
The Privacy Capability Labour Market Review

What 167 Australian job advertisements reveal about how organisations describe, hire and structure privacy work

Author: Fiona Chan

1. Executive summary

The Privacy Capability Labour Market Review at a glance

Between 11 May and 2 June 2026, this review captured 167 Australian privacy and privacy-adjacent job advertisements from LinkedIn, SEEK and Indeed across six manual collection sprints. Of those, 101 were coded as substantively privacy-focused. Around them sits a second sample of 55 privacy-adjacent near-miss roles that surfaced under privacy search terms without meeting the inclusion threshold: 43 where privacy is named as scope inside another function, and 12 where the advertisement does not name privacy at all.

The review reads the corpus as one window into how the Australian privacy market is articulating privacy capability in 2026, that is, what kind of capability organisations are buying when they describe a privacy role for the open market.

Several of the findings below use the four levels of my [Privacy Capability Framework](#): Rule Knowledge, Operational Fluency, Risk-Disciplined Judgement and System Shaping. A job advertisement cannot show what a practitioner can do, but it does signal the level the organisation is asking for, and the review coded those signals into four bands, L1 to L4: L1 for rules awareness, L2 for producing privacy work such as assessments, registers and policies, L3 for embedding privacy into business operations, and L4 for operating-model, contested-judgement and leadership work. Signals can mislead, and Section 6 tests them. At L4 they mostly do: advertisements that look as if they are asking for System Shaping are, read closely, asking for Risk-Disciplined Judgement.

Headline findings

1. Compliance and legal advisory framings dominate. The two largest archetypes, Compliance / Privacy Governance (27 roles) and Legal Privacy Counsel (22 roles), together account for 49% of the included sample. The framing language across these 49 roles is highly repetitive: policy and standard maintenance, framework alignment, legal advisory, regulator engagement, stakeholder partnership.

2. Embedded privacy capability is meaningful but secondary. The Operational Privacy archetype carries 17 roles (17%). The work is recognisably distinct from compliance stewardship and from legal advisory: practitioners are hired to embed privacy into product, vendor, data and AI workflows. But the volume is smaller, the L4 leadership roles are disproportionately project-shaped, and the language inside the L3 core tilts toward embedded expertise rather than embedded authority.

3. Independent privacy assurance is almost absent from advertised demand. The Risk / GRC Privacy archetype carries three of the 101 included roles. The independent second-line privacy challenge function the three-lines-of-defence operating model assumes is essentially missing from the privacy roles advertised in this corpus.

4. AI governance has not settled as a discipline. The corpus shows three different design patterns operating simultaneously: AI governance absorbed into privacy leadership (Sportsbet, QBE, MUFG); AI governance inside data leadership (AFCA, Helia, Capricorn, plus near-miss data-and-AI executives such as Pacific Blue); and standalone AI governance functions alongside privacy and data (Slade, Profectus, plus a near-miss executive role at IAG). None is winning. Sophisticated Australian organisations are making three different bets on the same underlying question at the same time.

5. Advertised demand concentrates in the L3 operational-embedding band (50% of includes) and the L4 leadership band (28%). The bands record the signals an advertisement gives for each framework level (see above). L4 concentrates in compliance (seven roles), AI Governance Hybrid (six), and legal, statutory officer and Operational Privacy leadership roles (five each). L4 is absent from privacy engineering and privacy assurance leadership. L1 is effectively absent from the included sample; L2 is small (10%) and concentrated in compliance and technical project work.

6. The L4 signal points to Level 3: no advertisement asks anyone to change how privacy decisions get made. The L4 band picks up four signals: operating-model design, contested judgement, senior leadership and authority. Read closely against the framework, none of the 101 advertisements puts the organisation's own decision conditions on the table, which is what System Shaping requires. That matters because System Shaping is the level [privacy as a decision capability](#) depends on: Risk-Disciplined Judgement brings judgement to a decision, and System Shaping shapes the system that determines whether such judgement can reach decisions repeatedly. Nine describe building governance inside conditions the organisation has already settled: a framework, a lifecycle, a set of standards, designed to fit a model nobody is being asked to change. Three of those nine are greenfields AI governance builds, where the architecture has to be described because none exists to inherit. Thirteen of the L4 and L3–L4 roles ask for risk-disciplined judgement on hard questions, which the framework places a level lower, and twelve are mixed. The corpus contains no Chief Privacy Officer mandate. Sections 3 and 16.4.2 set out how each band maps to a framework level; the original classifications are unchanged.

7. One in four privacy-search results is a role where privacy has been absorbed into another function. The 55 near-miss roles split two ways. Forty-three name privacy as scope inside an adjacent function and distribute across six absorption patterns plus a long tail: general counsel (privacy as one practice area), AI/ML/data leadership (privacy minor), cyber and security leadership (privacy absorbed), senior compliance and risk leadership (privacy as one sub-framework), FOI/Information access roles (privacy adjacent), and data governance leadership (privacy as one consideration). At 26% of raw captures, that is the absorption finding. The other twelve name privacy nowhere in the advertisement and reflect the width of the search terms rather than how work is designed. The cumulative near-miss share stayed between 29% and 36% as the corpus grew, closing at 33%; the absorption share alone moved between 20% and 28%, closing at 26%. Privacy capability in the Australian market disperses systematically into adjacent functions, and the included-only reading understates the dispersion.

8. When privacy arrives, and who has to answer it, are the under-described variables. Across the corpus, role designs name the expertise and leave out the two things that determine whether the expertise can matter: the point at which privacy enters a decision, and whether anyone is obliged to engage with the judgement before deciding. The advisory-framing pattern is dominant even inside Operational Privacy roles, which sit closest to the business and where the omission costs most. The independent-challenge pattern is largely absent. The operating-model implication is that the roles being advertised do not specify when privacy is expected to reach a decision or who must answer it, which is the part of the operating model that matters most when a decision is contested.

What this implies

The corpus is consistent with the proposition that the Australian privacy market has overbuilt compliance posture and underbuilt embedded privacy capability: compliance and legal roles dominate, embedded roles are fewer and often project-shaped, and independent assurance barely appears. Job advertisements are one source of evidence, so the corpus supports the proposition without settling it.

Underneath that sits a narrower observation about what the advertisements specify. The labour market is hiring people to perform privacy work, and it is much less precise about the capability required to change how privacy decisions are made. Read against the practitioner capability framework, this suggests that organisations are recruiting for the visible outputs of their existing operating models while leaving System Shaping implicit, conflated with seniority, or absent from the advertised mandate. That is a finding about what the advertisements ask for. They cannot show what an organisation already has, what it needs or how it sees itself. My own reading of why, which comes from practice rather than from the corpus, is that organisations asking for senior privacy capability describe Risk-Disciplined Judgement because System Shaping is not yet widely recognised: few have seen what good looks like at that level, so they ask for the most senior capability they can picture.

For organisations: the corpus surfaces five operating-model design decisions worth taking deliberately rather than by default: rebalancing toward operational embedding, choosing an AI governance pattern consciously, revisiting the cyber-privacy interface, choosing between concentrated and distributed Privacy Counsel designs, and addressing the Line 2 absence explicitly. The market default is the compliance-led shape; for some organisations that is the right answer, for others it is not, and AI governance and Privacy Act reform are adding pressure to that choice in 2026.

For privacy practitioners: AI governance scope is concentrated in the most senior privacy roles advertised, so AI governance capability layered onto senior privacy expertise is the combination the market most often asks for at the top. The Legal Privacy Counsel track is the most clearly defined senior privacy career path. The embedded privacy career path is real but smaller and more project-shaped. Statutory officer (DPO) work is its own track with its own jurisdiction-specific entry

channels. Privacy engineering and privacy assurance career paths exist but are currently labelled as something other than privacy.

For the discipline: the most important variable to watch over the next eighteen to twenty-four months is whether the labour market begins to describe when privacy enters a decision and who has to answer it. That language shift, if it happens, will signal that the operating-model rebalancing argument has reached the level of role design, which is the level at which the rebalancing actually happens. The current corpus does not show that shift yet. The next one might.

Methodology in one paragraph

The corpus was collected by hand from three job boards, LinkedIn (112 advertisements), SEEK (46) and Indeed (9), in six sprints between May and June 2026, using search windows that ranged from the same day to the past 30 days. Two collection modes, a broad sweep across multiple privacy query terms (sprints 01 to 04) and focused queries restricted to privacy-titled roles (sprint 05), with a top-up batch (sprint 06), produced complementary corpora. Every advertisement was coded against a nine-archetype typology, an L1–L4 capability framework, a 16-verb task taxonomy, and categorical AI governance and operational privacy signal strength fields. Coding was applied conservatively, with a documented QA process that surfaced 16 manual-review flags across the six sprints, all resolved. Full methodology, limitations, and reproducibility notes are in Section 16.

2. Why this review matters

The question this review answers

Australian organisations are visibly busy on privacy. They are publishing privacy policies, running training, appointing officers, building governance committees, and increasingly naming AI governance as a priority. Most large employers can point to substantial recent investment in privacy capability and many are recruiting actively for it.

The question this review answers is narrower and more specific: when organisations describe what privacy capability they want to purchase, what kind of capability are they actually asking for?

Neither of the two measures the market usually relies on closes the gap between *privacy activity* and *embedded privacy capability*. Hiring volume measures appetite. Compliance maturity assessments measure where the privacy program has been built out against a benchmark framework. Neither directly answers what the work the market is hiring for actually does.

Job advertisements are an imperfect but informative source for that question. They are the moment at which organisations translate their internal understanding of "what we need in privacy" into a marketing description for the open market. The advertisement reflects the design assumption: what the organisation thinks privacy capability looks like and how it should be packaged in a role. This review reads 101 included Australian privacy job advertisements as one window into that design assumption.

Why now

Three contextual shifts make this question more contested in 2026 than it would have been three years ago.

Privacy Act reform has changed the operational expectations. The first tranche of Privacy Act amendments, passed in late 2024, introduced a statutory tort for serious invasions of privacy (in force since June 2025), transparency obligations around automated decision-making (commencing December 2026), and a lifted enforcement toolkit for the regulator. The second tranche was released as an exposure draft, the Privacy Amendment (Personal Data Protection) Bill 2026, on 31 August 2026; consultation closed on 18 September and the Bill remains under government consideration. Organisations now sit between a regulatory regime that has tightened around them and one that is expected to tighten further. The operational capacity that compliance posture alone provided in the OAIC's first decade of enforcement is no longer sufficient. Privacy capability that lives only in policy documents, breach registers, and post-incident response can be evidenced as present and still leave an organisation operationally exposed in ways the regulator can now act on.

AI governance has not settled into a discipline. Many Australian organisations are now making AI deployment decisions without a stable answer for whose role it is to govern them. The corpus reflects this directly: the AI Governance Hybrid archetype contains roles where AI governance has been bundled with privacy, the near-miss sample contains roles where AI governance is being stood up as a discrete executive function, and the Legal Privacy Counsel archetype concentrates the explicit responsible-AI scope language in senior counsel roles. Three patterns, three different design assumptions, all in the same labour market at the same time. Section 10 takes this up in detail; it matters here because the AI governance question forces embedded privacy capability to the front of the queue in ways that the standard compliance program picture does not handle well.

The argument is ahead of the evidence. The view that privacy programs have overbuilt compliance and underbuilt the capability to embed privacy in the business is no longer a fringe one, but it has mostly been made without evidence about what organisations actually hire for.

This review supplies one piece of that evidence: whether the way organisations describe and buy privacy capability is consistent with the argument. Its answer, taken up across Sections 7 to 14, is yes, with qualifications about where the argument holds, where it doesn't, and what the market is doing about AI governance that the argument does not yet describe well.

Why job ads, specifically

There are at least four ways to gather evidence about how the Australian privacy market is conceptualising privacy capability. Each tells a different story and answers a different question. Naming the alternatives makes clear what this review's source choice is good for and what it isn't.

Practitioner surveys ask privacy practitioners what they do and what they think they need. They yield the inside view: how the work feels from inside the role. They are subject to selection effects (practitioners who respond to surveys are not a random sample of the practitioner population) and to social-desirability framing in how respondents describe their own work.

Maturity benchmarking exercises assess privacy programs against a framework (IAPP, ISO 27701, NIST Privacy Framework, OAIC guidance). They yield a program-design view: how organisations have structured their privacy capability against an external reference. They are useful for tracking the build-out of standard program components. They are less useful for surfacing the gap between described capability and operational capability, because the benchmarking framework defines what "mature" looks like, and a program can score highly on benchmark dimensions while leaving embedded privacy capability underbuilt.

Regulatory enforcement and breach data show where organisations have actually failed. The OAIC's notifiable data breach reporting and enforcement actions trace failure patterns over time. This evidence base is the most powerful for arguments about where privacy capability has been thin in practice, but it is lagging, redacted, and limited to the subset of failures that crystallised into a notifiable event.

Job advertisements show what organisations are buying. They are forward-looking (the role is the work the organisation expects to need over the next several years), externally articulated (the organisation has to describe the role in terms a candidate can recognise), and quantitatively tractable (a corpus can be assembled and coded against a common framework). They have one major weakness, which is that the description in the ad is not necessarily what the role does once filled. Section 3 deals with it.

The four sources answer different questions. This review takes job advertisements because the question at hand, what the market is buying when it buys privacy capability, is the question advertisements are the best instrument for.

The proposition this review checks

The review checks one proposition: a claim that follows from my thesis and that a corpus of job advertisements can test. I have been developing it across 2025 and 2026 through my research and practice:

The Australian privacy market has overbuilt compliance posture and underbuilt embedded privacy capability. Organisations have the artefacts. They do not consistently have the embedded capability to make privacy decisions inside operational and product workflows. The gap is most visible where AI governance, data governance and privacy intersect at decision points the standard privacy program framework was not designed to handle.

This proposition is one consequence of an argument I make elsewhere, on the [Privacy as a Decision Capability](#) page: that privacy fails in the decisions organisations make rather than in the documents they produce. The corpus cannot test that argument, which is about what happens inside organisations once a decision is live. It can test this one consequence of it, by showing what the labour market is buying when it buys privacy.

The proposition is a claim to be checked, not a conclusion assumed. The coding was set to underclaim, and where the corpus diverges from what the proposition would predict, notably in how AI governance is being organised and where statutory officer roles are appearing, the report says so and adjusts the picture.

What follows

The sections that follow read the corpus in three movements.

Sections 5 and 6 describe what the corpus contains: the seven role archetypes the market hires for, the capability level distribution across them. These are descriptive sections; they do not advance

the argument.

Sections 7 to 11 are the substantive findings sections. Section 7 takes up what the corpus over-indexes on (compliance and legal advisory work, and the framing language they share). Section 8 takes up what the corpus under-describes (independent privacy assurance and contested-judgement language). Section 9 profiles what privacy work inside the business actually looks like in the 17 roles that fall in that archetype. Section 10 reads what AI governance hiring patterns reveal about whether the discipline has settled. Section 11 reads the 55 near-miss roles for the "where privacy disappears" pattern.

Sections 12 to 14 take up implications: what the patterns mean for organisations designing privacy programs, for privacy practitioners planning capability development, and for the operating-model questions the broader proposition raises. These sections lift the findings into the design questions that organisations and practitioners face in 2026.

A reader interested in the substantive argument can read Sections 7, 8, 9, 10 and 11 as a block and arrive at the same conclusions as a full read. A reader interested in the operating-model implications can begin at Section 12 and use the earlier sections as evidence reference. The report is built to be read either way.

3. Methodology

What this review is

This review is a structured desktop study of how Australian organisations describe privacy capability when they hire for it. It analyses 167 job advertisements identified on LinkedIn, SEEK and Indeed during the collection window and codes 101 as substantively privacy-focused. The remainder are either privacy-adjacent near misses kept for separate analysis (55) or duplicate listings excluded from the sample (11).

It is not a labour-market census. It does not estimate the total number of privacy roles in the Australian market, measure salary, hiring volumes or time-to-fill, or infer what privacy teams actually do inside organisations. It describes how privacy capability is being articulated in job advertisements identified during the collection window.

That distinction is deliberate, and it governs every claim in the report. A job advertisement is a marketing document. It tells a reader what an organisation believes privacy capability looks like, framed as something the organisation wants to purchase. It does not tell the reader whether the role, once filled, actually produces the outcomes the ad describes. The corpus is therefore evidence of how privacy capability is described and purchased, not conclusive proof of how privacy programs operate in practice.

Source and collection window

The corpus was assembled from three job boards in six sprints between 11 May 2026 and 2 June 2026: LinkedIn (112 advertisements), SEEK (46) and Indeed (9). Of the 101 included advertisements, 63 came from LinkedIn, 29 from SEEK and 9 from Indeed. Searches were limited to Australia, with posting windows that ranged from the same day to the past 30 days depending on the search.

Collection was entirely manual. The project's method boundary, set before collection began, prohibits scraping, automated extraction, browser-extension automation, and any other form of bulk collection against LinkedIn or any other job board. Ads were identified by human search, opened individually, and their text and metadata captured for downstream coding. AI tools were used only after lawful manual collection to assist with intake, normalisation, classification, evidence extraction, and synthesis. No AI tool collected, scraped or generated any source content.

All employer references in this report derive from publicly advertised roles during the collection window. The captured text of every advertisement is retained as the evidence for each employer-specific claim.

Two-mode collection

The corpus was assembled in two collection modes, and the report distinguishes them where the distinction matters.

The first four sprints used broad search queries (privacy, privacy officer, privacy advisor, DPO, data protection, AI governance privacy, data governance privacy, GRC privacy, compliance manager privacy, and similar terms) and produced 141 ads at a 57% include rate. The fifth sprint used focused queries narrowed to privacy-titled roles only (Privacy Counsel, Privacy Officer, Privacy Manager, Privacy Compliance, DPO) and produced 21 ads at a 76% include rate. The sixth sprint was a small top-up batch of 5 ads, of which 4 were included.

The two modes produced complementary corpora. The broad-sweep sprints captured rich near-miss material: roles where privacy responsibility is absorbed into adjacent functions such as legal counsel, cyber leadership, freedom-of-information management, or AI governance. The focused-query sprint produced direct hits on the included sample. Both are necessary. A review that only ran focused privacy-titled queries would have missed the near-miss pattern and produced a tidier but misleading picture of where privacy capability sits in the Australian market.

This is the most important methodological disclosure in the review. The combined include rate across the two modes is 60%, but a single-mode collection would have produced a different number, and a different report. The corpus shape reflects the collection strategy.

What was coded for each advertisement

Every advertisement was coded on six dimensions:

Role archetype. Each ad was assigned a primary archetype from the nine-category typology: Legal Privacy Counsel; Compliance / Privacy Governance; Risk / GRC Privacy; Operational Privacy; Data / Technology Privacy; DPO / Regulatory Interface; AI Governance Hybrid; Privacy-Adjacent Near Miss; or Other / Mixed. Where the role was genuinely mixed, a secondary archetype was also recorded.

Capability level (L1–L4). Each included ad was assigned a level from L1 (rules awareness and support tasks) to L4 (operating model design, contested judgement, leadership). The level reflects what the ad evidences as the role's substantive capability requirement, not the seniority of the title. The bands record the signals an advertisement gives for each level of the Privacy Capability Framework; the next subsection explains how they relate, which matters for how the L3 and L4 numbers can be read.

Task orientation. Each ad was coded against the project's 16-verb task taxonomy (advise, interpret, review, draft, maintain, monitor, report, assess, train, implement, embed, challenge, assure, design, govern, lead), capturing the primary and top secondary verbs the JD body actually used.

AI governance signal. A categorical signal capturing whether the ad mentions AI governance, responsible AI, or AI-adjacent privacy work, and at what strength (Strong, Some, Mentioned, or None).

Operational privacy signal strength. A judgement of how strongly the role's described work points toward operational embedding (Strong) versus compliance posture (Weak), with Moderate and intermediate calls used liberally.

Capability and operating-model evidence phrases. Verbatim quotes from the JD body that anchor each archetype and capability-level call, so the coding is auditable against the source text.

The coding framework was applied conservatively. The QA guardrail was followed throughout: it is better to underclaim than to inflate operational privacy signals. Where role scope was genuinely mixed (for example, a senior counsel role where privacy is one of several practice areas), the ad was either coded as a near miss with a clear category name, or flagged for manual review.

How the bands relate to the framework

The L1 to L4 bands are signals of the four levels of the [Privacy Capability Framework](#): Rule Knowledge, Operational Fluency, Risk-Disciplined Judgement and System Shaping. The framework describes what a practitioner can do across a body of work. An advertisement cannot show that, because it is written before anyone is appointed, but it does show what each level looks like when an organisation writes it into a role: the rules to be known, the work to be produced, where the role sits in the business and the scope it is given. The bands code those signals.

Signals are imperfect, and the report tests them rather than assuming them.

- **L1 and L2** signal Rule Knowledge and Operational Fluency closely. Knowing the rules and producing privacy work are visible in how an advertisement describes the job.
- **L3** codes operational embedding: a role placed in the business, close to product, vendor, data and AI decisions. That is the position in which Risk-Disciplined Judgement is exercised, but position is not proof, and an embedded role can still be asked only to produce assessments.
- **L4** codes four signals together: operating-model design, contested judgement, senior leadership and authority. Only the first belongs to System Shaping. Contested judgement is Risk-Disciplined Judgement, and seniority and reporting lines are positions rather than capabilities.

Section 6 tests the L4 signal by re-reading all 28 L4 and 6 L3–L4 advertisements against the framework's definition of System Shaping. The result changes how the L4 count can be read: the advertisements that look most like requests for System Shaping are, read closely, requests for Risk-Disciplined Judgement, and none asks for System Shaping.

The bands were coded in May and June 2026, before the framework reached its final form in August, and the codes are kept unchanged so the corpus stays reproducible and the 2027 edition can be

compared with this one. Where the report means a framework level it writes the name out in full; L1 to L4 always mean the advertised signals. Section 16.4.2 sets out the mapping in full.

Quality assurance

Every advertisement passed through a four-stage process: intake, classification, QA review, and final coding. The QA step applied a separate checklist: whether the primary archetype is supported by evidence, whether the capability level is too high, whether operating-model signals are being over-read, whether vague terms like "strategic" or "stakeholder engagement" are being over-valued, and whether the role has enough evidence to be counted in the report.

Of the 167 ads processed, 140 were accepted as classified without a flag, 16 were flagged for human review and subsequently resolved or revised, and 11 were excluded as duplicate listings of the same role re-advertised by different recruiters or in different sprints.

What this review can and cannot claim

This review can describe how privacy capability is articulated in the Australian job market during the collection window. It can identify dominant archetype patterns, capability-level distributions, and recurring framing language. It can demonstrate where privacy responsibility is being absorbed into adjacent functions, and where it is being treated as a distinct discipline. It can show the shape of AI governance hiring across the corpus.

It cannot claim that the patterns observed in the corpus represent the whole Australian market. It cannot claim that what the ads describe is what the roles actually do once filled. It cannot generalise to hiring channels other than the three job boards searched. It cannot infer organisational behaviour from advertisement language alone.

Where the report draws conclusions, those conclusions are about how privacy capability is being described and purchased, not about how privacy work is being done. That is a narrower claim than the full operating-model picture, and it is the right claim for an evidence base of this kind.

4. Sample overview

At a glance

167 privacy and privacy-adjacent job advertisements were captured across six collection sprints. Of those, 101 were coded as substantively privacy-focused and form the report's included sample. Another 55 were coded as privacy-adjacent near misses and are analysed separately in Section 11. 11 were excluded as duplicate listings of the same role re-advertised by different recruiters or in different sprints.

STATUS	COUNT	SHARE OF TOTAL RAW
Included	101	60%
Near miss	55	33%
Excluded (duplicates)	11	7%
Total raw captured	167	100%

Two numbers in that table matter for the rest of the report.

The first is the included share at 60%. The other 40% of the corpus is not noise. The near-miss roles are the substantive evidence behind the "where privacy disappears" finding in Section 11, and the duplicates surface how the recruiter and direct-employer channels overlap.

The second is the near-miss share at 33%. One in three job advertisements that surfaced under privacy-related search terms was not a privacy role. Most of those, 43 of the 55, are roles where privacy responsibility is real but is carried by an adjacent function: legal generalists, cyber leadership, FOI managers, senior compliance executives, data leadership, IT GRC, AML, records management, or standalone AI governance. That is 26% of all raw captures, and it is the report's absorption finding. The remaining 12 name privacy nowhere in the advertisement and are search-index adjacency rather than absorption; Section 11 separates the two and explains why the distinction changes the reading. The cumulative near-miss share stayed between 29% and 36% as the corpus grew, whichever collection mode produced the batch, and closed at 33%.

Archetype distribution


Each included advertisement was assigned a primary role archetype from the project's nine-category typology. The distribution across the 101 included ads is:

PRIMARY ARCHETYPE	INCLUDED COUNT	SHARE OF INCLUDES
Compliance / Privacy Governance	27	27%

PRIMARY ARCHETYPE	INCLUDED COUNT	SHARE OF INCLUDES
Legal Privacy Counsel	22	22%
Operational Privacy	17	17%
AI Governance Hybrid	14	14%
DPO / Regulatory Interface	13	13%
Data / Technology Privacy	5	5%
Risk / GRC Privacy	3	3%
Other / Mixed	0	0%

Primary role archetype across the 101 included advertisements

Counts equal percentages; n = 101. LinkedIn, SEEK and Indeed, 11 May to 2 June 2026.



Three things stand out from that table even before the substantive findings sections.

First, **Compliance / Privacy Governance is the single largest archetype**. Twenty-seven roles, slightly over a quarter of the included sample. Combined with the 22 Legal Privacy Counsel roles, just under half of all included ads describe privacy work primarily as a compliance or legal advisory function. The implications of that ratio are taken up in Sections 7 and 12.

Second, **Operational Privacy sits at 17%**. A meaningful share, but clearly secondary to the compliance and legal framings, and smaller than the AI Governance Hybrid and DPO archetypes combined. The ads in this archetype describe embedding privacy capability into business operations, product, data, vendor and AI workflows rather than producing artefacts for legal defensibility. Those roles are profiled in Section 9.

Third, **Risk / GRC Privacy sits at three roles across the entire corpus**, not three per sprint but three across six sprints, both collection modes, and a deliberate effort to surface this archetype through GRC-related search queries. The denominator for an archetype is the 101 included roles, since near-misses carry the near-miss code by construction; what makes three a finding is that it did not move when the search strategy did. This is the corpus's most robust signal that the operating-model structure many privacy programs assume, embedded privacy capability in Line 1 with independent Line 2 challenge, is essentially absent from the labour market. Section 8 returns to this.

The 55 near-miss ads are not included in this table. Their distribution by absorption category is analysed in Section 11. The two largest categories are general commercial counsel with privacy as one of several practice areas, and AI governance leadership with privacy absent or minimal, each of which carries implications for the operating-model argument the report makes in later sections.

Capability level distribution

Every included advertisement was assigned a capability level from L1 (rules awareness and support tasks) to L4 (operating model design, contested judgement, leadership). The level reflects what the ad evidences as the role's substantive capability requirement, not the seniority of the title.

CAPABILITY LEVEL	COUNT	SHARE OF INCLUDES
L3 (operational embedding / business integration)	50	50%
L4 (operating model / contested judgement / leadership)	28	28%
L2 (artefact production / process administration)	10	10%
L2-L3 (mixed mid-band)	6	6%
L3-L4 (mixed senior band)	6	6%
Other	1	1%

Capability level of the 101 included advertisements

L2 artefact production to L4 operating model and contested judgement. N/A: one include not level-coded.




Chart titles come from the 2026 coding run. The bands are the advertised signals for each framework level; Section 6 tests the L4 signal against System Shaping. See Section 3.

L3 is the dominant level across the included sample, at exactly half of all included ads. The roles in this band describe practitioners who embed privacy work into product, vendor, data, AI and operational decisions, rather than producing artefacts to satisfy a separate review process. Adding the L2-L3 and L3-L4 mid-bands, 62% of the included sample sits in or astride the L3 capability level.

L4 is the second-largest band at 28%. These ads describe roles with framework leadership, executive advisory scope, contested-judgement scope or delegated decision authority; Section 6 shows that none asks the role to change the organisation's own operating model. The L4 roles spread across Compliance / Privacy Governance (seven, at Head, Director or Senior Manager level), AI Governance Hybrid (six), and Legal Privacy Counsel, Operational Privacy and the statutory officers of DPO / Regulatory Interface (five each). The L4 distribution skews to public sector and large-enterprise employers.

L2 sits at 10%, with the bulk in junior compliance, analyst and project-management roles where privacy is one input alongside other compliance work. L1 ads are effectively absent from the included sample: junior generalists with rules-awareness scope tend to fall into the near-miss bucket rather than the included denominator.

The headline reading from this distribution is not that the market is hiring senior. It is that the market is hiring people who are expected to *do* privacy work in the business, not merely *describe* it. Section 6 analyses this distribution in detail, including what L3 dominance implies for the privacy capability gap this review's proposition names, and why an L3 code does not by itself show that the work being bought is embedded privacy work.


AI governance signal across the corpus

Across all classified rows, the AI governance signal distribution is:

AI GOVERNANCE SIGNAL	COUNT	SHARE OF ALL CLASSIFIED
None	98	59%
Strong (AI governance is central scope)	40	24%
Some (AI in scope, not central)	19	11%
Mentioned only (boilerplate or single reference)	10	6%

AI governance signal across all 167 classified advertisements

Strong or Some in roughly one in three ads; within the 101 includes: 28 Strong, 11 Some, 4 Mentioned, 58 None.



Roughly 35% of all classified ads carry meaningful AI governance signal (Strong or Some). Within the included sample of 101, the share is similar: 28 Strong, 11 Some, 4 Mentioned, 58 None. The 14 AI Governance Hybrid roles in the included sample are the most concentrated form of this pattern, but the signal extends well beyond the hybrid archetype. Roles in Legal Privacy Counsel (Xero Senior Privacy Counsel, Optus Senior Corporate Counsel - Privacy, CBA Senior Legal Counsel Privacy and Data, ANZ Senior Lawyer Privacy & Data), Compliance / Privacy Governance (Bankwest Senior Manager Specialist Advisory Privacy and Data, Itoh Director Digital/Data/Information Governance, Deputy GRC Program Manager) and Operational Privacy (QBE Senior Privacy Lead, Sportsbet Head of Privacy/Data Enablement/RAI) all carry explicit responsible AI scope as a secondary frame.

Section 10 analyses what the corpus reveals about where AI governance is being organised, and the three-way split (standalone / inside data leadership / absorbed into privacy) that is the report's clearest evidence that AI governance has not cohered as a discipline.

Sprint contribution to the corpus


The six sprints contributed unevenly to the final corpus.

SPRINT	RAW	INCLUDED	NEAR MISS	EXCLUDED	INCLUDE RATE
sprint-01	66	45	19	2	68%
sprint-02	47	29	13	5	62%
sprint-03	14	4	10	0	29%
sprint-04	14	4	8	2	29%
sprint-05	21	16	4	1	76%
sprint-06	5	4	0	1	80%
Total	167	101 (avg 60%)	55	11	

A reconciliation note on this table: the per-sprint rows sum to 102 included and 54 near-miss because one advertisement counted as an include at its originating sprint was later reclassified as a near miss during downstream QA, and the sprint rows are preserved as the historical record rather than restated. The final counts are 101 included, 55 near-miss, 11 excluded, 167 raw.

What each collection sprint contributed


Include rate above each sprint. Sprint rows are the historical record; the final total is 101 / 55 / 11.



The wide range of sprint-by-sprint include rates (29% to 76%) reflects the deliberate two-mode collection strategy described in Section 3. The earlier sprints used broad search queries and surfaced rich near-miss material. The fifth sprint used focused privacy-titled queries to push the

included sample over the methodological minimum, with a corresponding jump in include rate. Sprint 06 was a five-ad top-up that crossed the 100-included threshold.

The near-miss share is more stable than the include rate, provided it is read cumulatively. Sprint-by-sprint, the near-miss share of each batch ranges from none in the five-ad sprint 06 to 71% in sprint 03, with batch composition: the senior-skewed sprints 03 and 04 were near-miss-heavy, the focused sprint 05 was not. The cumulative share is the meaningful measure, and it moved only between 29% and 36% across the whole collection, closing at 33%. That cumulative stability across both collection modes supports the near-miss finding. The absorption share alone, the 43 roles that name privacy inside another function, moved more, from 20% to 28% as the senior-skewed sprints arrived, and closed at 26%.



Sample composition: what the corpus is not

Some patterns of the Australian privacy hiring market are under-represented or absent in this corpus, and they belong in front of the reader here rather than only in the methodology appendix.

- **Government-platform-only roles.** Some state and federal privacy roles are advertised exclusively on JobsWA, I Work for NSW, NESAs Jobs, or similar government job platforms. These do not all cross-post to LinkedIn, SEEK or Indeed. The corpus under-represents lower-band public sector privacy work. Where state-platform-only ads were captured (typically because they cross-posted), the corpus disproportionately surfaced senior roles.
- **Internal-only roles.** Privacy roles filled internally without external advertising do not appear in the corpus at all. I expect this is a significant share of senior privacy capability movement, particularly in financial services.

- **Recruiter-private mandates.** Some senior privacy roles are filled by executive search without public advertising. The corpus captures recruiter advertisements but not direct-relationship search mandates.
- **NZ and trans-Tasman roles.** A small number of NZ-based or trans-Tasman roles were captured where a listing surfaced under Australian location filters. These were retained in the corpus because Australia and New Zealand are comparable privacy regimes for capability-language purposes.

These limitations narrow but do not invalidate the corpus. The findings are about what is observable on the three job boards searched during the collection window, with the corpus shape that the two-mode collection strategy produced.

5. Role archetypes observed

How the archetypes were defined

Every included advertisement was assigned a primary archetype from a nine-category typology. Seven of those categories carry roles in the included sample. The eighth (Privacy-Adjacent Near Miss) holds the 55 near-misses analysed separately in Section 11. The ninth (Other / Mixed) carries no roles.

The archetypes are not job-family labels. They are descriptions of what the advertisement says the role's privacy capability *does*. A solicitor with privacy as a major practice area sits in Legal Privacy Counsel regardless of whether the title says "Lawyer" or "Counsel". A senior compliance officer with privacy as a substantive scope sits in Compliance / Privacy Governance regardless of whether the title mentions privacy. A statutory FOI / Privacy officer reporting to a Privacy Commissioner sits in DPO / Regulatory Interface regardless of whether the role is described as DPO or as Principal Information Officer.

The archetype was assigned conservatively. Where a role genuinely straddled two categories, a secondary archetype was also recorded; the role's primary classification then went to whichever scope the JD body emphasised.

This section walks through each archetype in descending count order. For each archetype the section gives a working definition, the corpus count and capability level profile, a short list of representative roles from the corpus, and a one-paragraph analytical hook pointing forward to the sections where the archetype's implications are taken up.

Compliance / Privacy Governance: 27 roles (27%)

What this archetype captures. Roles whose primary privacy capability requirement is producing, maintaining and reporting on a privacy compliance program. These ads emphasise policy and standard maintenance, breach and incident registers, training rollouts, control attestation, regulator-facing reporting, ISO 27701 or similar framework alignment, and the administration of privacy artefacts (PIAs, ROPAs, vendor questionnaires). Privacy advice is part of the role but framed as compliance advisory, not legal advisory or operational design.

Capability level profile. Mostly L3 (operational embedding, but in the compliance function rather than the business), with a meaningful L4 tail of Senior Manager / Head of Compliance roles, and an L2 cluster of analyst and officer roles.

Representative roles in the corpus:

- **Cochlear, Privacy Manager (JD-026, L3)**. Program stewardship across a global medical-device manufacturer; privacy policy maintenance, training, breach response and PIA coordination as core scope.
- **NSW Education Standards Authority, Manager, Privacy, Data and Governance (JD-143, L4)**. Senior management of a privacy and data governance program in a state statutory authority; combines privacy compliance with broader information governance scope.
- **Colonial First State, Manager Financial Crime and Privacy Compliance (JD-067, L3)**. Privacy compliance as one limb of a combined financial crime and privacy program; AML-adjacent framing, regulator-reporting orientation.
- **Bankwest, Senior Manager Specialist Advisory - Privacy and Data (JD-157, L4)**. Senior compliance advisory inside a major bank's privacy function; explicit responsible AI scope is the standout secondary frame.
- **ING, Privacy Compliance Manager (JD-011, L3)**. Privacy program management in a bank, with AI mentioned as part of the regulatory horizon but not central scope.
- **AFCX (via Ink Recruitment), Head of Governance & Assurance (JD-167, L4)**. Whole-of-organisation governance and assurance head with privacy as a substantive scope; the corpus's clearest example of privacy compliance absorbed into a broader governance executive role.

Why this archetype matters. This is the corpus's largest archetype, and the language patterns inside it are the most repetitive across employers. The roles describe what most readers would recognise as the "standard" Australian privacy program: policy, training, breach register, PIA workflow, vendor questionnaire, regulator engagement. Section 7 takes up what these ads over-index on, and what they leave under-described.

Legal Privacy Counsel: 22 roles (22%)

What this archetype captures. Roles whose primary privacy capability requirement is legal advice on privacy and related data protection matters. These ads sit in legal teams or law firms, frame privacy as a practice area rather than an operational discipline, and require either a current practising certificate or recent in-house counsel experience. The work described is advisory, contract-drafting, regulatory interpretation, and incident-response legal advice rather than artefact production or business embedding.

Capability level profile. Predominantly L3 (in-house counsel with substantive privacy load) and L4 (Senior Legal Counsel, Special Counsel, Senior Lawyer). AI governance signal is markedly stronger in this archetype than in any other except AI Governance Hybrid: many senior counsel roles now name responsible AI explicitly.

Representative roles in the corpus:

- **Xero, Senior Privacy Counsel (JD-142, L3-L4)**. Pure privacy counsel role in a global software business; explicit Strong AI governance signal in scope.
- **CBA, Senior Legal Counsel, Privacy and Data (JD-097, L4)**. In-house bank counsel with privacy and data as a defined remit; Strong AI signal.
- **ANZ, Senior Lawyer, Privacy & Data (JD-063, L4)**. Senior in-house bank counsel role; framed as privacy-and-data rather than privacy-and-cyber. Strong AI signal.
- **Optus, Senior Corporate Counsel - Privacy (JD-139, L4)**. Senior in-house telco counsel role; Strong AI signal in scope language.
- **Canva, Senior Legal Counsel, Business Operations (JD-104, L3-L4)**. Privacy as one of several practice areas; AI scope explicit. A good worked example of how privacy is bundled with adjacent practice areas at scale.
- **Proximity, Privacy Lawyers (Mid Level) (JD-106, L3)**. Privacy specialist boutique work; a counterpoint to the in-house bank pattern.
- **Law Staff, TMT Lawyer (Cyber, Privacy & AI Focus) (JD-054, L3)**. Indicative of how recruiters are packaging privacy with TMT and AI as a single specialty.

Why this archetype matters. Combined with the 27 Compliance / Privacy Governance roles, Legal Privacy Counsel takes the proportion of the included sample whose primary privacy capability is advisory (legal or compliance) to 49%. The Senior Legal Counsel cluster also concentrates the most explicit AI governance language in the corpus, which feeds the analysis in Section 10 of where AI governance is being housed.

Operational Privacy: 17 roles (17%)

What this archetype captures. Roles whose primary privacy capability requirement is embedding privacy into the operations of the business. These ads emphasise vendor onboarding, product design partnership, data flow uplift, DPIA execution as a business-decision input rather than an artefact for legal sign-off, and ongoing relationship with product, engineering, marketing, customer and data teams. The language reads as "do privacy work with the business" rather than "advise the business on privacy".

Capability level profile. Heavily L3 (the band where embedding actually happens), with a strong L4 tail of Privacy Lead, Head of Privacy, and Principal Director roles that describe operational leadership rather than compliance leadership.

Representative roles in the corpus:

- **QBE Insurance, Senior Privacy Lead (JD-058, L4)**. 9-month fixed-term contract for senior privacy operational uplift; Strong AI signal. The standout L4 operational embedding role in the corpus.

- **Woolworths Group, Privacy Lead, Privacy and Customer Trust (JD-068, L3-L4)**. Embedded privacy leadership with a customer-trust framing rather than a compliance framing.
- **ASX, Data Privacy Specialist (JD-013, L3)**. Specialist privacy embedding in capital markets infrastructure; Some AI signal.
- **University of New South Wales, Data Privacy Lead (JD-034, L4)**. Higher-education embedded privacy leadership; statutory and policy interaction but described as operational, not statutory officer.
- **Sportsbet, Head of Privacy, Data Enablement and RAI (JD-126, L4)**. Bundled privacy, data enablement and responsible AI head; the corpus's clearest example of privacy leadership absorbing responsible AI as defined scope.
- **Picklebet, Data Privacy Manager (JD-033, L3)**. Smaller-employer embedded privacy lead; useful as the contrast point to the large-enterprise pattern.
- **Endeavour Group, Senior Privacy Officer (JD-047, L3)**. Retail operational embedding; Some AI signal.
- **Correlate Resources (recruiter), Business Analyst - Privacy, Data Flows & Transparency Uplift (JD-050, L2-L3)**. Project-shaped embedded privacy work; the most explicit "uplift" framing in the corpus.

Why this archetype matters. This is the archetype this review's proposition predicts the market needs more of. At 17% of the included sample it is meaningful but clearly secondary to the compliance and legal framings. Section 9 profiles what these roles actually describe doing, and where the L3 dominance translates (or fails to translate) into the operating-model outcomes the report's later sections argue for.

AI Governance Hybrid: 14 roles (14%)

What this archetype captures. Roles whose primary capability requirement is the governance of AI and adjacent data activity, with privacy named as a major scope component. These ads do not separate AI governance from privacy as cleanly as either of the standalone archetypes would suggest. Some are clearly led by privacy expertise that has absorbed AI; others are led by data and AI expertise that has absorbed privacy. The defining feature is that the role's design treats privacy and AI governance as one combined remit rather than two disciplines.

Capability level profile. Concentrated at L3 (specialist embedding) and L4 (head-of-function leadership). Almost every role in this archetype carries Strong AI signal by construction (13 of the 14; the exception is the RBA role at Some).

Representative roles in the corpus:

- **Profectus, AI Governance Advisor (JD-028, L3-L4)**. Pure AI governance advisory role with privacy as a named scope; representative of the standalone-advisor pattern.

- **Australian Financial Complaints Authority, Senior Manager Data & AI Governance (JD-082, L4)**. Combined data and AI governance leadership in a financial sector statutory body.
- **MUFG Pension & Market Services, Head of AI & Data Privacy Risk (JD-114, L4)**. AI and data privacy bundled as one risk leadership remit.
- **Slade Group, AI Governance Lead (JD-152, L4)**. Standalone AI governance leadership; signals where the market is sourcing dedicated AI governance heads.
- **Helia, Data and AI Governance Leader (JD-147, L4)**. Data and AI governance head; placed at the boundary of this archetype and standalone AI governance.
- **University of Sydney, Senior Privacy Specialist AI (JD-102, L3)**. Privacy specialist with AI scope as a defined deepening.
- **ctrl:cyber, Principal Consultant - Privacy, Data and AI Governance (JD-065, L4)**. Consultancy-led AI governance practice; combines three scopes into one role design.
- **Australian Criminal Intelligence Commission, EL1 Data Governance Specialist (JD-111, L3)**. Public-sector AI and data governance specialist; the corpus's clearest example of AI governance being absorbed into a public-sector data role.

Why this archetype matters. The hybrid archetype is the most analytically informative single category in the corpus. The three patterns inside it (privacy-led, data-led, standalone) are the evidence base for the report's claim in Section 10 that AI governance has not settled as a discipline in the Australian market. Read this archetype alongside the standalone AI governance roles in the near-miss sample (e.g. IAG, CommBank, Zip Co) and the AI-absorbed-into-privacy roles in Operational Privacy and Legal Privacy Counsel.

DPO / Regulatory Interface: 13 roles (13%)

What this archetype captures. Roles whose primary capability requirement is acting as the statutory or designated privacy officer for an organisation, including formal regulator engagement, statutory information access decisions (FOI, GIPA, RTI, PRIS), and decisions taken under delegated authority. These ads sit overwhelmingly in the public sector and in regulated industries with statutory privacy obligations beyond the Privacy Act 1988.

Capability level profile. Skews L3 and L4. L1 and L2 are effectively absent: statutory officer roles require capability against legislation and delegated decision authority, both of which assume substantive seniority.

Representative roles in the corpus:

- **State Solicitors Office (WA), Privacy and Responsible Information Sharing (PRIS) Manager (JD-093, L4)**. PRIS Act 2024 statutory role; the corpus's cleanest example of a state-level reform driving a new statutory officer profile.

- **Office of Information Commissioner, Investigations Officer (Privacy) & Senior Investigations Officer (Privacy) (JD-069, L3)**. Regulator-side privacy investigation roles; the reverse of the in-house DPO pattern but coded into this archetype because the work is regulator-interface.
- **Insurance Commission of WA, Principal Information Officer (JD-052, L4)**. WA statutory information officer; combined privacy and information-management scope.
- **Department of Education Queensland, Principal Privacy Officer - Privacy & Safer Technologies (JD-156, L4)**. State education department statutory privacy officer with technology scope.
- **University of Queensland, Right to Information & Privacy Manager (JD-042, L3)**. Combined RTI and privacy manager in a higher-education statutory context.
- **Johnson & Johnson MedTech, Dir, Priv & Data Protection Officer APAC (JD-123, L4)**. Private-sector designated DPO role with APAC regional scope; one of the few corporate DPO roles in the corpus.
- **VicTrack, Senior Privacy & Governance Officer (JD-163, L3)**. Victorian statutory body privacy officer with OVIC-facing scope.
- **Brisbane North PHN, Lead, Data Governance and Privacy (JD-017, L4)**. Primary health network privacy and data governance lead; combines statutory privacy with broader data governance.

Why this archetype matters. The DPO / Regulatory Interface roles concentrate in two specific patterns: state-level statutory reform (WA PRIS Act 2024, Qld IP Act, Vic PDP Act, NSW PPIP Act reforms) creating new officer roles, and higher-education / health-network organisations combining FOI / RTI / PRIS responsibilities with privacy under a single officer. Section 13 returns to the state-reform pattern as the driver of a distinct statutory officer career track.

Data / Technology Privacy: 5 roles (5%)

What this archetype captures. Roles whose primary capability requirement is the technical implementation of privacy controls inside data infrastructure: privacy engineering, data protection technology deployment, SAP / Salesforce / similar platform privacy configuration, technical PETs work, and privacy-by-design code review. Distinct from Operational Privacy because the work is described in engineering and technology terms, not business-embedding terms.

Capability level profile. L3, L2-L3 and one L2, with no L4 in the included sample. This is partly an artefact of how the work is advertised: senior privacy engineering roles tend to be advertised under broader data engineering or security engineering titles and therefore fall into the near-miss sample.

Representative roles in the corpus:

- **BHP, Principal Privacy Engineer (JD-012, L3)**. The corpus's clearest example of a privacy engineering role in a major Australian enterprise.

- **Avanade, Data Protection Manager (JD-037, L3)**. Consultancy data protection role with technical platform delivery scope; Some AI signal.
- **AustralianSuper, Privacy & Data Protection Technology Project Manager (JD-064, L3)**. Technology-side privacy project management in a major super fund.
- **EPI-USE Labs, Lead Consultant SAP Data Privacy (JD-091, L2-L3)**. Platform-specific privacy consulting; SAP-focused.
- **ANZ Banking Group, Privacy Risk Business Analyst - 12M FTC (JD-070, L2)**. Privacy risk and data project analyst.

Why this archetype matters. Five roles is a small absolute number. The report does not lean heavily on this archetype quantitatively, but it surfaces a point that matters beyond its size: most privacy engineering capability in the Australian market is not advertised as "privacy engineering". It is advertised as data engineering, security engineering, or data governance, with privacy as a small component. Section 8 returns to this when discussing privacy engineering as a named function.

Risk / GRC Privacy: 3 roles (3%)

What this archetype captures. Roles whose primary capability requirement is the independent challenge of privacy controls and operations from a second-line risk or assurance position. These ads describe risk-rated assessment of privacy, independent reporting lines, board-facing risk reporting on privacy issues, and a separation between the privacy practitioners who design and operate controls (Line 1) and the privacy assurance professionals who test them (Line 2).

Capability level profile. All three are L3.

Representative roles in the corpus:

- **HSBC Australia, Data Risk and Technology & Cyber Risk Specialist (JD-036, L3)**. Privacy as one input into a broader technology and data risk role with some second-line framing.
- **Protiviti Australia, Technology Risk and Resilience Manager (JD-016, L3)**. Consultancy second-line risk role with privacy as part of a broader technology risk scope.
- **Plenti, Risk & Compliance Manager (JD-081, L3)**. Combined risk and compliance role at a fintech lender; privacy named as a substantive scope item.

Why this archetype matters. Three of 101 included roles, across six sprints and deliberate query effort to surface them. Even taking those three at face value, the framing is not the clean second-line assurance pattern this archetype was set up to capture: each role bundles privacy into broader risk scope rather than treating privacy as a discrete risk discipline with its own assurance approach. Section 8 returns to this as the single strongest piece of corpus evidence that the three-lines-of-defence model many privacy programs assume on paper is not visible in advertised privacy demand.

What the archetype distribution shows

Read across all seven archetypes, the included sample concentrates in three places: compliance program stewardship, legal advisory, and the various forms of operational, hybrid AI, and statutory work. It thins markedly in the technical and assurance directions.

A reader expecting the Australian privacy job market to map cleanly onto a three-lines-of-defence picture would be wrong. A reader expecting it to be dominated by Privacy Officer / Privacy Manager titles would also be wrong: those titles exist but are outnumbered by Compliance, Legal Counsel, AI Governance, and DPO framings combined.

What the corpus describes is a market where privacy capability is mostly purchased as compliance program labour or as legal advisory labour, with operational embedding as a meaningful but secondary purchase, and with both independent assurance and dedicated privacy engineering essentially absent from how organisations advertise their privacy capability needs. The sections that follow take up what that distribution implies for how privacy operating models are being shaped, and where the gap between the work organisations advertise for and the work they need shows up most clearly.

6. Capability level analysis

The framework, briefly

A note on these bands before the numbers. The bands are the signals an advertisement gives for each level of the Privacy Capability Framework, coded in May and June 2026 and kept unchanged with the corpus. Where the report means a framework level it writes the name out in full. The L4 signal is the one that needs testing, and this section tests it below. Section 3 explains the mapping; Section 16.4.2 sets it out in full.

Each included advertisement was assigned a capability level from L1 to L4. The four levels describe what the role's substantive privacy capability requirement is, not the seniority of the title:

- **L1: Rules awareness.** Knows that privacy law exists, can identify issues, supports privacy work without owning it.
- **L2: Artefact production.** Produces and maintains privacy artefacts: PIAs, ROPAs, registers, policies, training material, vendor questionnaires.
- **L3: Operational embedding.** Embeds privacy into the operations of the business: vendor onboarding decisions, product design partnership, data flow uplift, DPIA execution as a business-decision input rather than as an artefact for sign-off.
- **L4: Operating model and contested judgement.** Designs privacy operating models, makes contested judgement calls under regulatory ambiguity, leads frameworks across multiple functions, exercises delegated authority.

Two intermediate bands (L2–L3 and L3–L4) were used where the role's described work genuinely straddled the boundary. They are reported separately rather than rolled into the higher level, to keep the distribution honest.

A role's capability level was set by the work described in the JD body, not by title or seniority. A Privacy Officer role that describes administering a PIA register sits at L2 even if the salary is senior. A Compliance Manager role that describes contested judgement on novel AI governance questions sits at L4 even if the title sounds mid-band.

The distribution

CAPABILITY LEVEL	COUNT	SHARE OF INCLUDES
L3	50	50%
L4	28	28%

CAPABILITY LEVEL	COUNT	SHARE OF INCLUDES
L2	10	10%
L2-L3	6	6%
L3-L4	6	6%
Other (N/A)	1	1%

The single N/A row is JD-040 (Reserve Bank of Australia, Senior Manager, Information and Data Governance), which was elevated from near miss to Include at sprint-01 manual review. A capability level was not coded for it before the corpus was frozen, so it is reported as N/A rather than retrofitted.

Three things to notice before the analysis.

L3 sits at exactly half the included sample. Adding the two mid-bands (L2-L3 and L3-L4), 62% of all included roles describe work at or astride the operational-embedding level. This is the largest single concentration in the corpus.

L4 is the second-largest band at 28%. More than a quarter of the included sample describes roles with framework leadership, delegated authority or contested-judgement scope. That share reflects how many of the advertised roles are statutory officer, head-of-function and senior in-house counsel posts.

L1 is effectively absent from the included sample. Junior generalist roles with rules-awareness scope tend to fall into the near-miss bucket rather than into the included denominator. The corpus does not give a defensible picture of L1 privacy capability supply in the market.

Reading the L3 dominance carefully

The headline number, half of all included roles at L3 and 62% at or astride L3, would, if read at face value, suggest the Australian market is hiring a lot of embedded privacy capability. It is hiring a lot of *L3-coded work*. The two are not the same thing, and the distinction is the most important reading-error to avoid in this section.

L3 is a level descriptor of *what the role does in the business*, not of *whether the role does embedded privacy work*. A Compliance Manager whose role description says "partner with business stakeholders to embed privacy controls into vendor onboarding and product launches" is coded L3 because the work language is embedding. A Senior Privacy Counsel whose role description says "advise product teams on privacy implications during design sprints" is coded L3 because the work language is embedded advisory. A Senior Privacy Lead in the Operational Privacy archetype whose role description says "uplift privacy operating practice across data, marketing and AI workflows" is also coded L3, but for substantively different work.

All three are L3. Only one is embedded privacy work in the sense this review's proposition names. The L3 share is therefore best read as a measure of how organisations *describe* privacy capability: the language has shifted from rules-and-artefacts toward embedding-with-the-business, without yet telling the reader how much of that embedding is being purchased as embedded privacy capability rather than as compliance-embedding or legal-embedding work.

Cross-tabulating L3 against the seven archetypes makes this concrete:

- In **Operational Privacy** (17 roles), L3 dominance describes the embedding work the archetype is built around.
- In **Compliance / Privacy Governance** (27 roles), L3 dominance describes compliance officers embedding control language into vendor and process workflows.
- In **Legal Privacy Counsel** (22 roles), L3 dominance describes in-house counsel embedding legal review into product and commercial pipelines.
- In **AI Governance Hybrid** (14 roles), L3 dominance describes specialist embedding into AI-adjacent data activity.
- In **DPO / Regulatory Interface** (13 roles), L3 dominance describes statutory officer work that combines decision-making authority with embedding into agency operations.




Chart titles come from the 2026 coding run. The bands are the advertised signals for each framework level; Section 6 tests the L4 signal against System Shaping. See Section 3.

The L3 rating is the same. The work it codes is not. Section 9 takes up what the *Operational Privacy* L3 sub-population actually describes, and where its size (17 roles, not 50) leaves the embedded

privacy capability supply in the Australian market.

What the L4 share concentrates

L4 sits at 28%, 28 roles. The distribution across archetypes is uneven:

- **Compliance / Privacy Governance:** 7 L4 roles (Head of Compliance, NSW Education Standards Authority Manager Privacy/Data/Governance, Senior Manager Specialist Advisory at Bankwest, Director Digital/Data/Information Governance at Itoh, AFCX Head of Governance & Assurance, Senior Advisor Governance at DCS, and a Data Security & Governance Manager advertised without the employer named).
- **Legal Privacy Counsel:** 5 L4 roles (ANZ Senior Lawyer, Carlyle Kingswood Senior Legal Counsel Privacy, CBA Senior Legal Counsel Privacy and Data, Dovetail Senior Legal Counsel Privacy & Data, Optus Senior Corporate Counsel Privacy).
- **Operational Privacy:** 5 L4 roles (Victoria Legal Aid Data Protection Lead, UNSW Data Privacy Lead, Randstad Principal/Director DPIA, QBE Senior Privacy Lead, Sportsbet Head of Privacy/Data Enablement/RAI).
- **AI Governance Hybrid:** 6 L4 roles (Protiviti Data & AI Governance Senior Manager, ctrl:cyber Principal Consultant, AFCA Senior Manager Data & AI Governance, MUFG Head of AI & Data Privacy Risk, Helia Data and AI Governance Leader, Slade Group AI Governance Lead).
- **DPO / Regulatory Interface:** 5 L4 roles (Brisbane North PHN Lead Data Governance and Privacy, ICWA Principal Information Officer, State Solicitors Office PRIS Manager, J&J MedTech DPO APAC, Dept of Education Queensland Principal Privacy Officer).
- **Data / Technology Privacy:** 0 L4 roles.
- **Risk / GRC Privacy:** 0 L4 roles.

Two readings follow.

First, L4 concentrates in **statutory officer roles and senior compliance / legal leadership**. The DPO archetype has the cleanest L4 cluster because the work *requires* delegated authority. The Compliance and Legal Counsel L4 roles are the head-of-function and senior-counsel posts that operate at framework-design and contested-judgement scope.

Second, L4 is **absent from the technical and assurance archetypes**. Data / Technology Privacy has no L4 roles in the included sample. Risk / GRC Privacy has no L4 roles in the included sample (and only three roles in total). The L4 work the market is hiring for is largely compliance, legal, AI governance, Operational Privacy and statutory, not privacy engineering leadership and not independent privacy assurance leadership. Section 8 returns to this as part of the broader pattern of what the corpus under-describes.

What the L4 band is actually holding

The 28 L4 roles are where the signal is weakest, and testing it changes what the number means.

As Section 3 sets out, the L4 band picks up four signals: operating-model design, contested judgement, senior leadership, and authority signals such as delegated authority or an executive or board reporting line. Under the practitioner framework only the first is System Shaping. Contested judgement is Risk-Disciplined Judgement, which is Level 3. A reporting line is a position rather than a capability. So the band cannot be read as 28 roles that require someone who can change how privacy decisions get made.

To find out how many do, I re-read all 28 L4 roles and all 6 L3–L4 roles against a single test: does the advertisement ask the role to shape the recurring conditions under which privacy decisions are made? Evidence had to show something like redesigning how privacy decisions are made, defining decision rights or escalation arrangements, shaping the enterprise operating model, changing cross-functional governance routines, building and distributing capability in other people, or creating a repeatable institutional arrangement instead of handling each difficult matter personally. A senior title, a team to lead, a board to report to, wide stakeholder engagement or difficult advice were not enough on their own. The original classifications were left untouched; this is a second reading of the same evidence.

READING AGAINST THE PRACTITIONER FRAMEWORK	ROLES	SHARE OF THE 34	SHARE OF ALL 101 INCLUDES
System Shaping evidenced	0	0%	0%
Governance construction inside fixed conditions	9	26%	9%
Mixed	12	35%	12%
Risk-Disciplined Judgement only	13	38%	13%

No advertisement in the corpus asks anyone to change the conditions under which the organisation’s privacy decisions are made.

Nine come closest, and they are worth looking at because of how they fall short rather than how near they get. Slade Group’s AI Governance Lead is asked to develop an operating model “including roles, responsibilities, decision rights, and escalation pathways”, which is the criterion almost word for word. Johnson & Johnson MedTech’s regional DPO is asked to “embed privacy capabilities in named privacy stewards from relevant functions”. Others describe designing frameworks, owning a governance lifecycle or establishing an enterprise approach.

What none of them describes is an organisation putting its own operating model on the table. In each case something is being built inside conditions the organisation has already settled: a framework, a lifecycle, a standard, a steward model. The appointee is expected to fit the arrangement, not to alter it, and a framework can be designed to a high standard without anything about how decisions get made changing at all. The nine fall into three groups.

Greenfields governance builds (Slade, Helia, AFCA). AI governance has no incumbent blueprint, so the role is a first build and the advertisement has to spell out the architecture because there is none to inherit. That produces the most Level 4-looking language in the corpus for reasons that have nothing to do with privacy. Shaping presupposes an existing system worked on through mandate, influence and sponsorship; a blank sheet is a different activity.

Framework builds inside a fixed model (UNSW, J&J MedTech, MUFG, Itoh). A framework is established, a steward network populated, a risk framework owned inside three lines already drawn. The model is the given.

Consultancy (Protiviti, ctrl:cyber). The capability is advertised; the authority to use it belongs to a client and is never evidenced in the advertisement.

The thirteen roles that read as Risk-Disciplined Judgement are not lesser roles, and the reclassification is not a criticism of them. They are largely senior counsel and principal officers doing exactly what Section 7 describes: giving the organisation a defensible position on hard questions. That is the capability the framework names at Level 3, and the market is buying a good deal of it. What those advertisements do not ask for is a change to the conditions that keep producing the hard questions.

I also scanned all 50 L3 roles for signals the original coding might have placed too low. One used operating-model language, and it describes implementing and embedding an operating model that already exists rather than shaping one. The L3 band is not concealing System Shaping; the ambiguity sat entirely inside L4, and it resolves the same way.

No advertisement in the corpus carries a Chief Privacy Officer or equivalent mandate. There are no CPO roles among the 101 includes. Four roles carry a "Head of" title and only one of those is a Head of Privacy, where privacy is bundled with data enablement and responsible AI. That is the same finding twice over: nothing in the corpus asks for the authority to change the operating model, and nothing asks for the work that authority would be used on.

There is a second test these roles have to pass, and it is the one that decides the Slade case. Shaping requires a position from which shaping is possible. An advertisement can describe the work and place the role nowhere near the arrangements it would be shaping.

The corpus is thin here, because seniority and reporting line were not systematically coded, but the JD bodies say enough. UNSW's Data Privacy Lead is "reporting to the Director of Data Privacy and Security, a dedicated specialist role with no direct reports". AFCA's Senior Manager reports to the Head of Data & Analytics, which is data leadership rather than privacy. Slade's AI Governance Lead states no reporting line at all, carries "Lead" rather than "Head of", and is recruiter-advertised with the employer hidden. Only Itoh's Director, reporting to a Deputy Executive Director, sits near the level at which operating models are actually changed, and that role spans six domains of which privacy is one.

So the advertisement that comes closest to naming the work does not indicate a position from which to do it, and the one placed highest is not principally about privacy. That is the same standard

the original coding applied when it ruled that a senior title alone is not enough for L4, turned around: a senior scope is not enough either.

Two boundaries belong on this reading. **It is a finding about advertisements, not about organisations.** The capability may well exist inside organisations already, in roles filled years ago, by executive search, by internal promotion, or held by someone whose title says nothing about privacy. Nothing here establishes what an organisation holds, needs, or understands about itself. **And a dissenting reading is available.** Someone who treats the shaping scope as sufficient, and the silence about placement as uninformative rather than as absence of evidence, will count one where this report counts none.

The mid-bands

The L2–L3 band (6 roles) and the L3–L4 band (6 roles) carry roles whose described work genuinely straddles a level boundary. These mid-band assignments are not a compromise position. They were used where the JD body described both lower-band and higher-band work substantively, and where collapsing the role into a single level would have lost information.

The L2–L3 band includes roles where artefact production sits alongside genuine embedding work: Vanguard Data Privacy Analyst, City of Parramatta Governance Officer, Mura Connect Senior Privacy Officer, EPI-USE Labs Lead Consultant SAP Data Privacy, Correlate Resources Business Analyst Privacy/Data Flows/Transparency Uplift, Talenza Senior Privacy Analyst. These are the analyst and senior-officer roles where the work is both "produce and maintain artefacts" *and* "work with the business to embed". A typical pattern for mid-tier privacy hiring in mid-size employers.

The L3–L4 band includes roles where embedding work sits alongside framework-leadership or contested-judgement scope: Canva Senior Legal Counsel Business Operations, CatholicCare Compliance and Privacy Officer, Profectus AI Governance Advisor, SS&C Technologies Data Governance and Protection Officer, Woolworths Group Privacy Lead Privacy and Customer Trust, Xero Senior Privacy Counsel. These are the senior privacy roles where the work is both "do embedded privacy work" *and* "lead frameworks and contested decisions". A typical pattern for the head-of-privacy and senior privacy counsel hiring in large enterprises.

The two mid-bands matter because they evidence where the market is asking for *more than one capability level in the same role*. That has implications for the operating model argument in Section 14: the corpus does not always cleanly separate doing-privacy-work from designing-how-privacy-work-gets-done, even at the senior end.

L1 and L2: what the absence says

L1 is absent from the included sample. L2 sits at 10%, ten roles. Combined, the bottom-band privacy roles make up around one in ten of the included corpus.

Two readings of that low share:

First, **junior privacy work is rarely advertised as privacy work**. Where it appears in the corpus at all, it tends to sit inside generalist compliance, records or GRC roles that fall into the near-miss sample, which carries no capability coding. The included L1/L2 share therefore says little about how much junior privacy work the market employs.

Second, **the L2 roles that are advertised are concentrated in compliance and data-technology archetypes**, not in privacy-titled archetypes. The ten L2 roles in the included sample sit predominantly in Compliance / Privacy Governance (analyst, officer, program administration) and Data / Technology Privacy (project-level technical privacy work). Operational Privacy and AI Governance Hybrid have very few L2 roles. The pattern suggests that where organisations advertise junior-band privacy work, they advertise it as compliance support or technical project support, not as embedded privacy support.

What the distribution implies

Four claims from this section that the report's argument sections return to:

1. **The L3 dominance is real but reflects how privacy is described, not what privacy capability is being purchased.** The 50% L3 share spans operational, compliance, legal, AI governance and statutory work: substantively different jobs that all read as L3 in the JD language.
2. **L4 concentrates in compliance / legal / statutory / AI governance leadership and is absent from privacy engineering and privacy assurance leadership.** This shapes where the market is and is not hiring for senior framework and judgement work.
3. **Read against the practitioner framework, no advertisement asks for System Shaping at all.** Nine come close and each builds governance inside conditions the organisation has already settled; thirteen of the L4 and L3–L4 roles read as Risk-Disciplined Judgement, and twelve are mixed. The market is buying difficult judgement far more readily than it is buying the capability to change the conditions that keep producing difficult judgements, and it carries no Chief Privacy Officer mandate anywhere.
4. **L1 and L2 absence in the included sample is partly a measurement artefact and partly a real signal.** The artefact: junior privacy work is advertised under non-privacy titles. The signal: where junior privacy work *is* advertised under privacy titles, it is overwhelmingly compliance and technical project support, not embedded privacy support.

The largest single analytical move the report makes from the capability level distribution is in Section 9, where the 17 Operational Privacy roles are profiled to show what L3 operational-privacy-embedding actually describes in the JD bodies, and how that picture sits against the 50 L3 roles overall. The level dominates the corpus; the work it codes does not.

7. What job ads over-index on

What "over-indexing" means here

"Over-indexing" in this report does not mean *too much*. It means *disproportionately concentrated*. A pattern over-indexes when it appears in the corpus at higher frequency than would be expected if the corpus reflected the full operating-model picture this review's proposition names as needed.

Four patterns over-index in the corpus. Each is described below with the evidence anchor in the corpus and the analytical reading that follows from it.

The four patterns are interrelated. The first two, compliance program stewardship and legal advisory framing, account for 49% of the included sample by archetype. The other two, framework alignment language and generic stakeholder-engagement language, cut across all archetypes and shape how privacy capability is described inside them.

Pattern 1: Compliance program stewardship as the default frame


Evidence in the corpus. The Compliance / Privacy Governance archetype carries 27 of the 101 included roles, the single largest archetype. The work language across these 27 roles is highly repetitive: policy and standard maintenance, breach and incident registers, training rollout, control attestation, regulator-facing reporting, ISO 27701 or similar framework alignment, the administration of privacy artefacts (PIAs, ROPAs, vendor questionnaires).

Worked examples:

- **Cochlear, Privacy Manager (JD-026, L3).** Program stewardship across a global medical-device manufacturer. The JD body emphasises maintenance of the privacy policy framework, training delivery, breach response procedure, and PIA workflow coordination.
- **NSW Education Standards Authority, Manager, Privacy, Data and Governance (JD-143, L4).** Senior state-statutory program management combining privacy with broader information governance scope. Framework alignment, regulator reporting, and policy ownership across multiple data classifications.
- **Australian College of Nursing, Compliance and Privacy Officer (JD-162, L2).** A junior compliance-and-privacy role where the work language is artefact maintenance, register management, and training rollout.
- **Westpac Group, Senior Manager (JD-146, L3).** Senior compliance manager in a major bank with privacy as a defined scope; framework alignment and regulator-reporting work dominate the JD body.

Privacy purchased as compliance or legal advisory work

Primary archetype of the 101 included advertisements.



What this over-indexes against. The compliance program stewardship frame describes privacy as a program to be built and run, with artefacts (PIAs, ROPAs, training, registers) as the unit of work and adherence to a framework as the success measure. It does not describe privacy as an embedded capacity to make contested judgement calls at the moment a decision is being made.

This is not a critique of the roles themselves. Compliance program stewardship is a necessary function, and the people in these roles are doing real work. The over-index claim is about *concentration*: 27% of all advertised privacy capability is being purchased as compliance program stewardship, with very little contested-judgement scope built into the role design.

A privacy program that is well-stewarded but does not have embedded privacy capability in the business is the program this review's proposition describes as overbuilt on compliance and underbuilt on operations. The corpus shows the labour-market shape of that pattern: the compliance roles get advertised, the contested-judgement roles often do not.

Pattern 2: Legal advisory framing for senior privacy work

Evidence in the corpus. The Legal Privacy Counsel archetype carries 22 of the 101 included roles. Seven of its 22 roles sit in the L4 or L3-L4 bands (5 L4 and 2 L3-L4), with the remaining 15 at L3. Senior privacy work in the corpus is more often packaged as senior counsel work than as senior embedded privacy work.

Worked examples:

- **CBA, Senior Legal Counsel, Privacy and Data (JD-097, L4).** Senior in-house bank counsel with privacy and data as defined remit; Strong AI signal. The JD body frames the work as legal advisory across the bank's privacy and data agenda.
- **Optus, Senior Corporate Counsel - Privacy (JD-139, L4).** Senior in-house telco counsel role; the JD frames the work as legal advisory rather than as embedded privacy leadership.
- **Xero, Senior Privacy Counsel (JD-142, L3-L4).** Pure privacy counsel role in a global software business; Strong AI signal. The mid-band coding reflects that the role straddles legal advisory

and operating-model leadership.

- **ANZ, Senior Lawyer, Privacy & Data (JD-063, L4)**. Senior in-house bank counsel; framed as privacy-and-data rather than privacy-and-cyber. Strong AI signal.

What this over-indexes against. Senior privacy work in this corpus is more often packaged as legal counsel than as operating-model leadership. The two are not the same. Legal counsel work is advisory: it tells the organisation what the law requires and what risk it should accept. Operating-model leadership is different in kind: it designs how privacy decisions get made inside the organisation, who has authority to make them, where contested judgement is escalated, and how the operating model accommodates new regulatory and technological pressure points.

A market that hires senior privacy capability primarily as legal counsel is buying advisory firepower at the top of the privacy function and leaving operating-model design as an implicit secondary scope. Where the advisory work is excellent, the organisation gets excellent advice. Where the operating model is itself the problem, advisory work is the wrong instrument for the job.

The over-index claim here is not that legal counsel roles are over-filled. It is that the framing absorbs senior privacy leadership into a primarily advisory function, which leaves the operating-model design work either undescribed in the advertisement or pushed into a separate (and largely unhired-for) capability.

Pattern 3: Framework alignment as the success metric

Evidence in the corpus. Across the Compliance / Privacy Governance, DPO / Regulatory Interface and AI Governance Hybrid archetypes, JD bodies regularly name external frameworks and reference points as the orientation for privacy work. The reference points named most often across all 167 JD bodies are Privacy by Design (30 ads), the GDPR (24), IAPP certifications as required or preferred credentials (22), the Australian Privacy Principles (17), and ISO 27001 (14). In AI-scoped roles, ISO/IEC 42001 and the NIST AI Risk Management Framework (7 ads each) and the NSW AI Assurance Framework (5) are the rising references. The two privacy-specific management frameworks are named far less often than their profile in privacy program literature would suggest: ISO 27701 appears in three ads and the NIST Privacy Framework in one. State-specific statutory frameworks (WA PRIS Act 2024, Vic PDP Act, Qld IP Act and RTI Act, GIPA) anchor the DPO roles in their respective jurisdictions.

The language pattern is: privacy capability is good when it is *aligned* to a framework. The framework is the external reference point. The role's job is to bring the organisation's privacy posture into alignment.

What this over-indexes against. Framework alignment is a measurable, observable property. It produces artefacts that can be inspected: gap assessments, alignment scorecards, control

inventories, regulator-readiness reports. It is consequently a tempting success metric, because it gives the privacy program a defensible reporting line into the executive and the board.

Framework alignment is not, however, equivalent to embedded privacy capability. An organisation can be well-aligned to ISO 27701 and still have product teams who do not know who to call when a new AI tool needs a privacy assessment. An organisation can score highly on its OAIC self-assessment and still have a marketing function that purchases customer-targeting tools without a meaningful privacy review. The frameworks are useful as scaffolding and as a common language for governance. They are not, in themselves, the operational capability this review's proposition names as the under-built layer.

The corpus over-indexes on framework alignment because frameworks are what privacy programs know how to describe. Operational capability is harder to describe; it is consequently described less.

Pattern 4: Generic stakeholder-engagement and strategic-advisory language

Evidence in the corpus. A meaningful share of mid-to-senior roles across all archetypes carry framing language that is high-level, generic, and difficult to anchor to specific work. These are verbatim examples from JD bodies across five different archetypes:

- "Partner with stakeholders to strengthen internal controls, supporting compliance documentation and governance processes" (Entain, Compliance & Privacy Analyst, JD-007).
- "Partner with business, legal, and technology teams to embed Privacy by Design and responsible data use across initiatives" (Breville, Privacy Specialist, JD-010).
- "Partner with key stakeholders to integrate differential privacy techniques" and "manage risk and drive strategic and operational outcomes" (BHP, Principal Privacy Engineer, JD-012).
- "You'll be the trusted advisor supporting our Data Protection Officer" (ING, Privacy Compliance Manager, JD-011).
- "You are a trusted advisor with strong expertise in privacy, data governance and risk management" (Brisbane North PHN, Lead - Data Governance and Privacy, JD-017).
- "Build and develop relationships with key stakeholders across the Group to influence and implement effective privacy and data ethics capabilities" and "drive privacy and responsible data use uplift" (Woolworths Group, Privacy Specialist, JD-023).

The QA process specifically flagged this language as a coding hazard: a role can describe itself entirely in terms like these and still be substantively a process-administration role at L2, or substantively a contested-judgement role at L4. The language does not discriminate between those levels of work.

What this over-indexes against. Generic stakeholder-engagement language leaves the role's purpose unstated. When the description of the work is at the level of "partner with the business" or "drive maturity uplift", the role's actual operating model implications are unspecified. The capability the role will produce, once filled, is contingent on what the candidate makes of it.

This is the language layer where the gap between *described capability* and *operational capability* is widest. A role described in fully generic language can deliver excellent operational capability if the candidate is strong and the organisation supports the work. It can equally deliver a year of stakeholder meetings and an updated framework document. The advertisement gives the reader no way to predict which.

The corpus over-indexes on this language because it is the lingua franca of professional services hiring across all functions, not just privacy. The over-index claim here is industry-wide rather than privacy-specific. It nevertheless matters for the privacy market because privacy roles are particularly vulnerable to it: the work is genuinely cross-functional, the framing has been borrowed from compliance and risk advisory, and the operating-model design questions are subtle enough that the language layer does not pick them up.

What this pattern lets organisations do, and doesn't

Read together, the four over-indexing patterns describe an Australian privacy labour market that knows how to articulate privacy capability as compliance program work, as legal advisory work, as framework alignment work, and as generic strategic stakeholder engagement.

That language and that role design support three things well: the build-out of standard program components (policies, registers, training, vendor processes); the response to regulator engagement and enforcement (legal counsel and DPO roles handle this directly); and the demonstration of executive-and-board-visible progress (framework alignment scorecards and maturity uplift narratives).

They support three things less well: the embedded capability to bring privacy judgement into product, data, marketing and AI decisions while those decisions can still change; the independent challenge of those decisions from a second-line risk or assurance position; and the operating-model design work that allocates privacy decision rights across functions in an organisation that has grown past the size at which a central privacy team can carry the work alone.

There is a plainer way to put what the concentration produces. Compliance and legal roles produce paperwork, and paperwork is not the same thing as preventing harm to a person. The two are related, and the first is often necessary to evidence the second, but a market that buys compliance and legal roles at nearly three times the rate it buys embedded privacy roles (49% against 17%) has decided, without deciding, which of them it is really purchasing.

Section 8 takes up the under-described patterns directly. Sections 12 and 14 take up what the over-indexing pattern implies for organisations designing their privacy operating model in 2026.

A note on what is not being claimed

This section's claim is about the *aggregate framing pattern* across the corpus, not about the quality or substance of individual roles. Many of the roles named as examples in this section are excellent roles, doing genuinely substantive work. The over-index claim is about structure: the market is *over-purchasing* the compliance-and-legal framing and *under-purchasing* the operational and assurance framings. Individual organisations who have made deliberate choices outside the pattern are not contradicted by the aggregate claim.

To say it directly: this review's proposition does not call for organisations to stop hiring compliance and legal privacy capability. Compliance program stewardship and senior legal counsel are necessary. The proposition calls for the operating-model imbalance to be corrected: for organisations to add the operational, assurance and operating-model design capability they have under-built, alongside the compliance and legal capability they already have. The corpus evidences the imbalance; it does not prescribe the correction. Sections 12 and 14 take up the design questions.

8. What job ads under-describe

What "under-described" means here

The over-indexed patterns in Section 7 are visible because they appear so often. Under-described patterns are visible the opposite way: they are present in the operating-model picture this review's proposition names as needed, and absent or minimal in the corpus.

Four patterns are under-described in the corpus. They are not absent from the Australian privacy operating environment. Organisations do, in practice, sometimes do this work. But they are absent from how the work is being advertised when organisations describe what privacy capability they want to purchase.

Two of the four are about *roles that don't appear*: independent privacy assurance and privacy engineering. The other two are about *language that doesn't appear inside the roles that do*: contested-judgement verbs and operational decision-making phrasing. Both kinds of absence matter.

Pattern 1: Independent privacy assurance, the three-role finding

Evidence in the corpus. Risk / GRC Privacy carries three roles in the entire included sample of 101: HSBC Australia Data Risk and Technology & Cyber Risk Specialist (JD-036, L3), Protiviti Australia Technology Risk and Resilience Manager (JD-016, L3), and Plenti Risk & Compliance Manager (JD-081, L3). All three are L3. None are L4. None describe privacy as an independent second-line discipline.

This finding held across all six collection sprints, across both collection modes (broad sweep and focused query), and across deliberate effort to surface this archetype through GRC-related search queries. That is three roles across six sprints and both collection modes.

Three is a small enough number that it should be read as a count and not converted into a rate. What makes it a finding is not the size of the fraction but the stability: the archetype did not grow when the search strategy changed, and it did not grow when the corpus did.

Each of the three is, on close reading, a partial example rather than a clean one:


- **JD-036 (HSBC)** describes a Data Risk and Technology & Cyber Risk Specialist where privacy is named as one of several scope items inside a broader risk role. The privacy second-line scope is genuinely present but not the role's primary identity.
- **JD-016 (Protiviti)** describes a consultancy Technology Risk and Resilience Manager where privacy sits inside a broader technology risk scope. The second-line framing is present because

the role sits in a consultancy practice that delivers second-line work for clients; it is not an in-house second-line privacy role.

- **JD-081 (Plenti)** describes a Risk & Compliance Manager at a fintech lender where privacy is named as a substantive scope item inside a combined risk-and-compliance role. The clean separation between Line 1 (operate the privacy control) and Line 2 (independently assure the privacy control) is not what this role's design assumes.

Risk / GRC Privacy: cumulative count across all six sprints

The archetype never moved past three, across both collection modes and deliberate GRC-targeted queries.



The nearest thing the included sample has to senior second-line privacy challenge sits outside the archetype: the MUFG Head of AI & Data Privacy Risk (JD-114, L4) is an explicitly second-line risk leadership role, with independent review and challenge of Privacy Impact Assessments in scope, coded AI Governance Hybrid because AI risk is its primary identity. Its existence sharpens rather than weakens the finding: where second-line privacy challenge is being hired at leadership level at all, it arrives bundled with AI risk, not as a privacy assurance discipline in its own right.

What this under-describes. The three-lines-of-defence model is a near-universal operating-model assumption in Australian regulated industries. Banks, insurers, super funds, energy retailers, public sector agencies and many large enterprises operate variants of it: Line 1 owns and operates controls, Line 2 sets policy and independently challenges, Line 3 provides assurance. For privacy capability, the model predicts an in-house second-line privacy assurance function: roles whose job is to independently test whether the privacy controls Line 1 is operating are actually doing what Line 1 says they are doing.

That function is not visible in the corpus. The closest the corpus comes is the DPO / Regulatory Interface archetype, whose statutory officers exercise delegated authority that has second-line characteristics. But DPO work is not the same as independent risk assurance. The DPO is the inside-the-tent regulator; the second-line risk function is the inside-the-tent challenger of how the rest of

the organisation runs its privacy controls. These are different jobs, and only one is being advertised at meaningful volume.

The implication reaches into how operating models actually run. Either Australian organisations are running their privacy operating models without the Line 2 challenge layer the three-lines-of-defence model assumes, which would mean the model is honoured on paper and absent in practice for privacy specifically. Or the Line 2 privacy challenge work is being absorbed into broader risk roles (technology risk, operational risk, conduct risk) where it is not labour-market-visible as privacy work. The near-miss sample shows that happening: Section 11's senior compliance and risk category holds an internal audit role covering model, data and AI risk with privacy as one audit dimension, and a general risk and governance manager role with privacy implicit rather than named. Either reading has implications for the operating-model argument the report takes up in Section 14.


Pattern 2: Contested-judgement language

Evidence in the corpus. The 16-verb task taxonomy used in this review distinguishes between verbs that describe routine privacy work (*advise, interpret, review, draft, maintain, monitor, report, train*) and verbs that describe contested-judgement work (*assess, implement, embed, challenge, assure, design, govern, lead*). The taxonomy was designed to surface where role descriptions cross from describing routine work into describing the work of *deciding what to do under genuine regulatory or operational ambiguity*.

Five of those eight carry the signal most cleanly, and the analysis below uses them as the marker: *challenge, assure, govern, lead, and design* in the specific sense of operating-model design. *Assess, implement* and *embed* sit in the second group for coding purposes, but they are not contested-judgement work in themselves, and nothing in the analysis rests on them. The five appear in the corpus at much lower frequency than the routine-work verbs. The task-verb coding confirms it: *advise* appears among the top task verbs of 84 of the 101 included ads, *challenge* appears in three, and *assure* appears in none. The pattern is sharpest in the Compliance / Privacy Governance archetype, where *advise, maintain, review, monitor, report* dominate, and *challenge, govern, design* are rare.

Top task verbs across the 101 included advertisements

Number of includes whose coded top task verbs contain each verb. Assure: zero. Challenge: three.



The contested-judgement verbs do appear in the corpus. They concentrate in the senior band: the L4 and L3–L4 roles hold 13 of the 17 *design* occurrences and 9 of the 13 *govern* occurrences while carrying only a third of the included sample, and they sit mainly in the L4 Compliance, DPO and AI Governance Hybrid roles. In other words, contested-judgement language is largely a senior-leadership phenomenon. Mid-band privacy roles do not, for the most part, describe themselves in terms that admit the possibility of contested judgement.

What this under-describes. The privacy work that organisations describe wanting at mid-band level is, in JD language, largely *not* contested-judgement work. It is process work, advisory work, framework alignment work, training work, and monitoring work. The contested calls (what to do when the regulatory position is unclear, when the operational impact of a decision is asymmetric, when the AI governance question has no good answer) are implicitly being deferred upward to senior leadership.

That deferral pattern is workable when senior leadership has capacity to take it, which the corpus cannot measure: the L4 band describes what advertisements ask for, not capacity in post. What the corpus does show is that the mid-band roles being advertised do not name contested judgement as part of the work, which leaves open who is expected to recognise a contested decision and escalate it.

The capability gap this language pattern produces is not a gap in seniority. It is a gap in how the work is described, which shapes what an appointee is told the job includes.

Pattern 3: Privacy engineering as a named function

Evidence in the corpus. Data / Technology Privacy carries five roles: BHP Principal Privacy Engineer (JD-012, L3), Avanade Data Protection Manager (JD-037, L3), AustralianSuper Privacy & Data Protection Technology Project Manager (JD-064, L3), EPI-USE Labs Lead Consultant SAP Data Privacy (JD-091, L2-L3), and ANZ Privacy Risk Business Analyst (JD-070, L2). No L4 roles in this archetype. Only one role (BHP) uses "Privacy Engineer" as the substantive role title.

What this under-describes. Privacy engineering, meaning the technical implementation of privacy controls inside data and product infrastructure, the deployment of privacy-enhancing technologies, the configuration of platforms to support privacy-by-design, and the code-level review of how personal data flows through systems, is a recognised discipline internationally. Major US tech companies hire substantial privacy engineering teams, and the IAPP's Certified Information Privacy Technologist (CIPT) credential exists to certify exactly this subspecialty. The discipline has been emerging in Australia for several years.

The Australian labour market evidence in this corpus is that the discipline is largely not being advertised under "privacy engineering" titles. Where the work exists, it is being purchased under broader data engineering, security engineering, data governance, or platform-administration titles where privacy is one of several scope items.

This is partly a real market signal: the Australian privacy engineering talent pool is small relative to North America, and small employers cannot economically justify dedicated privacy engineering roles. It is partly a measurement artefact: privacy engineering work that exists under other titles falls into the near-miss bucket rather than the included sample. The truth is probably both.

The implication for the operating-model argument is that organisations relying on privacy engineering capability for AI deployment, data minimisation programs, and privacy-by-design implementation cannot, by and large, hire that capability through privacy-titled job posts. They have to grow it inside data, security, or platform engineering teams. The corpus does not evidence the absence of the capability, only the absence of the labour-market category in which it is normally purchased.

Pattern 4: Specific operational decision-making language

Evidence in the corpus. Several specific phrasings that would mark embedded privacy capability appear rarely in the corpus, even where the role is coded as L3 or L4 Operational Privacy:

- "Set the boundary for" / "determine what is in scope for the privacy review"
- "Choose between competing privacy interests"
- "Recommend against a launch where the privacy posture is not defensible"

- "Independently rate the privacy risk of a proposed product or process"

These are the language patterns that would indicate the practitioner is expected to reach a judgement and put it to the business, rather than supply input that someone else weighs. They are largely absent. The more common framings are softer: "partner with product teams on privacy implications", "support the business with privacy advice", "input into privacy risk assessments".

What this under-describes. The language pattern reveals a positioning pattern rather than an authority one. Operational Privacy roles in the corpus are framed as advisory *to* the business rather than as working inside it. The business makes the call, which is where the call belongs. What the advertisements do not describe is any point at which the owner has to engage with the privacy judgement before making it, or any moment at which privacy is guaranteed to be in the room while the choice is still open.

This is a coherent operating-model design choice for some organisations, particularly those where business autonomy is high and central functions are advisory by design. It is a less coherent choice for organisations where the privacy posture genuinely is the operational risk, because advice that carries no obligation can be heard and set aside without anything being recorded, and the organisation learns nothing from having sought it.

The corpus does not allow the report to claim that this is a universal pattern, but the textual evidence is consistent: across all 17 Operational Privacy JD bodies, explicit decision-authority phrasing (approve, veto, sign-off, final decision, decision-making authority) does not appear at all, while every one of the 17 carries advisory or partnership framing. The QBE Senior Privacy Lead (JD-058) and the UNSW Data Privacy Lead (JD-034) come closer than most through their uplift and leadership scope, though neither ad names explicit decision authority. The absence of that phrasing is not, by itself, the problem: the business owner should hold the decision. What is missing alongside it is any language about timing or obligation, and that combination is what the operating-model design sections need to address.

What the combined silence implies

Three patterns, the near-absence of independent privacy assurance roles, the concentration of contested-judgement language at senior leadership, and the advisory framing of embedded privacy work, combine to describe a market that advertises privacy capability without describing when it arrives at a decision or what the business is expected to do with it.

Privacy capability in the corpus is, predominantly, capability to: - Build and run the compliance program - Advise the business on privacy implications - Embed privacy considerations into business workflows - Maintain alignment to external frameworks - Engage stakeholders on privacy matters

It is, less often, capability to: - Independently assure that the privacy controls actually work - Reach a privacy judgement the business owner is obliged to answer before deciding - Enter a launch, a

vendor onboarding or an AI deployment early enough to change its shape - Design the operating model that governs how all of the above is decided

The absence of the second list does not mean the work doesn't happen. It means the work is not being purchased through the labour market in the way the first list is. Some organisations handle this through statutory officer authority (DPO roles), some through ad-hoc executive escalation, some through external advisory engagement, some by accepting the privacy posture the business hands them. None of these is a reliable substitute for role designs that say when privacy enters a decision and who has to answer it.

Section 14 takes up the operating-model implications. The point this section makes is narrower and prior: the labour market does not, on aggregate, describe the privacy capability that the operating-model argument names as needed. Where organisations want it, they have to build it through means other than reading what the market is advertising and asking for the same things.

9. Where privacy works inside the business

What this archetype captures

Operational Privacy is the corpus's third-largest archetype at 17 of 101 included roles (17%). The archetype captures roles whose primary privacy capability requirement is *embedding privacy into the operations of the business*, distinct from compliance program stewardship (which embeds the compliance function into the business) and from legal advisory work (which embeds legal review into the commercial pipeline).

The defining test is what the JD body asks the practitioner to do day-to-day. Operational Privacy roles describe the practitioner working alongside product, data, marketing, customer, vendor and AI teams to embed privacy considerations into decisions those teams are making. The work is not advice given from a separate office. It is co-located with the decision-making.

The 17 roles in this archetype are the corpus's clearest evidence base for what embedded privacy capability looks like when the Australian market purchases it explicitly. This section profiles them in three bands (the five L4 leadership roles, the eight L3 and L3–L4 core roles, and the four L2 and L2–L3 entry roles) and then steps back to read what the combined picture shows.

The 49% of the included sample whose primary archetype is Compliance / Privacy Governance or Legal Privacy Counsel also embeds privacy work into the business, in different ways. This section does not relitigate that scope. It profiles the 17 roles where operational embedding is the archetype's primary identity, and reads what they tell the reader about the kind of capability the market knows how to advertise as operational.

The 17 roles at a glance

CAPABILITY BAND	COUNT	ROLES
-----------------	-------	-------

L4 (operating model / contested judgement)	5	Victoria Legal Aid Data Protection Lead (JD-032), UNSW Data Privacy Lead (JD-034), Randstad Principal/Director Data Privacy & DPIA (JD-048), QBE Insurance Senior Privacy Lead (JD-058), Sportsbet Head of Privacy / Data Enablement / RAI (JD-126)
L3–L4 (operational embedding + framework leadership)	1	Woolworths Group Privacy Lead, Privacy and Customer Trust (JD-068)
L3 (operational embedding)	7	Australian Retirement Trust Privacy Specialist (JD-008), ASX Data Privacy Specialist (JD-013), Woolworths Group Privacy Specialist 12-month contract (JD-023), Picklebet Data Privacy Manager (JD-033), Talent (recruiter) Privacy Specialist (JD-044), Endeavour Group Senior Privacy Officer (JD-047), Sharp & Carter Legal (recruiter) Privacy Specialist (JD-095)

CAPABILITY BAND	COUNT	ROLES
L2-L3 (mid-band entry)	2	Correlate Resources Business Analyst Privacy/Data Flows/Transparency Uplift (JD-050), Talenza Senior Privacy Analyst (JD-159)
L2 (artefact + early embedding)	2	Victoria Legal Aid Data Protection Specialist (JD-001), Hudson Project Manager Compliance & Privacy (JD-045)

The distribution skews L3, as the whole corpus does, though slightly less than the corpus average: 8 of 17 are pure L3 or L3-L4 (47%, against the corpus's 55%), and 5 of 17 are L4 (29%, marginally above the corpus's 28%). The archetype is not noticeably more senior than the corpus; what distinguishes it is where the work sits.

Sector spread across the 17 roles is wide: retail (Woolworths x2, Endeavour), gambling and wagering (Picklebet, Sportsbet), insurance and superannuation (QBE, Australian Retirement Trust), capital markets infrastructure (ASX), legal aid (Victoria Legal Aid x2), higher education (UNSW), and the recruiter and consultancy market (Randstad, Talent, Sharp & Carter Legal, Talenza, Correlate Resources, Hudson). Six of the 17 (35%) are advertised through recruiters rather than directly by the employer, above the 22 of 101 (22%) across all included roles and close to Legal Privacy Counsel's 8 of 22.

AI governance signal across the 17 is moderate: 2 Strong (QBE, Sportsbet), 5 Some (ASX, Woolworths Specialist, Endeavour, Correlate, Talenza), 1 Mentioned (Victoria Legal Aid Lead), 9 None. AI scope is not yet the defining feature of Operational Privacy hiring. It is a defining feature of *senior* Operational Privacy hiring, particularly at QBE and Sportsbet.

The L4 cluster: what senior embedded privacy looks like

The five L4 roles are the corpus's most analytically informative window into how embedded privacy *leadership* is being designed in the Australian market.

QBE Insurance, Senior Privacy Lead (JD-058, L4, AI:Strong). Nine-month fixed-term contract role for senior privacy operational uplift inside a major Australian insurer. The fixed-term shape is significant: this is project-defined senior privacy work with a contracting horizon, not a permanent head-of-privacy posture. The work is described as uplifting privacy operating practice across data, marketing and AI workflows, with explicit responsible AI scope. The standout senior operational embedding role in the corpus, and a clear example of "we need to fix our embedded privacy capability over the next year" being purchased as a discrete project.

Sportsbet, Head of Privacy, Data Enablement and RAI (JD-126, L4, AI:Strong). The corpus's clearest example of senior privacy leadership absorbing responsible AI as defined scope. The role design combines privacy, data enablement (the data team's product mandate) and RAI (responsible AI) into one head-of-function remit. This is pattern A of the three Section 10 names: privacy-led

operational leadership absorbing AI governance, rather than AI governance being stood up as a separate executive function.

UNSW, Data Privacy Lead (JD-034, L4). Higher-education embedded privacy leadership in a major Group of Eight university. The role is described as operational rather than statutory officer work; the statutory FOI / Privacy responsibilities are not the role's primary identity. UNSW has substantial data activity across research, education, partnerships and student services, and this role is the embedded privacy capability lead across that breadth.

Victoria Legal Aid, Data Protection Lead (JD-032, L4, AI:Mentioned). Legal aid is a non-trivial environment for embedded privacy work: sensitive case data, vulnerable clients, multi-jurisdiction work, partnership with private practice. The Data Protection Lead role sits above the Data Protection Specialist (JD-001, L2) in the same organisation, a rare two-tier structure inside this archetype.

Randstad, Principal / Director, Data Privacy & DPIA (JD-048, L4). Recruiter-advertised principal/director role with DPIA as a named scope item. The role's framing centres on senior embedded privacy expertise being deployed into client environments: consultancy or interim leadership work rather than in-house leadership.

What the L4 cluster shows. Three observations.

First, the L4 Operational Privacy hiring pattern is concentrated in large enterprises with substantive data activity (QBE, Sportsbet, UNSW, Victoria Legal Aid) and in the recruiter market (Randstad). It is not concentrated in mid-size employers. The corpus cannot show how mid-size organisations source embedded privacy leadership; internal growth, consultancy and going without would all leave the same gap in the advertisements.

Second, the L4 Operational Privacy roles are disproportionately *project-shaped* or *interim-shaped*. QBE is a nine-month fixed term. Randstad is an interim placement. Sportsbet is a head-of-function with three named scopes (privacy + data enablement + RAI), which suggests the role is partly a reorganisation absorbing functions that were previously separate. This pattern is consistent with embedded privacy uplift being treated by organisations as a discrete piece of work to be completed, rather than a permanent capability to be staffed.

Third, the L4 Operational Privacy roles carry more strategy, framework and uplift language than the others in the archetype, but none asks the role to change the organisation's own operating model. In Section 6's audit, QBE, Sportsbet and Victoria Legal Aid read as mixed. UNSW is the one Operational Privacy role among the nine that came closest, and it describes establishing a framework inside a model that is already set.

The L3 core: what mid-band embedded privacy looks like

The seven L3 roles plus the one L3–L4 mid-band role are the corpus's clearest window into what embedded privacy *practitioners* are being hired to do day-to-day in 2026.

The dominant pattern across these eight roles is the **embedded specialist** model. A senior practitioner sits inside the privacy function and partners with one or more business areas. Common scope items in the JD bodies: vendor onboarding privacy review, product launch privacy support, DPIA execution across business initiatives, privacy input into data team projects, partnership with marketing on customer data use, partnership with technology on access management and data minimisation, partnership with HR on workforce privacy.

The differentiation between roles is mostly in *which business areas the practitioner is embedded with*:

- **ASX (JD-013, L3, AI:Some)**: data privacy specialist in capital markets infrastructure; embedded with the technology and data teams, AI scope explicit.
- **Australian Retirement Trust (JD-008, L3)**: privacy specialist in a major super fund; embedded with customer, data and product teams.
- **Woolworths Group Privacy Specialist 12-month contract (JD-023, L3, AI:Some)**: fixed-term embedded privacy specialist; AI scope present.
- **Endeavour Group (JD-047, L3, AI:Some)**: senior privacy officer in retail; the JD body emphasises customer data work, vendor partnerships, and AI tooling scope.
- **Picklebet (JD-033, L3)**: embedded privacy manager at a wagering platform; the smaller employer pattern, where the role is the privacy function rather than embedded within it.
- **Sharp & Carter Legal (JD-095, L3)**: recruiter-advertised privacy specialist; client-deployment work.
- **Talent (JD-044, L3)**: recruiter-advertised privacy specialist; client-deployment work.
- **Woolworths Group Privacy Lead, Privacy and Customer Trust (JD-068, L3-L4)**: the L3-L4 mid-band role; combines embedding work with framework-leadership scope under a "customer trust" framing rather than a pure compliance framing.

What the L3 core shows. The work the Australian market hires for as an embedded privacy practitioner is recognisably consistent across these eight roles. The practitioner partners with multiple business functions, executes DPIAs and privacy reviews as part of business decisions, and is the in-business privacy interlocutor for the central privacy function.

What is mostly absent from the JD bodies is *language about decision authority*. The practitioner provides expertise; the business decides. This is the advisory-framing pattern Section 8 named. In the Operational Privacy archetype specifically, the pattern is the dominant design assumption. Even at L4, no advertisement in the archetype names explicit decision authority, and none asks for a change to the operating model (Section 6).

The implication is that the L3 embedded privacy capability the Australian market is hiring is, in its dominant form, embedded *advisory* capability rather than embedded *decision-making* capability. The distinction matters for the operating-model argument in Section 14.

The L2 and L2–L3 entry band

Four roles: Victoria Legal Aid Data Protection Specialist (JD-001, L2), Hudson Project Manager Compliance & Privacy (JD-045, L2), Correlate Resources Business Analyst Privacy / Data Flows / Transparency Uplift (JD-050, L2–L3, AI:Some), and Talenza Senior Privacy Analyst (JD-159, L2–L3, AI:Some).

These roles describe embedded privacy work at the analyst, business-analyst and project-management level. The Correlate role is the clearest worked example of project-shaped embedded privacy work: a business analyst hired specifically to map data flows and uplift transparency disclosures as a discrete project. The Hudson role is the clearest example of embedded privacy work being purchased as project management: privacy and compliance treated as project deliverables rather than as embedded capabilities.

The L2 / L2–L3 band is small (four roles), but it points to a real entry channel for embedded privacy work that exists alongside the more common analyst-and-junior-officer entry into the Compliance / Privacy Governance archetype. The two entry channels produce practitioners with different starting orientations: compliance-trained junior privacy people start by maintaining artefacts and registers, operationally trained junior privacy people start by mapping data flows and supporting project delivery. The downstream capability implications differ.

Where this work *doesn't* appear

Three absences from the Operational Privacy archetype need naming.

Mid-size employers are rare. The 17 Operational Privacy roles are concentrated in large enterprises and in the recruiter / consultancy market. Where mid-size Australian organisations hold embedded privacy capability, the corpus cannot show how: internal growth leaves no external advertisement, consultancy appears under broader titles, and a Compliance / Privacy Governance role can carry operational expectations without saying so.

Public sector outside statutory officer roles is rare. Most public-sector privacy hiring in the corpus sits in the DPO / Regulatory Interface archetype, not the Operational Privacy archetype. Where public-sector roles describe genuinely operational embedding work (e.g. UNSW, Victoria Legal Aid), it is more often higher education or statutory bodies with substantive operational data activity rather than the broader public sector.

Specialist sub-discipline framing is rare. The corpus does not show embedded privacy work being advertised as specialised by domain. There is no "Operational Privacy for Marketing" or "Operational Privacy for AI/ML pipelines" role pattern. The work is described generically and the practitioner is expected to apply it across whichever business areas they end up partnering with. This is partly an

artefact of the labour market's small size (specialisation requires scale) and partly a real signal that the Australian embedded privacy practitioner is advertised as a generalist.

What this profile shows

The Operational Privacy archetype is real, identifiable, and present in the corpus at meaningful volume: 17 roles, 17% of the included sample. The work it describes is recognisably distinct from compliance program stewardship and from legal advisory work. The practitioners hired through this archetype are doing the embedding work this review's proposition names as needed.

Two qualifications follow, and the report's broader argument depends on both.

First, the 17-role count is small relative to the 49% combined Compliance + Legal advisory share. The market is purchasing embedded privacy capability at meaningful but secondary volume. The proposition that embedded privacy capability is under-built relative to compliance posture is consistent with this distribution: there is some supply, but not the share that would be expected if organisations had rebalanced their privacy operating models toward operational embedding.

Second, the language inside the 17 roles still tilts toward advisory framing rather than decision-rights framing. The work being purchased is mostly *embedded expertise* rather than *embedded authority*. This is the language pattern Section 8 named, surfacing in its sharpest form inside the archetype where one would most expect decision-rights language to appear. The implication is not that the 17 roles aren't embedded privacy work; they are. The implication is that even at the leading edge of how the Australian market hires for privacy work inside the business, the question of when privacy reaches a decision and who must answer it is largely left to the L4 leadership layer, which is hired thinly and often as project-shaped work.

Section 10 takes up the AI governance dimension that intersects with this profile. Section 14 takes up the operating-model design implications.

10. What AI governance reveals

The question this section answers

A discipline has settled when the labour market consistently agrees on what work it covers, who does it, where it sits in the organisation, and what capability profile the people who do it should have. By that test, and this is my reading of the period rather than something this corpus measures, privacy as a discipline largely settled in Australia between 2014 and 2020: the work, the titles, the reporting lines and the capability expectations cohered into recognisable patterns despite cross-sector variation.

AI governance, by the same test, has not settled. The corpus shows three different design patterns for where AI governance work lives in the organisation, all of them visible in the Australian market simultaneously, and none of them clearly winning the design contest.

This section profiles the three patterns, names them as competing hypotheses about the future shape of the discipline, and reads what the corpus's combined picture implies about how organisations should think about the question over the next two to three years.

The section draws on two sub-samples of the corpus: the 14 AI Governance Hybrid roles in the included sample, and the AI-governance roles in the near-miss sample where privacy is absent or minimal. Together these surface the three patterns.

The 14 AI Governance Hybrid roles at a glance

PATTERN LEAN	COUNT	REPRESENTATIVE ROLES
Privacy-led (privacy expertise absorbing AI scope)	5	Breville Privacy Specialist (JD-010), University of Sydney Senior Privacy Specialist AI (JD-102), MUFG Head of AI & Data Privacy Risk (JD-114), ctrl:cyber Principal Consultant Privacy/Data/AI Governance (JD-065), ctrl:cyber Senior Consultant Privacy/Data/AI Governance (JD-099)
Data-led (data governance expertise absorbing AI scope)	6	Reserve Bank of Australia Senior Manager Information and Data Governance (JD-040), Capricorn Data & AI Governance Specialist (JD-090), Protiviti Data & AI Governance Senior Manager (JD-060), AFCA Senior Manager Data & AI Governance (JD-082), Helia Data and AI Governance Leader (JD-147), Australian Criminal Intelligence Commission EL1 Data Governance Specialist (JD-111)
Standalone-AI (AI governance as its own primary identity)	3	Profectus AI Governance Advisor (JD-028), Slade Group AI Governance Lead (JD-152), Talent AI Governance & Assurance Specialist (JD-166)


The "pattern lean" column reflects which of the three design assumptions the role's JD body emphasises. Several of the roles straddle two patterns (the ctrl:cyber consultancy roles, for

example, sit between privacy-led and standalone) and the assignment is a primary classification, not a clean partition.

Outside the AI Governance Hybrid archetype, AI governance signal appears strongly in several other archetypes. The Legal Privacy Counsel archetype concentrates explicit AI scope in senior in-house counsel roles (Xero JD-142, Optus JD-139, CBA JD-097, ANZ JD-063, Canva JD-104). The Operational Privacy archetype carries Strong AI signal in two L4 roles (QBE JD-058, Sportsbet JD-126), both privacy-led designs where AI scope has been added to embedded privacy leadership. The Compliance / Privacy Governance archetype carries Strong AI signal at the senior leadership level (Bankwest JD-157, Itoh JD-165, Deputy JD-150). These are all variants of pattern A (privacy absorbing AI) or pattern B (data/compliance absorbing AI), and they reinforce the cross-archetype reach of the three-way split.

Where the 28 Strong-AI-signal includes sit


Strong AI scope reaches well beyond the hybrid archetype, concentrating in senior counsel roles.



The near-miss sample also evidences the patterns, particularly pattern C (standalone-AI). The roles that surfaced under broad privacy-search queries but were classified as near-misses because privacy was absent or minimal include a standalone AI governance executive (IAG's Executive Manager AI Governance & Operations from sprint 03), two roles that organise AI work through product and operations leadership rather than through a governance function (CommBank's Gen AI product-owner role and Zip Co's technology, data and AI directorship, both from sprint 04), and a data-led executive variant (Pacific Blue's Head of AI and Data, where privacy appears as one minor control area). IAG and Zip Co do not name privacy anywhere, so they sit among the twelve near misses that Section 11 separates from the absorption count: they say nothing about where privacy work goes, but they are direct evidence of how AI governance work is being designed. These near-miss roles are not in the included denominator, but the pattern analysis depends on them.

Three simultaneous designs for where AI governance lives

Pattern lean across the 14 AI Governance Hybrid includes, plus the near-miss cluster that reinforces pattern C.



Pattern A: AI governance absorbed into privacy

The design assumption. The organisation has an existing privacy function with mature embedded privacy capability. AI governance is treated as a substantive extension of the privacy practitioner's scope rather than as a separate discipline. The privacy lead, with appropriate uplift, governs the AI work alongside the privacy work.

Worked examples.

- **Sportsbet, Head of Privacy, Data Enablement and RAI (JD-126, L4, AI:Strong).** The clearest example of pattern A at L4 in the corpus. The role design fuses privacy leadership, data enablement responsibility and responsible AI into one head-of-function remit. The bet the organisation has made: a single senior leader can hold all three because the design questions overlap.
- **QBE Insurance, Senior Privacy Lead (JD-058, L4, AI:Strong).** Operational privacy uplift project with explicit responsible AI scope. The privacy lead is being asked to embed both privacy and AI governance into the operating practice over the contract term.
- **MUFG Pension & Market Services, Head of AI & Data Privacy Risk (JD-114, L4, AI:Strong).** Head-of-function role where AI risk and data privacy risk are explicitly combined under one risk leadership remit. The reporting line is risk rather than privacy, but the capability profile asked for is privacy-led.
- **Breville, Privacy Specialist (JD-010, L3, AI:Strong).** A specialist-level pattern A: the privacy specialist role description names AI governance as a substantive scope area. Privacy expertise absorbing AI work at the practitioner level.

Where pattern A is strongest. Organisations with existing senior privacy leadership and an operational-not-compliance orientation. The Sportsbet, QBE and MUFG advertisements all read that way. The pattern A roles in the corpus are disproportionately L4, disproportionately senior, and disproportionately at organisations with substantial customer-data or product activity where the AI deployment volume justifies treating AI governance as a near-term operational priority for the privacy function.

The bet pattern A makes. Privacy practitioners can grow into AI governance work, and the cost of duplicate functional leadership exceeds the cost of stretching the privacy leader's capability. The bet is defensible when the privacy function is already operational rather than compliance-flavoured. It is more questionable when the privacy function is itself underbuilt: adding AI scope to a compliance-flavoured privacy program largely produces a compliance-flavoured AI governance program.

Pattern B: AI governance inside data leadership

The design assumption. The organisation has an existing data governance function (often inside CDO, CIO or risk reporting lines). AI governance is treated as a natural extension of data governance, on the reasoning that AI systems are trained on data, deployed against data, and governed through data lineage and quality controls. Privacy is one of several considerations the data-and-AI governance leader holds, alongside data quality, master data management, regulatory data obligations and AI risk.

Worked examples.

- **Australian Financial Complaints Authority, Senior Manager Data & AI Governance (JD-082, L4, AI:Strong)**. A statutory financial-services body where the senior leader holds combined data and AI governance scope. The reporting line is data leadership; AI governance sits inside that function.
- **Helia, Data and AI Governance Leader (JD-147, L4, AI:Strong)**. Head-of-function role where data governance is the primary identity and AI governance is a substantive extension. Financial services again.
- **Capricorn, Data & AI Governance Specialist (JD-090, L3, AI:Strong)**. A practitioner-level version of pattern B in a mid-size financial cooperative.
- **Protiviti Australia, Data & AI Governance Senior Manager (JD-060, L4, AI:Strong)**. Consultancy-side version: senior leadership of a Data & AI Governance practice that delivers pattern B designs to clients.
- **Reserve Bank of Australia, Senior Manager, Information and Data Governance (JD-040, AI:Some)**. Public-sector financial regulator example. Information and data governance is the primary identity; AI is named in scope but at "Some" rather than "Strong" signal level. A softer pattern B.

- **Australian Criminal Intelligence Commission, EL1 Data Governance Specialist (JD-111, L3, AI:Strong)**. Public-sector specialist version of pattern B; data governance with AI as named scope.
- **Pacific Blue, Head of AI and Data (JD-125, near-miss, AI:Strong)**. The near-miss executive variant of pattern B: an enterprise AI-and-data strategy and governance head at an energy retailer, with privacy appearing as one minor control area. Privacy is thin enough that the role fell outside the included denominator, which is itself informative about how pattern B treats privacy at the executive level.

Where pattern B is strongest. Financial services and public sector. Both sectors have mature data governance functions that predate the AI governance question, and both have organisational habits of treating data quality, data risk and data regulatory compliance as functions that the data leadership owns. Pattern B is the most "natural" design for organisations whose existing data governance maturity exceeds their privacy operational maturity.

The bet pattern B makes. AI governance is more like data governance than it is like privacy governance, and data leadership can absorb the AI dimension more naturally than privacy leadership can. The bet is defensible when the organisation's data leadership has both the regulatory and the ethical-judgement scope to handle AI questions, and when privacy is sufficiently mature elsewhere that AI governance does not have to do double duty as a privacy uplift exercise. It is more questionable when data leadership has historically been technical and lacks the ethical-judgement language or the regulatory familiarity to handle contested AI decisions.

Pattern C: AI governance as a standalone function

The design assumption. AI governance is its own discipline, distinct from both privacy and data governance, and deserves its own executive function with its own reporting line, its own framework, and its own leadership profile. The standalone AI governance leader does not report to the privacy leader and does not report to the data leader; the AI governance function sits alongside both.

Worked examples (included sample).

- **Slade Group, AI Governance Lead (JD-152, L4, AI:Strong)**. Standalone AI governance leadership role; privacy is named only in passing. The role identity is AI governance, not privacy-with-AI or data-with-AI.
- **Profectus, AI Governance Advisor (JD-028, L3-L4, AI:Strong)**. Advisory consultancy version of pattern C; the role identity is AI governance practice.
- **Talent, AI Governance & Assurance Specialist (JD-166, L3, AI:Strong)**. Recruiter-advertised standalone AI governance role; the assurance framing is notable, closer to a Line 2 AI governance assurance role than the privacy-or-data leadership versions of the pattern.

Worked examples (near-miss sample).

- **IAG, Executive Manager, AI Governance & Operations (JD-124, sprint-03)**. Major Australian general insurer hiring senior AI governance leadership with no privacy content anywhere in the JD body. The cleanest pattern C evidence in the corpus.
- **Commonwealth Bank, Principal Technology Product Owner - GenAI, Responsible AI (JD-129, sprint-04)**. Major bank organising AI governance work as product management rather than as risk or compliance management: the role builds the platforms, guardrails and safety capabilities for Gen AI delivery, with privacy as a stakeholder area rather than the role's identity.
- **Zip Co, Director Business Management, Technology, Data & AI (JD-134, sprint-04)**. Fintech executive role bundling technology, data and AI operations leadership, with privacy not named as substantive scope.

Where pattern C is strongest. Large enterprises with substantial AI deployment ambitions (CBA, IAG), executive search and recruitment markets where standalone AI governance heads are being sourced (Slade, Talent), and consultancy practices that have built standalone AI governance offerings (Profectus, ctrl:cyber to a degree). The pattern is more visible at the L4 / executive end of the market than at the practitioner end.

The bet pattern C makes. AI governance is sufficiently distinct from privacy and data governance that it needs its own executive owner. The bet is defensible when the organisation has the size to support a dedicated function and when the AI deployment portfolio is large enough that bundling AI scope into either privacy or data leadership would overload either function. It is more questionable in mid-size organisations, where standing up a standalone AI governance function may produce visible leadership without producing operational embedding.

The simultaneity argument

The three patterns are not sequential. They are not "old way / new way / next way". They are simultaneous, observable in the Australian market within a single collection window; whether they are converging is a question for the 2027 edition.

This is the corpus's clearest evidence that AI governance has not settled as a discipline.

If the market were settling on one pattern, the corpus would show one pattern dominant and the others as residual cases. If pattern A (privacy-absorbed) were winning, the AI Governance Hybrid archetype would skew heavily toward privacy-titled roles with AI scope added; instead it splits five privacy-led, six data-led and three standalone-AI. If pattern B (data-led) were winning, the senior AI governance hiring would concentrate inside data leadership; instead standalone AI governance appears at the senior end too, in IAG's executive role and Slade's AI Governance Lead, alongside the data-led roles at AFCA, Helia and Pacific Blue. If pattern C (standalone) were winning, the privacy-absorbing and data-absorbing patterns would be thin; they are not.

The market is making three different bets simultaneously, in different organisations, at the same time, on the same underlying question. The bets are not yet showing differential outcomes the market can read.

Two further patterns reinforce the non-settlement reading.

Public-sector AI governance is fragmented across all three patterns. The University of Sydney (JD-102) leans pattern A (Senior Privacy Specialist AI, privacy-led). The Australian Criminal Intelligence Commission (JD-111) leans pattern B (Data Governance Specialist with AI scope). The Reserve Bank (JD-040) leans pattern B softly. State-government AI assurance roles (the NSW Government AI Governance & Assurance Specialist visible in sprint 06, recruiter-advertised) lean pattern C. Public sector is not converging either.

Consultancy and recruiter market behaviour is informative. The ctrl:cyber roles (JD-065, JD-099) bundle Privacy, Data and AI Governance into a single consulting offering: the consultancy is hedging the discipline question by being able to deliver any of the three patterns to clients. Slade Group and Talent are advertising standalone AI governance heads, suggesting executive search is operating on the assumption that pattern C is hireable. Both are commercially rational responses to a market that has not picked.

What this implies about the discipline

Three implications follow.

First, organisations choosing a pattern in 2026 are choosing in conditions of genuine uncertainty. There is no settled "best practice" to copy. There is no peer cohort whose pattern can be defended as the obvious right answer. The right answer for any given organisation is contingent on the existing maturity of its privacy function, its data function, and its appetite for standing up new executive roles.

Second, the discipline is unlikely to settle within the next two to three years. Discipline settlement requires either a regulatory shock that forces a particular pattern (which is not currently on the legislative horizon at sufficient specificity), or a body of evidence that one pattern outperforms the others (which requires operational track records that 2026's roles will not produce until 2027 or 2028 at earliest). The current cohort of AI governance leaders are largely first-occupants of their roles; the discipline-settling effects of their decisions will not be visible for some years.

Third, the practical implication for organisations is that the AI governance design choice should be made deliberately and revisited regularly. Defaulting to whichever pattern an adjacent organisation has chosen, or to whichever pattern the recruiter market is most actively pitching, is the risk the simultaneity creates; the corpus cannot show how often it happens. The three-pattern simultaneity means organisations need to choose their pattern based on what their own privacy and

data functions can actually carry, and to expect to revise the choice as the discipline's evidence base builds out.

The practical implication for privacy practitioners is that AI governance capability is a tractable but non-trivial extension to the existing privacy skill base, and that the discipline-not-settled environment creates substantial career range for practitioners who can credibly operate across all three patterns. The Legal Privacy Counsel concentration of AI scope in senior in-house counsel roles (Xero, CBA, Optus, ANZ, Canva) reinforces this: AI governance capability layered onto senior privacy expertise is the combination the market most often asks for in its most senior privacy roles.

Section 14 takes up the operating-model implications for organisations choosing an AI governance pattern. Section 13 takes up the practitioner-capability implications.

A note on what is not being claimed

This section does not claim that one pattern is *correct*. The three-way split is being described, not adjudicated. I have views on the design trade-offs, but those views are not the report's findings; they are professional commentary that sits separately from the evidence base this corpus produces.

The section also does not claim that the AI Governance Hybrid archetype as a whole is a settled category. The archetype was created for this review to hold roles that genuinely combine privacy and AI governance scope, with the understanding that the categorisation itself may not survive the discipline's eventual settlement. If the standalone-AI pattern wins out, "AI Governance Hybrid" will look like a transitional artefact of the 2024 to 2027 period. If the privacy-absorbed pattern wins, "AI Governance" inside privacy roles will eventually stop being remarkable enough to need a hybrid archetype. The category does its work for this review without prejudicing the answer to the discipline question.

11. Near misses and adjacent roles

The 55-role denominator

The corpus contains 55 near-miss roles: job advertisements that surfaced under privacy-related search queries but did not meet the inclusion threshold for the main 101-role sample. The test applied was scope insufficient for inclusion, which covers two different situations, and the difference between them matters enough to report separately.

Forty-three roles name privacy as scope inside another function. These are the absorption cases: privacy responsibility is real and is being carried, but by someone whose job is primarily general counsel, cyber leadership, senior compliance, FOI, data governance or AI governance. At 43 of 167 raw captures, that is 26% of everything the searches returned.

Twelve roles do not name privacy at all. These surfaced under privacy-related queries because of adjacency in the search index rather than anything in the advertisement. Every one was checked against its captured JD body: nine contain no occurrence of the word "privacy" anywhere, and the remaining three contain it only in recruiter boilerplate about how the candidate's own data will be handled during the application. There is no privacy responsibility in these roles to be absorbed. At 12 of 167, they are 7% of the searches, and they are a finding about how thin the search-term boundary is rather than about how organisations design privacy work.

Together the two groups make the 33% near-miss share, and the cumulative share moved only between 29% and 36% as the corpus grew, across all six collection sprints and both collection modes, closing at 33%. Sprint-by-sprint, the share of each individual batch varied widely (from none in the five-ad sprint 06 to 71% in sprint 03) with batch composition, so the stability claim attaches to the cumulative share, not to any single batch. The absorption share alone moved more, from 20% to 28% as the senior-skewed sprints arrived, and closed at 26%: consistently about one in four or five of everything the searches returned.

This section reads the 55 near-misses together, because the categories cut across the two groups, but the absorption claims rest on the 43. The "where privacy disappears" pattern is the review's central finding about where privacy work goes: privacy capability in the Australian market is being purchased through job titles that do not name privacy, in functions whose primary identity is something else.


The pattern matters because it changes how the report's findings about the *included* sample need to be read. The 101 included roles describe how organisations purchase privacy capability *when they label it as privacy*. The 43 absorption roles describe how organisations purchase privacy capability *when they do not label it as privacy*. Both together describe the Australian privacy operating environment more completely than either alone.

The 55 roles consolidate into six major categories, which together hold 43 of the 55, plus a long tail of twelve roles across nine smaller categories. That 43 is not the same 43 as the absorption count:

nine of the twelve roles that do not name privacy sit inside the six major categories, and three sit in the long tail. Each category below therefore gives two counts, the roles in it and the absorption roles among them. The six major categories hold 34 absorption roles and the long tail holds nine.

Where privacy disappears: the 55 near-miss roles

43 name privacy as scope inside another function; 12 do not name privacy at all. Consolidated category coding.



Absorption pattern 1: General Legal Counsel with privacy as one practice area

Fifteen roles, fourteen of them absorption, the single largest absorption pattern. Senior in-house and law-firm counsel roles where privacy is named as one of several practice areas the role covers. Variants include: General Commercial Counsel with privacy as one of many areas; Tech Commercial Counsel where privacy is an adjacent function; Payments / fintech counsel where privacy sits inside a regulatory bundle; Capital markets counsel where privacy is one of several frameworks; senior in-house General Counsel with privacy as one of several oversight areas.

These roles are not Privacy Counsel roles, but they evidently carry meaningful privacy load. The JD bodies typically name "privacy law" or "Privacy Act" in a list alongside other regulatory areas (consumer law, competition, technology contracts, data protection, AML/CTF, financial services regulation, advertising law).

What this absorption means. Substantial privacy advisory capability in the Australian market is being purchased through General Counsel role designs where privacy is one of many specialisations. The candidate hired into such a role will, in practice, do privacy work. But the role's

accountability does not name privacy as a defined deliverable, and the candidate's career path is shaped by how the role's broader practice areas develop, not by how privacy capability develops.

This is the operating-model equivalent of distributing the privacy advisory load across the legal function rather than concentrating it in a Privacy Counsel role. The choice is defensible in mid-size organisations where the privacy load does not yet justify a dedicated Privacy Counsel. It is more questionable in large enterprises where the privacy load is substantial and the general-counsel-with-privacy-as-one-of-many design produces under-specialised privacy advice at scale.

Absorption pattern 2: AI governance, ML and data-leadership roles where privacy is absent or minor

Eight roles, five of them absorption, and the most analytically important pattern even though it is not the largest. Senior AI governance, AI ethics, ML product ownership, or data leadership work where privacy is either absent from the role's identity entirely or named as one minor control area among many. The corpus's specific evidence:

- **Standalone AI Governance executive roles with privacy absent:** the strongest single piece of corpus evidence for Section 10's pattern C (standalone AI governance as its own discipline). IAG's Executive Manager, AI Governance & Operations (JD-124) describes senior AI governance leadership with no privacy content anywhere in the JD body. It is one of the three roles in this category that do not name privacy at all, so it is evidence about how AI governance work is designed rather than about where privacy work is absorbed.
- **Commonwealth Bank, Principal Technology Product Owner - GenAI, Responsible AI (JD-129, near-miss).** AI governance work organised as product management rather than as risk or compliance management. Privacy is a stakeholder area, not the role's identity.
- **Zip Co, Director Business Management, Technology, Data & AI (JD-134, near-miss).** Fintech executive role bundling technology, data and AI operations leadership with privacy not named as substantive scope.
- **AI/Data strategy leadership with privacy as one minor control area:** Pacific Blue's Head of AI and Data (JD-125), the near-miss executive variant of Section 10's data-led pattern B.
- **Data/ML leadership where privacy is absent:** senior data and machine-learning leadership roles with no privacy scope.
- **CIO leadership with privacy as one governance domain among many:** privacy treated as one of several governance areas the CIO oversees rather than as a discipline with its own leadership.
- **Gen AI safety and AI-governance analyst variants:** an AI governance role with privacy as one design principle among several, and an AI governance/ethics analyst role with privacy as an awareness area.

What this absorption means. Senior AI governance work is being purchased in the Australian market through job designs that do not require privacy expertise. This is not a hypothetical concern. The corpus shows it happening at large, regulated, sophisticated employers.

The implication for the report's broader argument about AI governance settlement (Section 10) is significant. Pattern C (standalone AI governance) is being implemented partly through role designs that leave privacy out of the AI governance leader's identity. Where this happens, the AI governance function is being designed in a way that will not produce the integrated privacy-AI capability that some organisations (Sportsbet, QBE, MUFG) are building through pattern A. The two patterns are not interchangeable in their operating-model implications.

The implication for individual privacy practitioners is that a substantial share of the AI governance leadership the market is hiring will be filled by candidates without privacy backgrounds. The career path for privacy practitioners who want to lead AI governance work runs through pattern A (privacy-led organisations) or pattern B (data-led organisations with privacy as a substantive consideration), not through pattern C (standalone), where the advertisements give privacy little or no place.

Absorption pattern 3: Cyber and security leadership with privacy absorbed or absent

Eight roles, five of them absorption. Senior cyber security, security governance, IT/Security GRC, and cyber-tech-data risk leadership roles where privacy is either treated as a sub-element of cyber posture or is absent from the role's identity entirely. Variants include: cyber security leadership with privacy absent; cyber/security governance leadership with privacy as one of several risk types; cyber GRC consulting practice areas; physical/personnel security leadership; public sector security leadership where privacy is one of many frameworks the role manages.

What this absorption means. A meaningful share of senior security hiring in the Australian market treats privacy as a sub-element of security posture, hires security leaders who do not have privacy expertise, and expects the privacy implications of security decisions to be handled either through escalation to a separate privacy function (if one exists) or through generalist security judgement (if not).

When cyber and privacy capability are organised separately and the cyber function leads incident response, the privacy implications of the response can be handled late, handled inconsistently, or handled by people whose primary expertise is elsewhere. Large Australian breaches since 2022 have kept that interface in public view, but the corpus shows senior cyber-led roles still being hired with privacy as an implicit secondary or absent scope.

Absorption pattern 4: Senior compliance and risk leadership with privacy as one sub-framework

Six roles, five of them absorption. Senior compliance and risk leadership roles where privacy is named as one of several regulatory frameworks the role covers. Variants include: senior compliance leadership where privacy is one sub-framework; financial services compliance where privacy sits alongside AML/CTF, conduct, prudential; capital markets compliance; insurance compliance; internal audit roles covering model, data and AI risk where privacy is one audit dimension; general risk/governance manager roles where privacy is implicit.

What this absorption means. Senior compliance and risk roles in regulated industries, particularly financial services and insurance, distribute the privacy obligation across compliance leadership rather than concentrating it in a privacy-titled role. This is closely related to absorption pattern 1 (general counsel with privacy as one practice area) but with a compliance reporting line rather than a legal reporting line.

The candidate hired into such a role will hold privacy responsibility under a compliance umbrella. Their interaction with the embedded privacy function (if one exists) is typically through compliance-to-compliance interfaces. The operating-model implication is that, in these organisations, privacy questions reach compliance leaders whose primary expertise is elsewhere.

Absorption pattern 5: FOI / Information access roles adjacent to privacy

Three roles, all absorption. Statutory information-access roles (FOI Officers, GIPA Managers, RTI specialists) where privacy is named as an adjacent function but not the role's primary identity. One sub-finding from sprint 04 belongs here. GIPA / FOI Managers in the corpus appear in two organisational shapes: combined with privacy (the integrated officer model visible at JD-042 UQ, JD-046 Parramatta, JD-052 ICWA in the included sample) and partnered with a separate privacy function (the partnered model visible in the near-miss roles).

What this absorption means. The integrated FOI-plus-privacy officer model is mostly captured in the included sample as DPO / Regulatory Interface roles. The partnered model, where FOI and privacy are separate functions that partner, surfaces in the near-miss sample because the FOI role's identity is not privacy and the role does not meet the inclusion threshold.

This is a sub-finding about how Australian public-sector and quasi-public organisations are designing the FOI-privacy interface. The two organisational shapes have different operating-model implications: the integrated shape gives one officer authority across both regimes; the partnered shape requires the two officers to coordinate. The corpus does not allow a conclusion about which shape produces better operational outcomes, but it does evidence that the design choice is being made in both directions in the Australian market.

Absorption pattern 6: Data governance leadership with privacy as one consideration

Three roles, two of them absorption. Senior data governance roles where privacy is one of several regulatory references the role manages, or where privacy is absent from the identity entirely. Distinct from absorption pattern 2 (AI/ML/data leadership where privacy is also absorbed) because these roles are specifically data-governance-titled rather than AI-or-ML-titled.

What this absorption means. A share of senior data governance hiring treats privacy as one of several data regulatory considerations rather than as a discipline with its own leadership. The pattern is similar to absorption pattern 4 (senior compliance) but with a data leadership reporting line rather than a compliance reporting line.

This pattern feeds the AI governance pattern B (data-led) finding in Section 10: where privacy sits as a sub-element of data governance, AI governance has a ready home there too, which is the data-led design Section 10 describes. The corpus cannot show whether the same organisations make both choices.

Where else privacy disperses: the long tail

The remaining twelve near-miss roles, nine of them absorption, distribute across nine categories, each too small to justify a separate sub-section but collectively informative:

- **AML / Financial Crime / Fraud roles** where privacy is named or absent (three roles). Privacy implications of financial crime detection (large-scale customer data analysis) are not, in these roles, treated as discrete scope.
- **Academic research roles** (two roles): one in privacy-preserving ML, one in AI security and privacy research. Not organisational privacy roles; included in the near-miss sample for completeness but they do not contribute to the operating-model argument.
- **Records and Information Management** where privacy is absent or implicit (one role). Distinct from FOI roles because the work is operational records management rather than statutory access.
- **Junior compliance officer** with privacy as one of several areas (one role).
- **Customer remediation** where privacy incidents are one of several issue types being remediated (one role; the post-breach customer remediation function as it surfaces in the labour market).
- **Legal AI adoption / product management** where privacy is a familiarity area (one role).
- **Information governance / DLP** where privacy is implicit rather than explicit (one role).

- **Research ethics / integrity leadership** where privacy and data governance are one element (one role).
- **Administrative / legal support inside an FOI-and-privacy team** without substantive privacy scope (one role).

The report's central arguments do not rest on these smaller categories. They evidence the breadth of the absorption pattern: privacy capability disperses through more functions than the six major categories alone would suggest, without producing standalone sub-findings.

What the combined picture shows

Three claims close this section.

First, the labour-market picture of Australian privacy capability is incomplete without the near-miss sample. A reader who reads only the 101 included roles will see a market where privacy work is purchased through privacy-titled roles in seven recognisable archetypes. The 43 absorption roles show privacy responsibility being purchased through job titles that do not name privacy, in functions whose primary identity is something else: 43 roles, against 101 included. Both together describe how the market is actually buying privacy capability; either alone is partial.

Second, the absorption patterns are systematic, not random. Privacy disperses preferentially into general counsel work, AI/data leadership work, cyber and security leadership, senior compliance and risk leadership, FOI roles, and data governance roles. These six categories hold 34 of the 43 absorption roles; the other nine spread across a long tail of smaller categories. The dispersion is structured by which adjacent functions are most plausible homes for privacy responsibility, and the structure tells us something about which functions the Australian market treats as natural privacy-adjacent functions.

Third, the absorption pattern interacts with the discipline-settlement question for AI governance. Section 10 argued that AI governance has not settled because the labour market is making three different bets simultaneously. Section 11 shows that one of those three bets (pattern C, standalone AI governance) is being made in part through role designs that leave privacy out of the AI governance leader's identity. The near-miss AI/data/ML leadership cluster carries much of the evidence for the pattern C finding, and is the specific reason the discipline-settlement question matters for privacy practitioners: a settled pattern C would mean privacy is permanently outside the AI governance leadership track at a meaningful share of large employers.

The implications for organisations designing privacy operating models (Section 12), for individual practitioners (Section 13), and for the operating-model design question more generally (Section 14) all run through the absorption pattern. Reading the included sample alone systematically understates how much privacy capability the Australian market is dispersing into non-privacy roles, and consequently understates how widely privacy responsibility is spread.

A note on what the near-miss classification did and didn't do

The 55 near-miss roles were classified using a "scope insufficient for inclusion" test: the role surfaced under a privacy search but does not have enough privacy scope to count as a substantively privacy-focused role. For 43 that means privacy is named but subordinate; for 12 it means privacy is not named at all. The threshold is a judgement call. A different threshold would have produced different included and near-miss counts.

The threshold used in this review was conservative: when in doubt, classify as near-miss rather than include. This is consistent with the broader QA discipline of underclaiming rather than inflating operational privacy signal. The effect is that the 101 included sample is a relatively pure privacy-role sample, and the 55 near-miss sample is a relatively broad sample of privacy-adjacent work.

A re-run of the corpus with a more inclusive threshold (e.g. counting general counsel roles with substantive privacy scope as includes) would, on my estimate, shift roughly 8 to 12 roles from near-miss to include, changing the percentages in Section 4 and the archetype distribution in Section 5 modestly without changing the qualitative pattern in either direction. The absorption finding is robust to plausible threshold variation, because the patterns it evidences would still be visible at any reasonable threshold.

12. Implications for organisations

Reading the corpus as a buyer

For an organisation thinking about its own privacy capability, the corpus has one primary use: it shows what the market has been buying. That is not the same as what the market should be buying. It is a useful reference point because most organisations design their privacy hiring partly by reference to what other organisations are doing, and the corpus makes that reference point explicit.

This section names five decisions the corpus surfaces for organisations designing or revisiting their privacy operating model in 2026. Each is presented as a choice with the corpus evidence on what the Australian market is doing, what that choice produces operationally, and what the alternative would require.

The section does not recommend a single answer. The right choice for any given organisation depends on its existing privacy maturity, sector, regulatory exposure, AI deployment portfolio, and organisational risk appetite. The corpus is the reference point; the decision is the organisation's.

Decision 1: rebalance compliance posture toward operational embedding, or not

What the corpus shows. 49% of included roles describe primary privacy capability as compliance program stewardship or legal advisory work. 17% describe operational embedding as the primary scope. The Australian market is buying compliance and legal capability at roughly three times the rate it is buying embedded privacy capability.

What this design choice produces. An organisation that follows the market pattern hires for the compliance and legal capability, builds out the standard program components, and produces visible privacy artefacts. Privacy decisions inside business workflows are made by business teams, with privacy practitioner input when the business asks for it. The pattern is defensible when the business has low risk appetite for unilateral privacy decisions and when the privacy function is well-resourced enough to be consulted reliably.

Where this design choice fails. When the business does not ask for privacy input until late in the decision (after the vendor is signed, after the product launches, after the AI tool is procured), the compliance-and-legal capability cannot retrofit operational embedding into decisions that have already been made.

What the alternative requires. Rebalancing toward operational embedding means hiring fewer compliance and legal roles in proportion and more Operational Privacy roles, accepting that the operational roles will produce fewer visible artefacts and more visible-to-the-business decision moments. The corpus shows that the Operational Privacy hiring pattern exists (17 roles,

recognisable design) but is smaller, more often project-shaped, and more often advisory-framed than would be required to actually rebalance. Organisations choosing this alternative are leading the market, not following it.

Decision 2: pick an AI governance pattern deliberately

What the corpus shows. Three patterns operate simultaneously: privacy-absorbed (Sportsbet, QBE, MUFG), data-led (AFCA, Helia, Capricorn, plus the Pacific Blue near-miss executive), and standalone (Slade, Profectus, plus the near-miss IAG executive role). None is winning.

What this design choice produces.

- **Pattern A (privacy-absorbed)** produces an integrated privacy-AI capability inside the existing privacy function. It works when the privacy function is operationally mature. It under-performs when the privacy function is compliance-flavoured: a compliance-flavoured privacy program that absorbs AI governance largely produces a compliance-flavoured AI governance program.
- **Pattern B (data-led)** produces AI governance inside the existing data governance function. It works when data leadership has both regulatory familiarity and ethical-judgement capacity. It under-performs when data leadership has historically been technical and lacks the contested-judgement language to handle ambiguous AI decisions.
- **Pattern C (standalone)** produces a dedicated AI governance function alongside privacy and data. It works when the AI deployment portfolio is large enough to justify the function and when the organisation has the executive bandwidth to integrate three governance functions effectively. It under-performs in mid-size organisations, where standing up a standalone function produces visible leadership without operational embedding.

Where it goes wrong. With no settled pattern to copy, the risk is less choosing badly than not choosing deliberately: defaulting to whichever pattern adjacent organisations have chosen, or to whichever pattern the recruiter market is most actively pitching, without a clear-eyed view of which pattern fits the organisation's existing function maturity. The corpus shows the conditions for that default, not how often it happens.

What this implies practically. An organisation choosing an AI governance pattern in 2026 should do so by reading three things: the maturity of its existing privacy function, the maturity of its existing data governance function, and the realistic capacity of its executive layer to integrate any new governance function effectively. The right pattern is the one whose entry conditions the organisation actually meets, not the one its industry peers have most recently announced.

Decision 3: revisit the cyber-privacy interface explicitly

What the corpus shows. The senior cyber and security leadership roles in the near-miss sample treat privacy as a sub-element of security posture or leave it out (absorption pattern 3 in Section 11). Large breaches since 2022 have put this interface in public view, but the corpus shows no sign of cyber leadership roles being designed around it. Senior cyber-led roles are still being hired with privacy as implicit secondary or absent scope.

What this design choice produces. When cyber and privacy are organised separately and cyber leads incident response, privacy implications can be handled late, handled by people whose primary expertise is elsewhere, or escalated through a separate privacy function that has not been part of the incident timeline. The breach response, the regulator notification, and the customer remediation can all be operationally degraded by this pattern.

Where it goes wrong. The interface is rarely re-examined until an incident forces it. Organisations that have not had a public breach have less commercial pressure to revisit the design, even when their cyber-privacy interface has the same operating-model risk as organisations that have. The corpus does not capture how often the interface is being reviewed proactively versus reactively, and one collection window cannot show whether cyber-leadership hiring is changing.

What the alternative requires. Either elevating privacy to be co-equal with cyber inside the security function (which the corpus does not show happening at scale), or designing an explicit cyber-privacy interface with named decision authority on either side. The corpus does not produce the answer; it produces the question. Organisations that have not asked the question recently should.

Decision 4: privacy counsel concentration versus general counsel distribution

What the corpus shows. The Legal Privacy Counsel archetype carries 22 included roles where privacy is the role's defined remit. The General Counsel absorption pattern in the near-miss sample carries another 15 roles where privacy is one of several practice areas inside a broader legal role.

What this design choice produces.

- **The concentrated model** (dedicated Privacy Counsel) produces deep privacy advisory capability inside a focused role. The candidate is hired for privacy expertise; the career path is privacy-shaped; the advice is specialised.
- **The distributed model** (privacy as one practice area of a generalist counsel) produces broader commercial legal capability with privacy as one input. The candidate is hired for legal-generalist capability; the privacy work is part of the load but not the role's defining feature.

The labour-market signal. In the corpus, the concentrated model appears mostly at large employers in regulated industries (banks, telcos, major retailers, software platforms), and the distributed model more often at smaller employers and law firms, judged from the employers named rather than a coded size field. The threshold appears to sit at the scale where the privacy legal load justifies a full-time Privacy Counsel. The corpus shows where the two models cluster but not where the line falls: employer headcount was not a coded field, so any specific size figure would be an assertion rather than a finding. Regulatory exposure and data intensity move the threshold in ways a headcount number would obscure in any case.

What this implies practically. An organisation crossing the threshold from distributed to concentrated should expect that the transition is not automatic. The existing generalist counsel does not become a Privacy Counsel by adding privacy to their title. The transition typically requires hiring a dedicated Privacy Counsel and explicitly separating the privacy load from the generalist commercial work.

Decision 5: build the Line 2 privacy function, accept its absence deliberately, or find an alternative

What the corpus shows. Three Risk / GRC Privacy roles across the entire corpus. Of those three, none describe a clean second-line privacy challenge function. The independent privacy assurance capability the three-lines-of-defence model assumes is essentially absent from the privacy roles advertised in this corpus.

What this design choice produces.

- **Building it.** Hiring privacy-specialist risk practitioners into a second-line function alongside existing embedded privacy capability. Produces independent challenge of the operational function's decisions. Requires meaningful headcount investment and a clear reporting line that separates Line 2 from Line 1 in practice, not just on the organisation chart. One framing matters here: Line 2 challenge tests privacy decisions from outside the work; it cannot operationalise privacy on behalf of the business. Building Line 2 complements Line 1 operational capability. It does not substitute for it, and an organisation that builds the challenge layer while leaving Line 1 operational capability underbuilt has produced independent challenge of decisions no one is equipped to make well. The corpus shows almost no Australian employers hiring for privacy-specific Line 2 either way.
- **Accepting its absence.** Acknowledging that privacy will not have the second-line challenge the three-lines model otherwise assumes, and designing the privacy operating model around concentrated Line 1 capability with strong DPO-level statutory authority and executive-level governance oversight.
- **Finding an alternative.** Absorbing the Line 2 privacy challenge function into broader risk roles (technology risk, operational risk, conduct risk) that have privacy as one of several scope items.

The pattern visible in the absorption category in Section 11, and in the MUFG AI-and-privacy-risk leadership case in Section 8. Produces partial second-line coverage without standing up a privacy-specific Line 2 function.

Where it goes wrong. Listing three lines of defence in the operating model document and then operating, in practice, with one and a half. The corpus does not surface this directly, since operating-model documents are not in the labour-market evidence base, but the labour-market absence of Line 2 privacy roles is consistent with it.

What the alternative requires. Either building Line 2 explicitly (very few Australian employers are currently hiring for this) or naming the alternative in the operating-model document directly. Pretending to have a three-lines model when the labour-market hiring is consistent with a one-and-a-half-lines model is the least defensible position.

The counter-reading: maybe the market is buying exactly what it needs

One objection runs underneath all five decisions and deserves a direct answer rather than a footnote. The labour market's hiring pattern is revealed preference. If operational embedding and independent privacy assurance were worth buying, organisations would be buying them. On that reading, the corpus does not evidence an under-built capability layer; it evidences a market that has weighed those roles and declined them, and the imbalance this report describes is the market allocating correctly.

The objection is partly right, and Section 12 is built to accommodate it: for an organisation with low privacy decision volume, stable processes and a competent compliance function, declining to buy operational embedding can be the correct call, which is why each decision above is framed as a choice with entry conditions rather than a prescription.

As a general reading of the corpus, revealed preference assumes conditions the privacy market does not meet. The logic works where buyers get fast, attributable feedback on what their purchase failed to do. Privacy failure feedback is slow, and it lands elsewhere first: the cost of a weak privacy operating model falls on the individuals whose data is mishandled well before it falls on the organisation, and when it does reach the organisation it is rarely traced back to role design. Hiring patterns also reproduce themselves for ordinary institutional reasons. Position descriptions are copied from the previous hire and from peers, recruiters pitch what has previously filled, and the operating-model question is settled implicitly at requisition rather than deliberately at design.

The corpus cannot prove the market is wrong, and this report does not claim it does. What the corpus shows is consistent with the dominant pattern being a default rather than a decision, and the five choices above exist so that an organisation's answer is the second rather than the first.

What these five decisions have in common

The five decisions sit downstream of one prior choice the corpus does not make for any organisation: the choice about *what kind of privacy function the organisation wants to have*.

Three coherent shapes are visible in the corpus, none of them necessarily better than the others:

- **The compliance-led shape**, dominant in the corpus, where the privacy function is primarily a program-stewardship and advisory function with the business as the decision-maker. Decision 1 stays with compliance; Decision 2 leans pattern B or pattern C; Decision 3 accepts the cyber-led security interface; Decision 4 follows the market threshold; Decision 5 accepts Line 2 absence.
- **The operational shape**, smaller in the corpus, where the privacy function is primarily an embedded decision-input function with named authority over specific decision moments. Decision 1 leans operational; Decision 2 leans pattern A; Decision 3 elevates privacy inside security; Decision 4 concentrates; Decision 5 builds Line 2 or names the alternative explicitly.
- **The statutory officer shape**, visible in the DPO archetype, where the privacy function is led by a designated statutory officer with delegated authority across the organisation. Decision 1 hybridises compliance and operational under the officer's authority; Decision 2 sits inside the officer's remit; Decision 3 is mediated through the officer; Decision 4 still operates; Decision 5 partly substitutes for Line 2 through the officer's authority.

An organisation that has not consciously chosen between these three shapes is, in practice, getting whichever one its labour-market hiring patterns produce.

There is a prior question the corpus raises about all five decisions, and it is a question about who is being hired to answer them. Each of the five is a change to the conditions under which privacy decisions get made: where privacy sits, which AI governance pattern applies, how the cyber interface works, how the legal advisory load is distributed, whether Line 2 exists. Section 6's secondary audit found no advertisement in the 101 that puts an organisation's own decision conditions on the table, and none carrying a Chief Privacy Officer mandate. Nine describe building governance inside conditions already settled. The five decisions above are not being written into the roles the market is advertising.

That does not mean organisations lack the capability. It may sit with an executive appointed years ago, with a general counsel, with a consultancy, or with someone whose title says nothing about privacy. It does mean an organisation working through these five decisions should establish where that capability actually sits before assuming a hire will supply it, because the roles being advertised across the market are largely scoped to run a privacy operating model rather than to change one.

Two things follow, and neither is a hiring problem.

The capability is not gated on title. A practitioner does not need to be a Chief Privacy Officer to see the pattern across individual decisions and change how the function handles it. What they need is a mandate or executive sponsorship that makes the change possible. The corpus is silent on

capability and reasonably clear on position: it is not advertising roles placed where operating models get changed. An organisation that wants the work done has a placement question to answer before it has a recruitment one.

And it can be grown rather than bought. The capability develops out of the level below it, from practitioners who already exercise judgement inside consequential decisions and have started to notice the same decision arriving repeatedly for the same structural reason. An organisation with strong practitioners and a weak operating model is often closer to this than it looks, and closer than the labour market would suggest, because the market is not where the capability is being formed. What it takes is exposure to the pattern, sponsorship to act on it, and someone senior enough to make a change of arrangement stick.

The market default is the compliance-led shape. For some organisations that is the right answer. For others it is not, and AI governance pressure in 2026 raises the cost of getting that choice wrong.

Section 14 takes up the operating-model design questions in more depth. Section 13 takes up what the same evidence base implies for individual privacy practitioners building their careers in this market.

13. Implications for privacy practitioners

Reading the corpus as a candidate

For a privacy practitioner thinking about their next role, their career direction, or the capability they should build over the next two to three years, the corpus has one primary use: it shows what the Australian labour market is currently rewarding. Career strategy that follows the labour-market signal is not the only sensible strategy. Practitioners with strong views about where the discipline should go are within their rights to invest against the market. But the signal deserves a deliberate read before anyone chooses to invest against it.

This section names six readings of the corpus that matter for individual practitioners. None of them constitute career advice in any narrow sense; they are descriptive observations about what the Australian privacy labour market is buying and what that implies for which capability profiles are currently being most rewarded.

The section is honest about labour-market reality. Some of the readings will be uncomfortable for practitioners who have invested heavily in capability profiles the market is not currently rewarding. The corpus does not promise that the market is correct; it describes what the market is doing.

Reading 1: AI governance scope is concentrated in the most senior privacy roles

The labour-market signal. Across the corpus, the highest concentration of Strong AI governance signal sits inside the Legal Privacy Counsel and AI Governance Hybrid archetypes at L4 level. The included roles that pair senior privacy expertise with explicit responsible AI scope include Xero, Optus, CBA, ANZ, Canva (the Senior Legal Counsel route); QBE, Sportsbet (the Operational Privacy L4 route); MUFG, AFCA, Helia (the AI Governance Hybrid L4 route); and Bankwest, Itoh, Deputy (the Compliance / Privacy Governance senior route). These roles sit at the top of the market's advertised seniority range: 18 of the 34 roles in the L4 and L3–L4 bands carry strong AI governance signal, against 10 of the 67 roles below them.

What this implies for practitioners. A senior privacy practitioner who adds genuine AI governance capability matches the profile the market most often asks for in its most senior advertised privacy roles. That demand appears across the privacy-led and data-led patterns, so a practitioner who can operate in more than one has the widest set of roles to apply for.

The risk. AI governance is a fast-moving discipline. Capability built today may need re-anchoring as the discipline settles (or fails to settle). The practitioner who invests in deep, framework-specific AI governance expertise (e.g. ISO/IEC 42001, NIST AI RMF, EU AI Act compliance language) is making a

more conservative bet than the practitioner who invests in operating-model-design capability for AI governance; the latter is more portable across pattern shifts but harder to evidence in a CV.

What this does not mean. It does not mean that practitioners without AI governance capability are unhireable. The 58 includes coded "AI signal None" are real roles with real candidates filling them. It does mean that, among the most senior advertised roles, strong AI governance scope is more common than not.

Reading 2: the embedded privacy career path is real but smaller and more project-shaped than the compliance or legal paths

The labour-market signal. The Operational Privacy archetype carries 17 included roles. Of those, 5 are L4 leadership roles, 8 are L3 core practitioner roles, and 4 are entry-band roles. The L4 leadership roles are disproportionately project-shaped (9-month FTCs, interim placements, reorganisations absorbing previously separate functions).

What this implies for practitioners. The embedded privacy career path is identifiable and has a defined L3 to L4 progression visible in the corpus. It is smaller than the compliance or legal paths, more contract-shaped at the senior end, and concentrated in large enterprises and the recruiter market. A practitioner choosing this path should expect more interim and project-shaped engagements than they would in a compliance or legal track, and more variability in where their capability is most valued.

The advantage. Embedded privacy capability is the capability layer this review's proposition names as the under-built part of the Australian privacy operating environment. Practitioners building this capability are investing against current market scarcity. Where organisations decide to rebalance their privacy operating model toward operational embedding (Decision 1 in Section 12), they need embedded privacy practitioners, and the supply of them is small.

The trade-off. The senior embedded privacy market is small enough that practitioners reach L4 leadership through fewer routes than in the compliance or legal tracks. The corpus shows five L4 roles in this archetype, against seven L4 in Compliance, five L4 in Legal Counsel, and six L4 in AI Governance Hybrid. Operational privacy L4 is hireable but the route is narrower.

Reading 3: the Legal Privacy Counsel track is the most clearly defined senior privacy career path in the Australian market

The labour-market signal. Legal Privacy Counsel is the corpus's second-largest archetype at 22 roles, with 5 L4 roles plus 2 L3–L4 mid-band roles. Senior in-house privacy counsel positions at major banks (CBA, ANZ), telcos (Optus), and platform employers (Xero, Canva) anchor the upper end. The

track has Strong AI signal concentration: almost every L4 Legal Privacy Counsel role in the corpus carries explicit AI governance scope.

What this implies for practitioners. For practitioners who hold a current Australian practising certificate or have substantial in-house counsel experience, the Legal Privacy Counsel track carries five L4 roles, behind Compliance / Privacy Governance (seven) and AI Governance Hybrid (six). The senior in-house counsel route is well-defined, the career path is clear, and four of its five L4 roles carry strong AI governance scope.

The constraint. The track requires a legal qualification and a practising certificate (or recent in-house counsel experience). Practitioners without that background cannot enter the track directly. The pattern produces a division in the Australian privacy practitioner population: legally-qualified practitioners have a well-defined route to senior roles that non-legal practitioners do not.

What this implies more broadly. The most clearly defined senior privacy career path in the corpus is a legal career path with privacy specialisation, rather than a privacy career path with legal qualification.

Reading 4: the statutory officer / DPO track is its own thing

The labour-market signal. DPO / Regulatory Interface carries 13 included roles, with a high L4 share (5 of 13). The track concentrates in the public sector (state agencies, regulators, statutory bodies), in higher education, and in regulated industries with statutory privacy obligations beyond the Privacy Act (financial services, health). State-level reforms (WA PRIS Act 2024, Vic PDP Act, Qld IP Act) are visibly creating new statutory officer roles.

What this implies for practitioners. The statutory officer track is a defined career path with its own capability profile (delegated authority, regulator engagement, statutory information access experience, FOI/RTI/PRIS familiarity). It is largely non-fungible with the corporate embedded privacy track: the capability requirements differ enough that practitioners typically build their careers in one track or the other, not both.

Where the track is growing. State-level reforms are creating statutory officer roles in jurisdictions that did not previously have them. Practitioners with FOI / RTI / GIPA / PRIS experience are positioned to take advantage of this growth. The track is also one of the few where mid-band and senior public-sector privacy hiring is consistently visible.

The constraint. The track is largely state-bound. WA statutory officer experience is most portable to WA roles; Vic to Vic; Qld to Qld. National regulator interface roles (OAIC, ACCC) have their own narrower entry channels. Practitioners building this track typically settle into one jurisdiction.

Reading 5: privacy engineering and privacy assurance career paths exist but are not currently labelled as privacy

The labour-market signal. Data / Technology Privacy carries five roles. Risk / GRC Privacy carries three. Both archetypes are small absolute numbers, and both are missing L4 roles. The L4 privacy engineering leadership role and the L4 privacy assurance leadership role do not appear in the corpus as privacy-titled roles.

What this implies for practitioners. Privacy engineering capability and privacy assurance capability are real disciplines that the Australian market needs, but the labour-market access to them currently runs through non-privacy titles. Privacy engineers build their careers in data engineering, security engineering, or platform engineering roles with privacy as one of several scope items. Privacy assurance professionals build their careers in technology risk, operational risk, or internal audit roles with privacy as one audit dimension.

The career strategy implication. Practitioners who want to build privacy engineering or privacy assurance capability should expect their early career to be labelled as something else. The senior privacy engineering or senior privacy assurance leadership roles do not yet exist as a labour-market category in Australia. They may emerge over the next three to five years (I expect they will, as the embedded privacy capability conversation matures), but practitioners cannot yet plan around their existence.

A note for hiring organisations. This is one of the places where the labour-market default produces the wrong answer. An organisation that needs privacy engineering capability and reads only the privacy-titled labour market will conclude that the capability does not exist. It does, under other titles. The practitioner is the same person; the role design is the difference.

Reading 6: the capability the market does not advertise is not gated on the titles it does

The labour-market signal. Section 6's secondary audit found no advertisement in the corpus asking anyone to change the conditions under which privacy decisions are repeatedly made, and no Chief Privacy Officer or equivalent role anywhere in the 101. A practitioner reading the market for a route into that work will not find one advertised.

What this implies for practitioners. It does not mean the work is closed off, and it does not mean the capability belongs to people with senior titles. Seeing the pattern across individual decisions, and changing how the function handles it rather than handling each instance personally, is capability rather than rank. It develops out of the level below it: from a practitioner who is already exercising judgement inside consequential decisions and has begun to notice the same decision arriving repeatedly for the same structural reason. The corpus contains a great many roles that would provide that exposure, because roles at that level are what the market is mostly buying.

The constraint is not capability, it is position. Noticing the pattern is not the same as being able to change it. That takes a mandate, or executive sponsorship willing to alter how decisions are routed, and the corpus suggests those are rarely attached to advertised privacy roles. A practitioner building this capability should expect to have to find the sponsorship rather than to be given it, and to make the case for a change of arrangement in an organisation that did not hire them to make it.

What this does not license. Not every privacy decision needs an operating-model response, and reaching for one where a clear regulatory answer would do is a misuse of the capability rather than a demonstration of it. The value lies in recognising which problems are recurring and structural, and which are simply hard.

A note on what these readings are not

This section does not tell any individual practitioner what to do with their career. It describes what the labour market is currently rewarding. Practitioners have their own reasons for choosing capability profiles, work patterns, and career trajectories that may or may not align with what the market is buying. Those reasons are valid; this section does not adjudicate them.

It also does not predict the future. The six readings describe the Australian privacy labour market in 2026. AI governance discipline settlement, Privacy Act reform pace, large-employer breach response, and operating-model rebalancing across the market will all shift the labour-market signal over the next several years. Practitioners building their careers against the current signal should expect to recalibrate every two to three years.

Finally, this section does not romanticise either the compliance-led market default or the operational-privacy alternative. The compliance-led roles are real privacy work done by real practitioners who keep organisations within their regulatory obligations and protect customers from concrete harms. The Operational Privacy roles are also real privacy work done by real practitioners who embed privacy into business decisions. Both contribute. The section's claims about market rewards do not imply that one form of practice is more valuable than the other; they describe what the labour market is currently paying for, which is a different question.

14. What this means for operating model design

The operating-model question this report can answer

Most operating-model conversations about privacy in 2026 are conducted at the level of organisational chart and reporting line. The privacy function sits under Legal or under Risk or under the COO. The DPO reports to the General Counsel or to the CEO. AI governance reports to the CTO or to the privacy lead. These are real decisions, but they are not the operating-model decisions that shape whether privacy capability actually works.

The operating-model decisions that shape whether privacy capability actually works are about *where privacy sits and when it arrives*: whether privacy meets a consequential choice while credible alternatives are still open, who has to respond to the privacy judgement rather than merely receive it, where an unresolved disagreement goes, and what work the operating model expects each role to do versus what it expects to be done by escalation.

This section reads the corpus for what it tells us about those decisions. It does not propose a new operating model. It identifies four patterns in how the Australian market is designing privacy operating models implicitly through its labour-market hiring, and names what those patterns produce when carried into practice.

The four patterns are not exclusive. Most organisations exhibit several. The section's claim is that the labour-market evidence makes the trade-offs between them legible in ways the operating-model document does not.

Where the work belongs

One question sits underneath all four patterns, and the corpus speaks to it directly.

Privacy risk is created in the first line, in the ordinary decisions a business makes about products, vendors, data use, AI, marketing and customer operations. Those choices rarely look like privacy decisions while they are being made. If privacy judgement is going to reach them while they can still change, the capability has to sit in the first line as well, close enough to the choice to be present when it is made rather than consulted once it has been.

This is not an argument for privacy holding authority over the business. The business owner makes the call and carries its consequences. It is an argument about position. A practitioner who is there when the purpose is set, the vendor shortlisted and the data fields chosen can put a viable alternative in front of the person who decides. A practitioner who arrives after those things are settled can only attach conditions to what has already been chosen.

The corpus makes the position question visible in a way an organisation chart does not. Section 9 profiles the 17 roles where the market buys privacy work inside the business. What is striking about them is not their seniority but their framing: every one is written as an adviser to the business rather than as a participant in it. The market has located the capability in roughly the right place and then described it as though it sits outside.

Read the four patterns below against that. The question each has to answer is not who can overrule the business. It is whether privacy is close enough to the decision to change it, and whether anyone is obliged to answer when it does.

Pattern 1: Authority concentrated in the central privacy function

What the corpus shows. Compliance / Privacy Governance and DPO / Regulatory Interface together account for 40 of the 101 included roles. Both archetypes describe central-function work: the privacy program is run by a central team, the central team holds the methodology and the artefacts, and the central team is the organisation's main interface with the regulator and with privacy-implicated business decisions.

Authority structure implied. Privacy decision-making authority sits with the central function. Business teams escalate privacy questions to the central function and receive privacy decisions back. The central function has the standing to delay or modify a business decision when the privacy posture requires it; in practice, this standing is exercised through compliance escalation channels rather than direct decision authority.

Strengths. Predictable, auditable, consistent across the organisation. Regulator engagement is concentrated and professional. The privacy program can be evidenced at the executive and board level through the central function's reporting. Easier to defend in incident response: there is a clear "who owns this" answer.

Weaknesses. The central function becomes a bottleneck for privacy decisions inside high-velocity business workflows. Product, data, marketing, and AI teams that need privacy input *during* a decision (rather than as a sign-off after the decision) find the central function operating at a different tempo than the business decision they need to make. The mismatch produces either delayed decisions, late escalations, or business teams making the call without privacy input.

When this pattern fits. Organisations with stable business processes, slower decision tempo at the level where privacy decisions matter, and clear compliance obligations that can be discharged through a program-stewardship model. Public sector, regulated industries with stable product portfolios, mid-size organisations with limited privacy decision volume.

Pattern 2: Authority embedded in the business with central function as advisory

What the corpus shows. Operational Privacy carries 17 included roles. The dominant JD-body framing is *partnership* with business teams rather than *authority over* business teams. Phrases visible across the archetype: "partner with product to embed privacy"; "support business stakeholders in privacy decision-making"; "input into business risk assessments". The central function is structured as an embedded advisory capability; the business holds the decision.

Authority structure implied. Privacy decision-making authority sits with the business team. The privacy practitioner is positioned as expert input, sometimes embedded full-time with a business area, sometimes shared across several areas. The practitioner makes the privacy case; the business decides whether to act on it.

Strengths. Decisions move at business tempo. Privacy considerations enter the decision early because the practitioner is present at the decision moment. Business teams develop their own privacy literacy through working alongside an embedded practitioner. Operational integration is high.

Weaknesses. Presence is not the same as obligation. Nothing in the pattern as advertised requires the business owner to engage with the privacy judgement before deciding, so where the business has high autonomy and low privacy literacy, embedded advisory work can produce the same operational outcome as having no privacy capability in the room: the practitioner is there, the judgement is offered, and the decision proceeds without anyone having to answer it. Where the embedded practitioner's position is contested, the only escalation channel is back to the central function, which then operates as Pattern 1 from a position of having missed the decision moment.

When this pattern fits. Organisations with mature business teams, high cross-functional privacy literacy, and operating cultures that genuinely value embedded advisory work rather than treating it as a permission-seeking step. Software platforms, customer-data-intensive retail, mature digital-product organisations.

Pattern 3: Authority hybridised through statutory officer

What the corpus shows. DPO / Regulatory Interface carries 13 included roles, of which 5 are L4. The statutory officer pattern combines features of Pattern 1 (central function) and Pattern 2 (embedded advisory) with an additional element: the officer may hold delegated decision authority, mainly over information access, that does not require escalation through general management. The officer's standing is statutory rather than organisational.

Authority structure implied. Where a statute or delegation gives the officer decision-making power, typically over information access applications under FOI, RTI or GIPA regimes, those decisions

are made under the delegation and are reviewable externally, by an information commissioner or tribunal, rather than through the line management chain. For most privacy questions, the officer's role is advisory and monitoring, and the officer's standing comes from the designation rather than from a power to decide.

Strengths. The officer has standing that does not depend on general management, and within an access-decision delegation the decision is the officer's and is reviewed externally rather than in line. The designation gives contested privacy questions a named place to go, held by someone whose role exists to make calls under regulatory ambiguity.

Weaknesses. The pattern relies on designations that exist mainly in the public sector. The Privacy Act 1988 does not require most private-sector organisations to appoint a privacy officer. Australian Government agencies must appoint one under the Australian Government Agencies Privacy Code, state public-sector arrangements vary, and statutory decision powers sit mainly in the information access laws. Private-sector organisations choosing this pattern have to construct the officer's standing through internal delegation rather than statutory grant, which weakens the design's advantage.

When this pattern fits. Public sector. Higher education. Regulated industries with statutory privacy obligations beyond the Privacy Act. Private-sector organisations operating under foreign data protection regimes (EU GDPR, equivalent) that require a designated DPO, noting that a GDPR DPO advises and monitors rather than decides.

Pattern 4: Authority distributed through three-lines-of-defence (the absent pattern)

What the corpus shows. Risk / GRC Privacy carries three roles. None describe a clean second-line privacy challenge function. The independent privacy assurance layer the three-lines-of-defence model assumes is essentially absent from the privacy roles advertised in this corpus.

Authority structure implied. The three-lines model distributes authority across three functions: Line 1 (the business and the central privacy team operating controls), Line 2 (the independent privacy risk and assurance function challenging the controls), and Line 3 (internal audit providing periodic independent assurance over both). Each line has its own authority basis: Line 1 owns the operational decision, Line 2 owns the independent challenge of that decision, Line 3 owns the periodic assurance over both.

Where this pattern would be strong if it existed. Large regulated organisations where the three-lines model is the default for other risk types (financial crime, prudential, conduct, technology). The infrastructure is already in place: independent reporting lines, board risk committee oversight, second-line policy ownership. Adding privacy as a fourth or fifth risk type inside the existing model is tractable.

Why the pattern is absent for privacy. Two reasons surface from the corpus.

The first is that privacy second-line work has been historically absorbed into broader risk roles (technology risk, operational risk, conduct risk) rather than treated as a discrete privacy risk discipline. The absorption pattern in Section 11 evidences this, and the MUFG AI-and-privacy-risk leadership case in Section 8 shows the one channel where second-line privacy challenge is being hired at leadership level: bundled with AI risk. The capability exists; it is not labour-market-visible as privacy work.

The second is that many Australian organisations have not historically treated privacy as a risk type at the executive and board level with the same status as financial crime, prudential risk, or technology risk. The escalation to risk committee oversight that would create demand for an independent Line 2 privacy function has been weaker. Where Privacy Act enforcement has tightened over 2024 to 2026, the gap is becoming more visible; the labour-market response has not yet caught up.

One boundary matters for anyone reading Pattern 4 as the fix. Line 2 challenge tests privacy decisions from outside the work; it cannot operationalise privacy on behalf of the business. A Line 2 function built over a weak Line 1 produces independent challenge of decisions no one is equipped to make well. The pattern's value assumes Line 1 operational capability exists to be challenged.

The implication for operating-model design. Organisations that nominally operate a three-lines-of-defence model and list privacy as a risk type may in practice be operating with one-and-a-half lines: Line 1 (privacy and business), partial Line 2 (privacy issues handled through broader risk roles), Line 3 (internal audit periodically covering privacy as one of many topics). The labour-market evidence is consistent with that gap between the documented model and the operational reality; it cannot confirm it, because the operating-model documents are not in the corpus.

What the four patterns combine to show

No single Australian organisation operates a pure version of any one of the four patterns. The advertisements suggest most organisations operate variants of Patterns 1 and 2 in combination, with Pattern 3 where statutory authority is available and Pattern 4 in nominal form on the operating-model document.

The labour-market evidence makes three observations possible.

First, when privacy arrives and who has to answer it are the under-described variables.

Operating-model documents describe reporting lines, functional scope and accountability structures. What they rarely describe is the moment privacy enters a decision, or what the decision owner is obliged to do with the privacy judgement once it arrives. The advertisements are no better: an org chart can show privacy as a central function with clear standing while the JD bodies of every Operational Privacy role describe partnership and advisory framing, and neither document says

when the practitioner is actually in the room. The labour-market evidence is one check on the operating-model document rather than a corrective to it. Both describe intent; what happens at the decision itself is visible only from inside the organisation.

Second, the four patterns are not equally good at getting privacy to a decision while it can still change. That is a different test from asking who can overrule the business, and it reorders them. Pattern 2 puts the practitioner closest to the choice, which is the part of the problem hardest to retrofit, while leaving the weakest obligation on the owner to engage with what the practitioner says. Pattern 1 has clear escalation, but the central function usually sees the decision after the business has shaped it. Pattern 3 gives the officer standing and, within an access-decision delegation, the call itself, which is real but jurisdiction-bound, and does not by itself get privacy into the room any earlier. Pattern 4 tests decisions from outside the work once they have been made.

On that reading the dominant patterns are not the worst-placed ones. Pattern 2 is the closest thing the corpus shows to privacy meeting a decision while credible alternatives are still open, and it is also the most improvable: what it lacks is not authority over the business but any requirement that the owner engage with the judgement before deciding. Pattern 4's absence matters, but it is the absence of a check on the first line rather than the absence of the capability being checked, and a check built over a weak first line does not produce better decisions.

Third, the rebalancing argument this review's proposition makes is, at its base, an argument for shifting the pattern mix. It is not an argument for hiring more privacy people, although that may follow. It is an argument for role designs that put privacy at the decision while it can still change and oblige the owner to answer the judgement: a stronger version of Pattern 2 first, since that is where the corpus already shows the capability sitting; Pattern 3 where statutory authority is available; Pattern 4 as a check on a first line strong enough to be worth checking.

And there is a fourth observation the capability analysis forces. Shifting the pattern mix is itself a change to the operating model, which means somebody has to be capable of making it and authorised to do so. Section 6's audit found no advertisement in the 101 that asks for work of that kind. Nine describe building governance inside conditions the organisation has already fixed, and none carries a Chief Privacy Officer mandate. The patterns above describe how privacy operating models are being run. The corpus is much quieter about who is being recruited to redesign them, which is a different capability and, on the evidence of these advertisements, a rarely specified one. Whether that capability exists inside organisations already, held by people whose roles were never advertised in these terms, is a question the corpus cannot answer.

The labour-market evidence does not show this shift currently happening at scale. It shows the existing pattern mix in place, with AI governance pressure being absorbed into the dominant patterns rather than reshaping them. Whether the pressure ultimately forces the shift, or the patterns absorb the pressure successfully without restructuring, is the operating-model question that 2026 to 2028 will resolve.

A practical close

Three things the report's findings imply for organisations actively designing their privacy operating model in 2026:

1. **Name the pattern you are operating, not the pattern your operating-model document claims.** The labour-market evidence in this report is one tool for surfacing the gap. The hiring patterns of comparable organisations, the language inside your own privacy job descriptions, and the authority your privacy practitioners actually have are others.
2. **If you are operating Pattern 1 or Pattern 2, which the corpus suggests most organisations are, name what each gets wrong and decide whether you have another way to cover it.** For Pattern 1, it is decisions made by the business before the central function sees them. For Pattern 2, it is a decision owner who receives the privacy judgement without ever having to answer it. Pattern 1's fix is timing; Pattern 2's is obligation, and neither is fixed by adding headcount.
3. **Treat the AI governance pattern decision (Section 10, Decision 2 in Section 12) as an operating-model decision, not just a hiring decision.** Where it lands shapes which of the four patterns governs your AI governance work.

Section 15 closes the report by reading the proposition against the corpus's combined evidence.

15. Conclusion

What the corpus shows

This review read 167 Australian privacy and privacy-adjacent job advertisements collected from LinkedIn, SEEK and Indeed over six manual collection sprints between May and June 2026. Of those, 101 were coded as substantively privacy-focused. The included sample distributes across seven role archetypes, with Compliance / Privacy Governance, Legal Privacy Counsel and Operational Privacy as the three largest. The band distribution skews toward L3 operational embedding (50%) with a substantial L4 leadership tail (28%); the bands record what advertisements signal, and Section 6 shows the L4 signal points mostly to Level 3. The AI governance signal reaches across multiple archetypes, with the strongest concentration in the AI Governance Hybrid and senior Legal Privacy Counsel roles.

Around the included sample sits a second sample of 55 near-miss roles: advertisements that surfaced under privacy search terms but did not meet the inclusion threshold. Forty-three of them name privacy as scope inside another function, 26% of everything the searches returned; the other twelve do not name privacy at all. The 43 are the report's clearest evidence that privacy capability in the Australian market disperses systematically into adjacent functions: legal counsel, AI and data leadership, cyber and security leadership, senior compliance and risk leadership, FOI roles, and data governance roles.

What the corpus implies

The corpus is consistent with the proposition that the Australian privacy market has overbuilt compliance posture and underbuilt embedded privacy capability. It does not adjudicate the proposition definitively, since labour-market evidence is one input rather than a complete operating-model picture, but it produces four findings that the proposition's argument depends on.

First, the labour-market shape of embedded privacy supply is small relative to compliance and legal advisory supply. Operational Privacy at 17% of included roles is meaningful but secondary to the 49% combined Compliance and Legal Counsel share. Even within Operational Privacy, L4 leadership hiring is disproportionately project-shaped and the language inside the L3 core tilts toward embedded expertise rather than embedded authority. Where embedded privacy capability is hired, it is hired in volumes and shapes that suggest organisations have not yet rebalanced their operating models toward operational embedding at scale.

Second, independent privacy assurance is almost absent from advertised demand, and the absence is not a sampling accident. Three Risk / GRC Privacy roles among the 101 included, none describing a clean second-line privacy challenge function. The finding held across all six collection sprints, across both collection modes, and across deliberate effort to surface this archetype. The

implication is that the Line 2 challenge layer the three-lines-of-defence operating model assumes is not being hired for as privacy work: either because the model is honoured on paper and absent in practice, or because the Line 2 work is being absorbed into broader risk roles that the corpus does not capture under privacy search terms. Either reading matters for how privacy operating models behave under contested decisions.

Third, no advertisement asks anyone to change how privacy decisions are made. The L4 band carries 28 roles, but its signals point mostly to Risk-Disciplined Judgement: read closely against the framework, none of the 101 advertisements puts the organisation's own decision conditions on the table. Nine come close and each describes building governance inside conditions already settled: a framework, a lifecycle, a standard, designed to fit a model nobody is being asked to change. Three of those are greenfields AI governance builds, where the architecture is spelled out because none exists to inherit. Thirteen of the L4 and L3-L4 roles ask for risk-disciplined judgement on hard questions instead, and twelve are mixed. The corpus contains no Chief Privacy Officer mandate at all, which is the same absence from the other side: no authority advertised, and nothing advertised for authority to be used on. This is a finding about what is written down, not about what organisations hold; the capability may well exist inside organisations in roles never advertised in these terms. My own reading of why, which comes from practice rather than from the corpus, is that organisations asking for senior privacy capability describe Risk-Disciplined Judgement because System Shaping is not yet widely recognised: few have seen what good looks like at that level, so they ask for the most senior capability they can picture.

Fourth, AI governance has not settled as a discipline. The corpus shows three different design patterns operating simultaneously: privacy-led absorption, data-led absorption, and standalone AI governance. None is winning. The non-settlement matters because the design choice an organisation makes between the three shapes which operating model carries AI governance work, who has authority over contested AI decisions, and whether privacy expertise is inside or outside the AI governance leadership track. The labour-market evidence makes the trade-offs between the three patterns legible but does not resolve them.

What this report does not claim

The corpus describes how Australian organisations articulate privacy capability when they hire for it during one collection window. It does not describe what privacy work actually gets done inside organisations once roles are filled. It does not describe the universe of Australian privacy hiring, only the portion visible on those three job boards. It does not measure salary, time-to-fill, hiring intent, or organisational privacy program maturity directly. It is one window into the operating environment, accurate within its scope, partial outside it.

A complete picture of the Australian privacy operating environment in 2026 would integrate this labour-market evidence with practitioner survey data, program maturity benchmarking, regulatory enforcement and breach data, and direct review of organisational operating-model documents. This

review is one contribution. The findings it produces are claims about the labour market and reasoned inferences about the operating-model implications, not unconditional claims about the Australian privacy operating environment as a whole.

The proposition revisited

I entered this review holding a proposition: the Australian privacy market has overbuilt compliance posture and underbuilt embedded privacy capability. The corpus evidence is consistent with that proposition. It is not, by itself, sufficient to confirm it, but the patterns it surfaces (compliance-and-legal dominance, operational capability at meaningful but secondary scale, near-absent independent assurance, unsettled AI governance, systematic absorption of privacy into adjacent functions) are exactly the patterns the proposition predicts would be visible in labour-market evidence.

The proposition does not call for organisations to stop hiring compliance and legal privacy capability. It calls for the operating-model imbalance to be corrected: for organisations to add the operational, assurance and operating-model design capability they have under-built, alongside the compliance and legal capability they already have. The corpus evidences the imbalance. The correction is an organisational design choice; the labour market shows what is currently being chosen by default.

The most important finding of this review, if a single one had to be named, is not in any one of the seven archetypes or six absorption patterns or three AI governance bets. It is that the advertisements describe privacy capability without describing when it reaches a decision or what the business is obliged to do with it. Role designs name the expertise and leave out the two things that determine whether the expertise can matter: timing and obligation. Advertisement language under-describes both for most professions, so this is a finding about how privacy work is articulated rather than a claim that privacy fares worse than other functions; testing that comparison would need a corpus this review does not include.

Set beside the capability finding, that omission has a shape worth naming. The market specifies rules, artefacts, advice and operational embedding readily, and is much less precise about who is expected to change the arrangements that produce them. Read against the argument I make on the [Privacy as a Decision Capability](#) page, the compliance loop reproduces itself through hiring: organisations recruit people to produce, maintain and embed visible privacy work without specifying who can redesign the decision system that produces that work. The loop is not a conspiracy and nobody in it is doing their job badly. Position descriptions are copied from the last hire, recruiters pitch what has filled before, and the requisition settles the operating-model question implicitly because nobody experienced it as an operating-model question.

That reading comes from the framework and the thesis rather than from the corpus. What the corpus contributes is narrower and checkable: the language for redesigning the system is rarely present in the advertisements, and where it is present it is more often attached to AI and data governance than to privacy.

Whether the labour market starts naming timing and obligation is the most useful thing to watch over the next eighteen to twenty-four months. It will not settle the question underneath, which is whether privacy judgement actually changes what an organisation decides. Nothing visible from outside an organisation settles that, and advertisements are the outside. What they can show is whether the design has started to make room for it. This review is a starting reference for that watching.

16. Methodology appendix

16.1 Research question and scope

Research question. When Australian organisations describe what privacy capability they want to purchase, what kind of capability are they actually asking for, and what does the pattern reveal about how the Australian privacy operating environment is being structured in 2026?

Scope. Privacy and privacy-adjacent job advertisements identified on LinkedIn, SEEK and Indeed (Australia) during the collection window 2026-05-11 to 2026-06-02.

Out of scope. Salary distribution, hiring intent forecasting, internal-only roles and roles not advertised on the three job boards searched, role-by-role outcomes after hiring, organisational privacy program maturity measured by means other than advertised role design.

Working proposition (a claim that follows from the thesis and that the corpus can check, not a predetermined finding):

The Australian privacy market has overbuilt compliance posture and underbuilt embedded privacy capability. Organisations have the artefacts. They do not consistently have the embedded capability to make privacy decisions inside operational and product workflows. The gap is most visible where AI governance, data governance and privacy intersect at decision points the standard privacy program framework was not designed to handle.

16.2 Source platform and collection method

Three job boards. LinkedIn (112 advertisements, 63 included), SEEK (46, 29 included) and Indeed (9, 9 included), searched with Australian location filters and posting windows from the same day to the past 30 days. Government job boards, employer careers pages and executive search were not searched; see 16.7.

Manual collection only. No scraping, browser-extension automation, or bulk download against LinkedIn or any other job platform. Every advertisement was identified by human search, opened individually, and its text and metadata captured for downstream coding. This method boundary was set before collection began, recorded in the project's canonical method document, and was not deviated from at any point during the six sprints.

AI tool usage post-collection. AI tools were used only after lawful manual collection for intake normalisation, classification assistance, evidence extraction, and synthesis. No AI tool collected, scraped, or generated any source content.

Employer references. All employer references in this report derive from publicly advertised roles during the collection window. The captured advertisement text is retained as the evidence trail.

Two-mode collection. Six sprints used two distinct query modes:

- Sprints 01 to 04 (broad sweep): query terms included privacy, privacy officer, privacy advisor, privacy counsel, privacy risk, GRC privacy, data governance privacy, data protection, AI governance privacy, responsible AI privacy, DPO, cyber privacy, information governance privacy, compliance manager privacy. Produced 141 raw captures at an aggregate 57% include rate.
- Sprint 05 (focused queries): query terms restricted to Privacy Counsel, Privacy Officer, Privacy Manager, Privacy Compliance, DPO. Produced 21 raw captures at a 76% include rate.
- Sprint 06 (top-up batch): mixed query terms, 5 raw captures at 80% include rate.

The two-mode disclosure is methodologically important: the corpus shape reflects the query strategy. A single-mode (focused-only) corpus would have suppressed the near-miss finding (Section 11). A single-mode (broad-only) corpus would have produced too few direct privacy hits to support archetype-level conclusions for some smaller archetypes.

16.3 Corpus summary

METRIC	VALUE
Raw captures	167
Included	101
Near miss	55
of which privacy named as scope in another function	43
of which privacy not named in the advertisement	12
Excluded (duplicates)	11
Included share	60%
Near-miss share	33%
Absorption share (43 of 167)	26%
Collection sprints	6
Collection window	2026-05-11 to 2026-06-02
Source platforms	LinkedIn (112), SEEK (46), Indeed (9); Australia

16.4 Coding framework

16.4.1 Role archetype typology

Nine archetypes. Seven carry roles in the included sample; one carries the 55 near-misses; one carries no roles.

1. Legal Privacy Counsel. Roles whose primary privacy capability requirement is legal advice on privacy and related data protection matters. Sit in legal teams or law firms. Require current practising certificate or recent in-house counsel experience. Work described is advisory, contract-drafting, regulatory interpretation, incident-response legal advice.

2. Compliance / Privacy Governance. Roles whose primary privacy capability requirement is producing, maintaining and reporting on a privacy compliance program. Policy and standard maintenance, breach and incident registers, training rollouts, control attestation, regulator-facing reporting, ISO 27701 or similar framework alignment.

3. Risk / GRC Privacy. Roles whose primary privacy capability requirement is the independent challenge of privacy controls and operations from a second-line risk or assurance position. Independent reporting lines, board-facing risk reporting on privacy issues, separation between Line 1 and Line 2 in practice.

4. Operational Privacy. Roles whose primary privacy capability requirement is embedding privacy into the operations of the business. Vendor onboarding, product design partnership, data flow uplift, DPIA execution as a business-decision input rather than an artefact for sign-off.

A note on this label. "Operational Privacy" is the coding label the 2026 corpus was classified under, and it is held fixed so the 2027 edition can be compared against this one without re-coding. It is used throughout the report as an archetype name and is capitalised for that reason. It is not the report's term for the capability itself, which is described in plain language as embedded privacy capability. Readers coming from my writing elsewhere should not read the archetype label as a category claim.

5. Data / Technology Privacy. Roles whose primary privacy capability requirement is the technical implementation of privacy controls inside data infrastructure: privacy engineering, data protection technology deployment, platform privacy configuration, technical PETs work, privacy-by-design code review.

6. DPO / Regulatory Interface. Roles whose primary privacy capability requirement is acting as the statutory or designated privacy officer for an organisation, including formal regulator engagement, statutory information access decisions (FOI, GIPA, RTI, PRIS), and decisions taken under delegated authority.

7. AI Governance Hybrid. Roles whose primary capability requirement is the governance of AI and adjacent data activity, with privacy named as a major scope component. The defining feature is that the role's design treats privacy and AI governance as one combined remit rather than two disciplines.

8. Privacy-Adjacent Near Miss. Roles that surfaced under privacy-related search queries but whose privacy scope was insufficient for inclusion. Used for the 55 near-miss classifications. The category covers two situations, separated in Section 11 because they support different claims: 43 roles where privacy is named as scope inside another function, and 12 where the advertisement does not name privacy at all. The second group was verified against the captured JD bodies on 2026-08-19. Nine contain no occurrence of the word "privacy" anywhere; the remaining three contain it only in recruiter boilerplate about how the candidate's own data will be handled during the application. Candidate-privacy boilerplate is a recurring source of false positives in a privacy job search and is worth anticipating in any replication.

9. Other / Mixed. Roles whose scope is genuinely mixed and does not cleanly map to one of the above. No roles in the included sample.

16.4.2 Capability level framework (L1–L4)

L1: Rules awareness. Knows that privacy law exists, can identify issues, supports privacy work without owning it. Effectively absent from the included sample.

L2: Artefact production. Produces and maintains privacy artefacts: PIAs, ROPAs, registers, policies, training material, vendor questionnaires.

L3: Operational embedding. Embeds privacy into the operations of the business: vendor onboarding decisions, product design partnership, data flow uplift, DPIA execution as a business-decision input rather than as an artefact for sign-off.

L4: Operating model and contested judgement. Designs privacy operating models, makes contested judgement calls under regulatory ambiguity, leads frameworks across multiple functions, exercises delegated authority.

Two intermediate bands (L2–L3 and L3–L4) were used where role descriptions genuinely straddled boundaries. They are reported separately, not rolled into the higher level.

Crosswalk to the canonical Privacy Capability Framework

Provenance. These bands were designed and applied in May and June 2026, while the practitioner framework was still developing, and were frozen with the corpus. The framework reached its settled form in August 2026, published as [The Privacy Capability Framework](#), where its fourth level is named System Shaping. The bands code the signals an advertisement gives for each framework level. They were built while the framework was in progress and keep their original codes so the corpus stays reproducible and the 2027 edition remains comparable.

Reading rule. L1 to L4 always mean the advertised signal bands. The practitioner levels are written out in full: Rule Knowledge, Operational Fluency, Risk-Disciplined Judgement, System Shaping.

The framework, set out in [The Privacy Capability Framework](#), describes capability demonstrated across a body of work; the bands record how each level shows up in an advertisement. The table maps each band to the level it signals and how well the signal holds.

FRAMEWORK LEVEL	ADVERTISED SIGNAL BAND	HOW WELL THE SIGNAL HOLDS
Level 1 Rule Knowledge	L1 Rules awareness / support	Closely. Both describe accurate identification of the obligation and disciplined escalation.
Level 2 Operational Fluency	L2 Artefact production / process administration	Closely. Both describe applying professional method to produce credible privacy work.
Level 3 Risk-Disciplined Judgement	L3 Operational embedding / business integration	Partly. The framework level asks whether the practitioner can locate the consequential choice, reason from effects on people, develop a workable alternative and name who owns the unresolved risk. The band signals where the work sits. Embedding is a proxy for access and exposure, not proof of judgement.
Level 4 System Shaping	L4 Operating model / contested judgement / leadership	Weakly; tested in Section 6. The framework level requires shaping the recurring conditions under which decisions are made. The band also picks up contested judgement (Level 3), senior leadership, and authority signals such as an executive or board reporting line (not a capability). Read closely, none of the 28 L4 advertisements signals System Shaping.

The mixed signal is visible in the project's own classifier, whose eight L4 signals mix canonical Level 4 ("designs privacy operating model"; "defines decision rights, governance, escalation or accountability"; "establishes assurance or control effectiveness model"; "sets privacy risk appetite or governance forums") with canonical Level 3 ("challenges business decisions or risk acceptance"), with a reporting line ("reports to executive/board"), and with two ambiguous signals ("leads enterprise-wide privacy strategy"; "leads cross-functional capability").

Consequence for reading this report. The 28 L4 roles cannot be cited as 28 roles requiring System Shaping. Section 6 reports a secondary audit that separates them. The original classifications are unchanged: the original codes remain the corpus of record and the 2026 distribution (50 L3 / 28 L4 / 10 L2 / 6 L2–L3 / 6 L3–L4 / 1 N/A) stands for reproducibility and 2027 comparison. The secondary audit is an additional reading of the same evidence.

Method of the secondary audit. All 28 L4 and 6 L3–L4 roles were re-read against a canonical Level 4 test: evidence of redesigning how privacy decisions are made, defining decision rights or escalation arrangements, shaping the enterprise privacy operating model, changing cross-functional governance routines, building and distributing practitioner capability, altering how Line 1 capability and independent challenge interact, or creating repeatable institutional arrangements instead of personally handling each difficult matter. Senior title, team leadership, board reporting, broad stakeholder engagement, difficult advice and framework maintenance for compliance were treated

as insufficient on their own. All 50 L3 roles were scanned for concealed architectural signals; one used operating-model language and described implementing an existing model rather than designing one, so the L3 band was left undisturbed. The audit was performed by a single reader against the evidence-phrase columns, and inherits the single-coder limitation at 16.7 more sharply than the original coding does, because it is interpretive rather than mechanical.

16.4.3 Task verb taxonomy (16 verbs)

The task-orientation column captures the primary and top secondary verbs the JD body actually used. The 16 verbs split into two groups:

- **Routine / advisory work:** advise, interpret, review, draft, maintain, monitor, report, train.
- **Contested-judgement / operating-model work:** assess, implement, embed, challenge, assure, design, govern, lead.

Train was grouped with the second set at coding and is reported with the first here. Delivering training is routine delivery work, and Section 8's analysis had already been treating it that way; the grouping is presentational and no coded value changed. Within the second group, five verbs carry the signal most cleanly and Section 8 uses them as the marker: *challenge*, *assure*, *govern*, *lead*, and *design* in the specific sense of operating-model design. *Assess*, *implement* and *embed* sit in the second group for coding purposes, but they are not contested-judgement work in themselves, and nothing in the analysis rests on them.

The distinction is described in Section 8 in the discussion of contested-judgement language scarcity.

16.4.4 AI governance signal

Four-level categorical: **Strong** (AI governance is central scope), **Some** (AI in scope, not central), **Mentioned** (boilerplate or single reference), **None**.

16.4.5 Operational privacy signal strength

Four-level categorical: **Strong**, **Moderate-Strong**, **Moderate**, **Weak**. Applied conservatively; the QA guardrail was that it is better to underclaim operational signal than to inflate it.

16.4.6 Capability and operating-model evidence phrases

Verbatim quotes from the JD body that anchor each archetype and capability-level call. They make each classification auditable against the source text.

16.5 Classification and QA process

Each advertisement passed through four stages.

Stage 1: Intake. Manual capture from LinkedIn, SEEK or Indeed of each advertisement's source URL, employer, role title, location, state, work arrangement, seniority, employment type, salary (where shown), reporting line, privacy specificity, inclusion candidacy, and full JD body text.

Stage 2: Classification. AI-assisted coding against the framework above: primary archetype, secondary archetype, L1–L4, confidence, AI signal, operational signal strength, task orientation, top task verbs, capability evidence phrases, operating-model evidence phrases, missing capability notes.

Stage 3: QA review. A separate checklist was applied to each batch. The checklist tested whether the primary archetype is supported by evidence, whether the capability level is too high, whether operating-model signals are being over-read, whether vague terms like "strategic" or "stakeholder engagement" are being over-valued, and whether the role has enough evidence to be counted in the report.

Stage 4: Final coding. Manual review resolution where flagged. Inclusion status and, for near-misses, the near-miss category set.

Manual review case counts per sprint (16 total)

SPRINT	MANUAL REVIEW CASES	OUTCOME
sprint-01	7	All resolved; JD-040 elevated from Near Miss to Include with AI Governance Hybrid primary
sprint-02	5	JD-074, JD-076, JD-083, JD-101, JD-111: all resolved
sprint-03	1	JD-124: resolved as near-miss (standalone AI governance, privacy absent)
sprint-04	1	JD-141: resolved as Compliance/Privacy Governance Include
sprint-05	2	JD-144 (Macquarie Data Protection Mgr): confirmed as near miss, privacy listed as advantageous rather than required; JD-147 (Helia Data and AI Governance Leader): resolved as Include with AI Governance Hybrid primary
sprint-06	0	(none)

16.6 Authority of the corpus

Every count in this report comes from the coded corpus and was cross-checked against it in July and again in October 2026.

16.7 Limitations

Stated in section order, with primary section reference.

Platform coverage (Sections 3, 11). Advertisements were captured from LinkedIn, SEEK and Indeed, with LinkedIn the largest source (112 of 167). Roles advertised only through other channels (internal channels, executive-search-private, government-only platforms such as JobsWA, I Work for NSW, NESA Jobs) are under-represented or absent. The DPO / Regulatory Interface and Data / Technology Privacy archetype counts are likely lower-bound figures.

Collection window dependency (Section 3). The corpus reflects what was on the three job boards between 2026-05-11 and 2026-06-02. Findings about archetype distribution are contingent on that window. A re-run six months later could shift the AI Governance Hybrid share materially.

Two-mode artefact (Section 3). The 57% / 76% / 80% sprint-level include rates reflect the chosen search queries. A different query mix would produce different numbers. The blended 60% include rate is corpus-specific, not market-general.

Description versus reality (Sections 3, 15). The corpus describes how organisations articulate privacy capability when hiring. It does not measure whether the work the advertisement describes actually happens once the role is filled.

Manual judgement in classification (Sections 3, 5, 6, 11). Archetype, capability level, and near-miss classifications involve human judgement against the project's definitions. Coding was performed by a single analyst with AI assistance and a separate QA checklist; there was no independent second coder, so no inter-rater reliability statistic can be reported. The compensating controls are the conservative-classification guardrail, the QA pass (16 manual-review flags surfaced and resolved across six sprints: 7 / 5 / 1 / 1 / 2 / 0 per the table in 16.5), and the verbatim evidence-phrase columns that make every classification auditable against the source text. Residual judgement risk is heaviest on the L3 / L4 boundary in mixed-scope roles and on the Include / Near-miss threshold for general counsel roles with substantive privacy scope.

Near-miss composition (Sections 1, 4, 11). The 55 near-misses are not a single population. 43 name privacy as scope inside another function and carry the absorption argument; 12 do not name privacy anywhere and reflect search-term width rather than role design. The 12 were verified against their captured JD bodies on 2026-08-19: nine contain no occurrence of the word, three contain it only in recruiter boilerplate about candidate data.

Near-miss category consolidation (Section 11). Each near-miss was recorded with a free-text reason for falling short of inclusion, and these were grouped into the six major categories plus the long tail used in Section 11.

Salary data not systematically captured (Section 13). Where salary is shown in an advertisement it was recorded, but the corpus does not produce systematic salary distribution claims. Reward-

differential claims in Section 13 are inferred from role distribution and capability level, not from a salary dataset.

Three-lines-of-defence framework assumption (Section 14). The Pattern 4 framing assumes three-lines is the relevant operating-model baseline. Some agile-aligned organisations explicitly reject three-lines; for those organisations the framing needs adjustment.

Australia-specific (Section 13). Career-strategy implications in Section 13 read the Australian 2026 labour market. UK, US, EU, and NZ privacy markets will read differently.

Time-of-corpus snapshot (Section 10). AI governance is the fastest-moving discipline in the corpus. Findings about the three-way split are accurate for the collection window; a six-month re-run could show meaningful shifts.

16.8 Sources and references

External regulatory and statutory context referenced.

- Privacy Act 1988 (Cth) and 2024 reform amendments (Tranche 1, passed late 2024; Tranche 2 released as the exposure draft Privacy Amendment (Personal Data Protection) Bill 2026 on 31 August 2026, consultation closed 18 September 2026, not introduced as at October 2026).
- OAIC enforcement guidance and notifiable data breach reporting.
- WA Privacy and Responsible Information Sharing (PRIS) Act 2024.
- Vic Privacy and Data Protection Act 2014 (as amended).
- Qld Information Privacy Act 2009 (as amended).
- NSW Privacy and Personal Information Protection Act 1998 (PIPP Act, recent reforms).

Standards and frameworks named in JD-body language across the corpus (number of advertisements mentioning each, across all 167):

- Privacy by Design (30).
- GDPR (24).
- IAPP certification language, CIPP / CIPM / CIPT / AIGP (22).
- Australian Privacy Principles / APP guidelines (17).
- ISO/IEC 27001 (14).
- ISO/IEC 42001, AI Management System (7).
- NIST AI Risk Management Framework (7).
- NSW AI Assurance Framework (5).
- ISO/IEC 27701, Privacy Information Management System (3; the ads reference the 2019 edition's role as the ISO 27001 privacy extension; the standard became standalone and independently

certifiable in its October 2025 revision).

- NIST Privacy Framework (1).
-

16.9 Reproducibility note

Reproducing this review's results requires the coded corpus, the coding framework documented in 16.4, the QA checklist described in 16.5, and a willingness to apply the same conservative classification approach.

The corpus is private to the project and is not redistributed with the report. The methodology is fully documented and can be re-applied to a fresh corpus collected under the same constraints. A reviewer who applies the methodology to a different corpus may produce different specific counts but, if the methodology is followed conservatively, should produce findings consistent with the report's qualitative claims about the Australian privacy labour market.

The single largest reproducibility hazard is the Include / Near-miss threshold. The conservative threshold used in this review produces a relatively pure included sample; a more inclusive threshold (counting general counsel roles with substantive privacy scope as includes) would, on an estimate rather than a re-run, shift roughly 8 to 12 roles from near-miss to include without changing the qualitative pattern. Reproducers should pick a threshold consciously and document it.
