AI Representation Risk Framework™ are independent risk categories, but they rarely appear independently. The audit work surfaced six recurring cross-pattern observations that matter strategically. Each speaks to how the layers compound, and where the leverage points sit for businesses considering remediation.

01 Layers compound rather than offset

Many businesses showed several layers active simultaneously. The layers reinforced each other rather than offsetting each other. A business with thin authority signals was also more likely to be hallucinated about, more likely to be supplanted by competitors, and more likely to be described from outdated sources. Strength on one layer did not protect against weakness on another.

02 Authority Gap and Source Decay drive Discovery Suppression

When current authoritative content is thin and old content dominates, AI systems are more likely to hedge or omit entirely. The combination of weak authority signals on current content with strong peripheral signals on outdated content was a recurring driver of complete suppression in category-level recommendation queries. Discovery Suppression often presents as the symptom; the underlying issue is structural authority and recency.

03 Entity Confusion is frequently the root cause of Factual Hallucination

When AI cannot cleanly separate the business from a similar entity, it imports facts from the wrong source. Many of the hallucinations observed in the audit work were not invented from nothing. They were facts about another business that AI had attached to the audited one. Resolving entity definition often resolved the hallucination at the same time.

04 Competitor Drift is rarely about competitor quality

It is almost always about competitor legibility. Smaller, less established competitors with tighter category positioning frequently outperformed larger, more credentialed audited businesses in AI representation. Authority and quality of work are not what AI systems weigh first. They weigh signal legibility. Businesses confident in their market quality were sometimes the most exposed to Competitor Drift, because confidence had not been translated into structurally extractable evidence.

05 The website is the primary truth source AI systems calibrate against

Where website signals are strong, peripheral inaccuracies have less weight. Where website signals are weak, peripheral sources dominate. This makes the website the highest-leverage surface for AI representation work. It is also the surface the business has the most direct control over. Many of the patterns observed across the seven layers traced back to a website that was not stating the business's primary facts in extractable form.

06 Owners and operators are usually unaware

Direct queries to AI systems by the business owner often produced different results than queries from a buyer-perspective starting point. Owners typing the business name as a query received responses calibrated to a known entity. Buyers asking AI for recommendations in a category received responses calibrated to whoever AI considered the most legible providers. The two sets of results differ structurally, and the divergence is one of the main reasons representation issues persist undetected.

WHAT THIS MEANS IN PRACTICE

AI representation is a system, not a metric. The seven layers are connected, and the leverage points sit upstream of the symptoms. Solving the visible problem, omission or misdescription, usually means solving an underlying structural one: authority, entity, or recency.

Commercial implications

The business consequences of misrepresentation are not abstract. Independent research published in the last eighteen months quantifies the cost across three dimensions: trust, revenue, and brand perception. Each is amplified when the misrepresentation is occurring on a surface the business does not control and may not be aware of.

Lost trust at the point of inaccurate information

Syndigo's State of Product Content survey, reported by Retail Customer Experience, found that 73 per cent of consumers would think less of a brand when they find incomplete or inaccurate online product information, up from 62 per cent the previous year. Fifty per cent had abandoned a purchase because of missing information, 35 per cent returned a product because it didn't meet expectations, and 83 per cent would abandon a site with insufficient product information²⁰.

eCommerce News, reporting Meltwater and YouGov research from April 2026, found that 89 per cent of Australians want stronger government regulation of AI-generated content, 86 per cent want AI-generated content clearly labelled, and 62 per cent say misleading AI content would damage their trust in a brand. Seventy-three per cent worry about AI misinformation²¹.

73%

Of consumers think less of a brand when they find incomplete or inaccurate online information (Syndigo, reported via Retail Customer Experience).

Lost revenue at the point of failed search



Eleven per cent will permanently abandon a brand after a single failed search.

Elastic, 2026 Australian Search Survey

Elastic's 2026 search survey, drawing on 1,600 consumers, found that 55 per cent are willing to pay more on a competitor's site to avoid poor search experiences. Sixty-two per cent use external search engines when on-site search fails, and 78 per cent of those end up purchasing from a competitor. Eleven per cent will permanently abandon a brand after a single failed search, and 60 per cent perceive brands with irrelevant search results as technologically behind¹⁹.

Although Elastic's research focuses on on-site search and external search engines, the same dynamic now applies to AI-mediated search. When AI returns inaccurate, incomplete or hedged information about a business, the buyer turns to alternatives. The buyer who arrives at a business via AI is the buyer most willing to leave it for a competitor.

The hallucination cost layer

Korra's August 2025 analysis estimated that hallucination-related misinformation cost enterprises US\$67.4 billion in losses in 2024, with 47 per cent of executives reporting that they had acted on faulty AI content¹⁶. While this estimate originates from a commercial vendor and should be treated as directional rather than rigorous, the directional point is clear: organisations are increasingly acting on AI-generated content that is wrong, and absorbing the cost of doing so.

The trust environment is hostile



65 per cent of Australians believe AI creates more problems than it solves.

Roy Morgan Research, October 2025

Roy Morgan's October 2025 research found 65 per cent of Australians believe AI creates more problems than it solves, with only 35 per cent thinking it solves more problems than it creates. Women and older Australians were more sceptical²². The buyer environment for AI-mediated representation is therefore not neutral. Australians are simultaneously using AI more, and trusting it less. Misrepresentation in this environment compounds: the buyer assumes some AI output may be wrong, and applies that assumption selectively, often to the brand they are evaluating.

THE COMMERCIAL READING

Inaccurate AI representation is not a soft cost. It is a measurable contributor to lost trust, abandoned purchases, perception of technological inadequacy, and brand defection to competitors. The data on these consequences predates AI. The arrival of AI as a primary discovery surface concentrates them onto a channel the business does not directly control.

The policy gap

Australia's policy and guidance environment for AI is active, but it places the responsibility for safe and accurate AI use squarely on individual businesses. There is no regulatory mechanism that protects a business from being misrepresented by AI systems, and no compliance regime that compels remediation. Representation governance, in practical terms, sits with the business itself.

Government guidance focuses on internal AI use

The National Artificial Intelligence Centre's October 2025 Guidance for AI Adoption sets out six practices: assigning accountability through a senior AI governance lead, assessing impacts and planning to prevent bias, measuring and managing risk in proportion to AI type, sharing essential information including disclosure of AI use, testing and monitoring continuously, and maintaining human control with override capability²³.

The ACSC's 2025 small business cyber security and AI guidance reinforces the same direction: verify AI outputs, train staff, establish internal policies, and maintain human oversight¹⁷. Both documents focus, appropriately, on how a business should use AI internally.

The National AI Plan defers to existing law

The National AI Plan, summarised by the Association of Corporate Counsel, aims to harness AI for economic and social potential by establishing a National AI Centre and an AI Adopt Program to help SMEs implement AI. It emphasises using existing legal frameworks for privacy, consumer protection, competition and workplace relations rather than introducing new AI-specific laws, and encourages firms to create internal AI governance frameworks²⁴.

The Conversation, in a December 2025 critique, argued that the plan's safety measures are minimal, that relying on existing laws may be inadequate, and that determining liability when AI causes harm remains structurally difficult under current arrangements²⁵.

Why this matters for AI representation

None of the existing policy or guidance addresses how AI systems represent a business externally. The frameworks address how a business uses AI; they do not address how AI uses the business. There is no mechanism for a business to compel an AI platform to correct an inaccurate description, and the existing legal frameworks for privacy, consumer protection and competition are not well suited to enforcing accuracy in dynamically generated AI outputs.

The practical implication is simple. AI representation is now a strategic risk that businesses must govern themselves. The signal accuracy, entity definition, source quality and authority structure that AI systems use to describe a business are within the business's influence, but only if the business takes ownership of them. Waiting for regulation to address misrepresentation is not a viable strategy.

THE GOVERNANCE REALITY

No current Australian regulatory mechanism directly ensures that AI-generated business descriptions are corrected. The policy direction is to enable AI adoption and rely on existing law for downstream harm. Representation governance is now an executive responsibility, not a compliance one.

What this means for your business

This report is not a how-to. The methodology behind the audit work is proprietary, and the corrective response varies materially by business, sector and signal profile. What follows are the strategic priorities the evidence supports for any Australian business that takes AI representation seriously.

01 Treat AI as a representation surface, not a channel

Channel thinking optimises for traffic. Representation thinking optimises for accuracy, completeness and authority. Channel thinking measures volume. Representation thinking measures fidelity. The first strategic shift is recognising that AI is not a new performance channel to optimise. It is a new representation surface to govern. The KPIs are different. The cadence is different. The team responsibility is different.

Why traditional SEO is not enough

Search engine optimisation remains important, but it does not fully solve AI representation risk. A business can rank well in traditional search and still be omitted from AI recommendation queries, described incorrectly, compared unfavourably, or presented with hedging language that reduces buyer confidence.

Traditional SEO is primarily concerned with ranking, crawlability, keyword relevance, backlinks and page performance. AI representation is concerned with signal integrity: whether AI systems can understand the entity, extract the right facts, distinguish the business from competitors, describe its value accurately, and support that description with credible public evidence.

The corrective work is broader than adding keywords or publishing more content. It includes clearer primary-source statements on the website, consistent entity information across public surfaces, structured data, proof content, third-party authority signals and removal or correction of outdated sources.

Generative Engine Optimisation and AI search visibility are related, but citation volume alone is not the same as representation accuracy. Being mentioned by AI is not enough if the description is wrong, weak, outdated or commercially misleading.

02 Audit before you act

Most businesses do not know how they are being represented in AI outputs. Direct name queries are not enough. Buyer-perspective category queries, comparison queries and qualifier queries surface the actual representation pattern. Before any remediation, the question worth answering is: what is AI saying about us, to the buyer, when the buyer does not lead with our name.

Practical next step

Before changing website content or chasing AI visibility, check how AI is currently representing the business. For deeper evidence, use the [AI Representation Audit](#).

**AI REPRESENTATION
AUDIT**

03 Prioritise the website over peripheral surfaces

The website remains the primary truth source AI systems calibrate against. Investments in directory profiles, social presence and PR have value, but they sit downstream of website authority. Where website signals are weak, peripheral signals dominate, often inaccurately. The first leverage point in nearly every audit was the business's own website.

04 Make accuracy structurally extractable

AI systems do not reward marketing copy. They reward extractable evidence. Plain statements of category, services, locations, methodology and credentials, written in declarative language, structured for machine readability, repeated consistently across surfaces. The discipline required is editorial, structural and evidential. It is not about more content. It is about content that AI can confidently extract from.

05 Resolve entity definition early

Where Entity Confusion is present, it tends to be the root cause of multiple downstream problems. Resolving entity definition, consistent legal and trading names, consistent founder attribution, consistent location signals, and consistent disambiguation from similarly named entities, is high-leverage work. It rarely appears on a marketing roadmap, but it materially changes how AI describes the business across every other layer.

06 Invest in authority before you invest in volume

Authority is the layer where the most confident competitors win. Volume of content does not produce authority. Substantive evidence, credible third-party citation, structured methodology content, proof points and credential surfaces produce authority. AI systems hedge less when the evidence is concrete and extractable. This is also the layer where premium-positioned businesses pay the highest cost when it is weak: the gap between what they can deliver and what AI says they can deliver is widest.

THE STRATEGIC POSTURE

AI representation is governed, not optimised. The work is editorial and structural, not promotional. The leverage sits with accuracy, authority and entity definition, not volume or activity. Businesses that treat this as a strategic surface now will compound the advantage. Businesses that wait will pay an asymmetric cost.

About this report

The audit work

This report draws on observations from 320 AI Representation Audits across 56 businesses conducted by Iconic Marketing during the build and testing phase of its AI representation methodology. Each audit tested AI platform outputs across repeated runs, with three runs per AI platform per audit, against a structured set of buyer-language queries calibrated to the audited business's category, location and proposition.

About Iconic Marketing

Iconic Marketing was built from experience in marketing for multinational organisations, where large budgets, experienced teams, advanced technology and sophisticated strategy were the standard. When we looked at what was available to businesses in regional Australia, we saw a clear gap.

The same level of marketing strategy, digital capability and commercial thinking that larger businesses and Sydney-based organisations often have access to was not always accessible, affordable or available to businesses in Newcastle, the Hunter and Penrith. That is why Iconic Marketing exists.

Founded by Renee Gersteling, with over 25 years of senior sales and marketing leadership across national and international organisations, Iconic Marketing was created to bring that level of strategic thinking and practical execution to businesses that want to grow without the overhead of a large internal marketing team.

We help businesses strengthen how they are positioned, discovered and understood across traditional search, digital channels and AI-generated results. Because in today's market, visibility alone is not enough. Your business needs to be found, accurately interpreted and trusted before a buyer ever makes contact.

If you want to test how AI is representing your business

START HERE · FREE

AI Representation Snapshot

An early view of how AI may be representing your business, and where signals are unclear.

[Learn more ›](#)

EVIDENCE · SCORING

AI Representation Audit

Risk scoring and prioritised corrective actions across the seven layers.

[Learn more ›](#)

IMPLEMENTATION

AI Representation Sprint

Hands-on correction of weak, outdated or inconsistent public signals.

[Learn more ›](#)

WEB

iconic.marketing

EMAIL

connect@iconic.marketing

LOCATIONS

Newcastle | Sydney | Penrith

ON THE AI REPRESENTATION RISK FRAMEWORK™

The seven-layer framework presented in this report is proprietary to Iconic Marketing. It is presented here as the organising lens for the audit observations and the wider evidence base. The framework is the foundation of Iconic Marketing's diagnostic and remediation work for clients across Australia.

References

Sources cited throughout this report. Where a source originates from a vendor or commercial publication, that context is included in the body where the source is referenced.

- 1 **Australian Bureau of Statistics. Counts of Australian Businesses, including Entries and Exits, July 2021 to June 2025. Feb 2026.**
<https://www.abs.gov.au/statistics/economy/business-indicators/counts-australian-businesses-including-entries-and-exits/latest-release>
- 2 **Australian Small Business and Family Enterprise Ombudsman. Small Business Data Portal: Number of small businesses in Australia. 2024 data.**
<https://www.asbfeo.gov.au/small-business-data-portal/number-small-businesses-australia>
- 3 **Australian Bureau of Statistics. Development of a Composite Indicator for Business Digital Intensity in Australia (Digital Intensity Index). 2023.**
<https://www.abs.gov.au/articles/development-composite-indicator-business-digital-intensity-australia-digital-intensity-index>
- 4 **RMIT, Swinburne and Telstra. 2025 Australian Digital Inclusion Index. Sept 2025.**
https://digitalinclusionindex.org.au/wp-content/uploads/2025/10/ADII-Report-2025_V6-Remediated.pdf
- 5 **auDA. The Digital Lives of Australians 2025. Apr 2025.**
<https://www.auda.org.au/news-insights/research-reports/digital-lives-of-australians/>
- 6 **Adobe. From Assistants to Agents: The AI Evolution in Australia. July 2025.**
<https://news.adobe.com/news/downloads/pdfs/2025/07/07022025-assistants-to-agents-the-ai-evolution-in-australia-report.pdf>
- 7 **PayPal / Commsroom. PayPal report: 48% of Aussies now use AI for product searches. 2026.**
<https://commsroom.co/paypal-report-shows-48-of-aussies-now-use-ai-for-product-searches/>

- 8 Reuters / EBU / BBC. AI assistants make widespread errors about the news, new research shows. Oct 2025.**
<https://www.reuters.com/business/media-telecom/ai-assistants-make-widespread-errors-about-news-new-research-shows-2025-10-21/>
- 9 Reserve Bank of Australia. Bulletin: Technology Investment and Artificial Intelligence, What Are Firms Telling Us? Nov 2025.**
<https://www.rba.gov.au/publications/bulletin/2025/nov/technology-investment-and-ai-what-are-firms-telling-us.html>
- 10 Indeed Hiring Lab. Nothing Artificial About Australian AI Adoption. Apr 2026.**
<https://www.hiringlab.org/au/blog/2026/04/01/nothing-artificial-about-australian-ai-adoption/>
- 11 Governance Institute of Australia. AI Deployment and Governance Survey. Aug 2025.**
<https://www.governanceinstitute.com.au/app/uploads/2025/04/AI-deployment-and-governance.pdf>
- 12 UTS Human Technology Institute. Safe Artificial Intelligence Adoption Model Report. Jan 2025.**
<https://www.uts.edu.au/news/2025/02/hti-report-reveals-ai-exceeding-expectations-smes-many-face-range-adoption-barriers-they-need-help-overcome>
- 13 BizCover. Small Business AI Report 2025. 2025.**
<https://www.bizcover.com.au/blog/ai-transforming-australian-small-business-sector/>
- 14 Australian Bureau of Statistics. Characteristics of Australian Business 2021-22. Aug 2024 release.**
<https://www.abs.gov.au/statistics/industry/technology-and-innovation/characteristics-australian-business/2021-22>
- 15 The Guardian. Google AI Overviews put people at risk of harm with misleading health advice. Jan 2026.**
<https://www.theguardian.com/technology/2026/jan/02/google-ai-overviews-risk-harm-misleading-health-information>
- 16 Korra. The \$67 Billion Warning: How AI Hallucinations Hurt Enterprises. Aug 2025.**
<https://korra.ai/the-67-billion-warning-how-ai-hallucinations-hurt-enterprises-and-how-to-stop-the-m/>

- 17** ACSC / ASD. Artificial Intelligence for Small Business: Managing Cyber Security Risks. 2025.
<https://www.cyber.gov.au/business-government/secure-design/artificial-intelligence/artificial-intelligence-for-small-business>
- 18** ACCC. Consumers warned about 'ghost stores' imitating Australian businesses. July 2025.
<https://www.accc.gov.au/media-release/consumers-warned-about-ghost-stores-imitating-australian-businesses>
- 19** Elastic. 2026 Australian Search Survey: Predictions from 1,600 Consumers. 2026.
<https://www.elastic.co/blog/australian-search-research-2026>
- 20** Syndigo (via Retail Customer Experience). State of Product Content 2024. 2024.
<https://www.retailcustomerexperience.com/news/lack-of-information-poor-product-top-reasons-for-abandoning-purchase/>
- 21** eCommerce News (Meltwater & YouGov survey). Australians demand AI labelling as trust concerns grow. Apr 2026.
<https://ecommercenews.com.au/story/australians-demand-ai-labelling-as-trust-concerns-grow>
- 22** Roy Morgan Research. Growing majority of Australians believe AI creates more problems than it solves. Oct 2025.
<https://www.roymorgan.com/findings/10048-ai-press-release-october-2025>
- 23** National AI Centre (Australia). Guidance for AI Adoption. Oct 2025.
<https://www.ai.gov.au/staying-safe-and-responsible/essential-ai-practices/guidance-ai-adoption-implementation-guidance>
- 24** Association of Corporate Counsel. Australia has a National AI Plan. Now What? Dec 2025.
<https://www.acc.com/australia-has-national-ai-plan-now-what>
- 25** The Conversation. Australia's official plan for AI safety isn't much more than a single dot point. Will it be enough? Dec 2025.
<https://theconversation.com/australias-official-plan-for-ai-safety-isnt-much-more-than-a-single-dot-point-will-it-be-enough-276962>