

The Privacy Capability Labour Market Review

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

Fiona Chan

2026
16 October 2026

167 advertisements captured
101 substantively privacy-focused roles
LinkedIn, SEEK and Indeed · 11 May–2 June 2026

Contents

Research findings and organisational implications.

The findings

3–4

The central argument and eight findings

Research scope and role archetypes

5–10

Question, corpus and archetypes

Capability signals and audit

11–15

Original bands, crosswalk, audit and closest cases

How privacy work is described

16–23

Recurring language, gaps and embedded practice

AI governance and adjacent functions

24–29

AI scope, hybrid designs and absorption

Organisational implications

30–35

Organisational decisions and operating models

Practitioner implications and conclusion

36–38

Career routes, capability and closing judgement

Methodology and sources

39–42

Collection, coding, limitations and full research

For executives

Read the findings, then the five organisational choices and the operating-model comparisons.

For practitioners

Read the embedded work and AI sections, then the career routes and capability implications.

For a critical reading

Read the band crosswalk, the audit and its closest case, then the methodology and limitations.

The research boundary

101 included roles within 167 raw captures. Advertisement evidence, not a workforce survey or an evaluation of employers.

The market describes privacy work more clearly than how privacy reaches a decision

Role descriptions name the expertise. They are much less explicit about when privacy enters a decision and who must answer the judgement.

01

49%

Compliance and legal advice dominate

49 of 101 included advertisements sit in Compliance / Privacy Governance or Legal Privacy Counsel.

02

17

Embedded work is a smaller market

17 of 101 roles sit in Operational Privacy. Senior roles are often project-shaped; expertise is clearer than authority.

03

3

Independent challenge barely appears

3 of 101 roles sit in Risk / GRC Privacy. None describes a clean second-line privacy challenge function.

Hiring an expert leaves a design decision unresolved

The corpus is consistent with compliance capability being bought more readily than embedded privacy capability. It does not settle what an organisation needs. Before the next hire, establish which business decisions require privacy input, when that input must arrive and who is obliged to engage with it.

Base: 101 included advertisements; 49% is rounded. The wider collection contains 167 raw captures, including 55 near misses and 11 excluded duplicates. These are observations about advertisements, not a representative workforce survey or an assessment of employers.

The supporting findings qualify the headline pattern

The evidence extends beyond the included privacy-focused roles and requires careful interpretation of senior capability signals.

04 AI governance has three designs

Privacy-led, data-led and standalone functions coexist. The corpus establishes no winning pattern. The choice depends on what existing functions can carry.

05 Senior signals need careful reading

50 roles are coded L3 and 28 L4. L3 signals operational embedding, not proof of Risk-Disciplined Judgement. L4 combines leadership, authority and judgement signals, not proof of System Shaping.

06 L4 does not establish System Shaping

The secondary reading finds no evidenced mandate to change the organisation's own decision conditions under the report's test. This single-reader interpretation is examined on pages 11–15.

07 Privacy disperses into other functions

43 roles name privacy inside adjacent functions: 26% of 167 raw captures, rounded. The other 12 near misses do not support the absorption finding.

08 Timing and authority are under-described

Even roles closest to the business rarely specify when privacy enters a decision or who must engage with the judgement. Operational access does not establish authority to change how the decision is made.

The original classifications remain unchanged. Advertisements cannot establish what organisations already hold internally, how roles operate once filled or which design produces better outcomes.

Hiring describes the capability an organisation thinks it needs

Job advertisements show a design assumption, rather than an operating outcome.

The question

When Australian organisations describe the privacy capability they want to purchase, what kind of capability are they asking for? Hiring volume cannot answer that question. Neither can a maturity score whose framework defines in advance what a developed program looks like.

Why advertisements

An advertisement makes an internal understanding of the work visible to candidates. It describes expected scope, the expertise to be supplied and sometimes the role's place in the organisation. That makes a corpus of advertisements useful for examining what employers articulate as privacy capability.

The proposition being checked

The review entered with a proposition: the Australian privacy market has overbuilt compliance posture and underbuilt embedded privacy capability. The proposition is checked against the corpus, with conservative coding and attention to findings that complicate it, including AI governance and statutory officer work.

What this evidence can establish

The corpus can show whether advertised demand is consistent with that proposition. It cannot show whether privacy judgement changes a live business decision, whether a role performs as described once filled, or whether an employer already holds the missing capability elsewhere.

The findings concern advertised demand. They do not measure operating effectiveness or demonstrate the capability of appointed practitioners.

One collection window, three job boards

167 raw captures were assembled manually in six sprints from 11 May to 2 June 2026.

The observable market

LinkedIn supplied 112 raw captures, SEEK 46 and Indeed 9. Of the 101 included advertisements, 63 came from LinkedIn, 29 from SEEK and 9 from Indeed. Australian location filters and posting windows from the same day to the past 30 days shaped what the searches returned.

Coverage limits

Internal appointments, private executive-search mandates, employer-only careers pages and government-only job boards were outside the search strategy. Some New Zealand or trans-Tasman listings surfaced through Australian filters and were retained, as disclosed in the full review.

Two complementary search modes

The first four sprints used broad queries and returned 141 captures. The fifth used focused privacy-titled queries and returned 21. A final mixed top-up returned 5. Broad searching captured adjacent functions; focused searching added direct privacy-role evidence.

A description of demand

The corpus is neither a census nor a representative workforce survey. It does not measure salary distribution, time-to-fill, team size or operating effectiveness. Search strategy and the short collection window govern the interpretation of every count.

Manual capture preceded AI assistance. No AI tool collected, scraped or generated advertisement content. Employer examples refer to advertisements, not evaluations of employer performance.

The near misses are part of the evidence

The final sample distinguishes substantive privacy roles, adjacent work and duplicate captures.




Figure 1. Final classification: 167 = 101 + 55 + 11. Base: all raw captures. Source: 2026 job-advertisement corpus.

Two distinct near-miss groups

43 name privacy as scope inside another function. The other 12 contain no substantive privacy responsibility: 9 have no occurrence of the word and 3 mention candidate data only in recruiter boilerplate.

Use the right denominator

Archetype and band counts use the 101 included roles. Absorption is 43 of the 167 raw captures, or 26% rounded. Near misses are 55 of 167, or 33% rounded.

Search results are not market shares

The blended 60% inclusion rate is specific to these queries. Broad-only and focused-only searches would produce different corpus shapes.

Almost half the included roles buy compliance or legal advice

Seven populated archetypes classify primary role scope, rather than job titles.

Primary archetype · Included advertisements




Figure 2. Primary archetypes, n = 101. Counts sum to 101; every advertisement has one primary archetype. The combined 49% is rounded. These counts cannot establish how much capability employers already hold or need. Source: 2026 job-advertisement corpus.

The default is advisory

The two largest archetypes describe program stewardship and legal advice. Embedded privacy work is recognisably different, but smaller.

The organisational choice

Decide whether advice on request is sufficient, or whether privacy input needs a defined place in product, vendor and data decisions.

Keep scope and capability separate

The L3 band's 50 roles signal scope, not proof of Risk-Disciplined Judgement. They span several archetypes, distinct from the 17 Operational Privacy roles.

Two advisory routes account for 49 of the 101 roles

Compliance stewardship and legal advice buy different capabilities, even when both sit close to the business.

Compliance / Privacy Governance · 27

The largest archetype describes producing, maintaining and reporting on a privacy program. Policies, training, registers, impact-assessment workflows and regulator reporting are recurring units of work. Advice is framed through the compliance program rather than primarily through legal interpretation.

Representative advertisements

Cochlear's Privacy Manager (JD-026) describes program stewardship. Colonial First State's Manager Financial Crime and Privacy Compliance (JD-067) combines privacy with financial crime. NESAs Manager, Privacy, Data and Governance (JD-143) combines privacy with wider information governance.

Legal Privacy Counsel · 22

The legal archetype describes privacy advice, contract drafting, regulatory interpretation and incident-response legal work. It sits in legal teams or law firms and commonly requires a practising certificate or substantial in-house counsel experience. This is a well-defined route into senior privacy work.

A different design question

CBA's Senior Legal Counsel, Privacy and Data (JD-097), ANZ's Senior Lawyer, Privacy & Data (JD-063) and Optus's Senior Corporate Counsel, Privacy (JD-139) name specialist privacy scope. The question is whether that advisory expertise is enough for the organisation's consequential decisions, or needs an accompanying operational capability.

The combined 49% is rounded from 49 of 101 included advertisements. Concentration describes the corpus; it does not establish that an individual employer has hired the wrong mix.

Operational, statutory and AI work have distinct identities

The next three archetypes hold 44 of the 101 included advertisements.

Operational Privacy · 17

The coding label identifies work whose primary scope is embedding privacy into business operations. Product, vendor, data and marketing decisions are the setting for the work, rather than a separate compliance workflow. QBE (JD-058), Woolworths (JD-068) and ASX (JD-013) provide recognisable examples.

AI Governance Hybrid · 14

These advertisements treat privacy and AI or adjacent data governance as a combined remit. The category holds privacy-led, data-led and standalone-AI identities. AFCA (JD-082), MUFG (JD-114) and Slade (JD-152) show how different the arrangements can be inside one analytical category.

DPO / Regulatory Interface · 13

Statutory or designated officers carry regulator engagement, information-access responsibilities and sometimes delegated decisions. Examples include the WA State Solicitors Office PRIS Manager (JD-093), UQ's Right to Information & Privacy Manager (JD-042) and J&J MedTech's regional DPO (JD-123).

Primary scope governs classification

An archetype is assigned from the advertisement's substantive work, rather than the title alone. Secondary archetypes record genuine overlap. Every included advertisement has one primary archetype, so the distribution sums to 101 without treating mixed work as multiple jobs.

Operational Privacy is the frozen 2026 archetype label. Embedded privacy capability is the report's plain-language term for the capability itself. The AI hybrid category is an analytical instrument, not a prediction about where the discipline will settle.

Senior signals need a second reading

The original L1–L4 coding records what advertisements signal. It cannot establish capability demonstrated by a practitioner.

Original advertised signal band · n = 101




Figure 3. Original classifications are preserved, including both mixed bands. L1 is absent from the included sample. The audit on page 13 tests the 28 L4 and 6 L3–L4 advertisements. Source: 2026 job-advertisement corpus.

What a senior appointment must clarify

Specify whether the role will run an existing arrangement or change the recurring conditions under which privacy decisions are made.

L3 signals access to the work

The operational-embedding band describes where privacy work sits. Exposure to decisions is not proof of **Risk-Disciplined Judgement**, which requires reasoning through the consequential choice.

L4 combines different signals

Operating-model language, contested judgement, leadership and authority can all contribute to L4. They do not individually establish **System Shaping**.

The bands signal role scope; the framework describes capability

The original May–June codes remain unchanged. The practitioner framework reached its settled form in August 2026.

Original advertised band	Canonical practitioner level	How the signal should be read
L1 · Rules awareness	Rule Knowledge	Accurate identification of requirements and disciplined escalation. Advertisements cannot demonstrate practitioner performance.
L2 · Artefact production	Operational Fluency	Application of professional method to produce credible work and run established processes.
L3 · Operational embedding	Risk-Disciplined Judgement	Partial signal. Business access provides exposure; it does not prove judgement, alternatives or unresolved risk ownership.
L4 · Operating model and contested judgement	System Shaping	Weak signal. The band combines operating-model, judgement, leadership and authority language. The audit tests whether it evidences shaping recurring decision conditions.

Canonical progression: know → operate → judge → shape. Levels are cumulative capabilities, not titles. Mixed L2–L3 and L3–L4 bands remain separate. Historical classifications and counts are retained for reproducibility.

What the secondary audit finds

34 advertisements were re-read: 28 coded L4 and 6 coded L3–L4. The audit adds an interpretation; it does not replace the original codes.



0 System Shaping evidenced under the report's test

Figure 4. Secondary audit, n = 34. Counts: 9 + 13 + 12 = 34. A further scan of all 50 L3 roles found no concealed System Shaping signal. Source: 2026 job-advertisement corpus.

What the test asks

Does the advertisement ask the role to shape the recurring conditions under which privacy decisions are made?

Senior advice, a team to lead or a board reporting line is insufficient on its own. The interpretation also asks whether the advertisement evidences a position from which the shaping work is possible.

A finding about advertisements

This is a single-reader interpretive audit. Silence in an advertisement does not show that an organisation lacks the capability, mandate or sponsorship internally. The closest case and a dissenting reading follow.

L4 combines operating-model, contested-judgement, leadership and authority signals; it is not a measure of System Shaping. L3 signals operational embedding, not proof of Risk-Disciplined Judgement.

The closest case turns on placement

Slade Group's AI Governance Lead (JD-152) names much of the work. The advertisement does not establish a position from which to reshape the organisation's own arrangements.

The wording comes close

“including roles, responsibilities, decision rights, and escalation pathways”

Slade Group · AI Governance Lead · JD-152

The review records this operating-model scope. The shaping test also asks what conditions the appointee would be able to change.

What the advertisement leaves unresolved

No reporting line is stated. The title is Lead rather than Head of; the role is recruiter-advertised and the employer is hidden. Those facts do not prove authority is absent, but they leave placement unestablished.

The report's reading

The role describes a greenfields governance build inside conditions the organisation has already settled. The report does not treat that as evidence of changing its own decision conditions.

A dissenting reading remains available

A reader who treats the shaping scope as sufficient, and silence about placement as uninformative, could count one where this report counts none.

What to make explicit in a role

State both the recurring arrangements the appointee can change and the mandate or executive sponsorship that makes the change possible.

An interpretation of advertisement evidence, not an employer assessment. L4 is not a measure of System Shaping; the original classifications are unchanged.

The nine closest cases build within settled conditions

The distinction is between constructing governance and changing the organisation's own recurring decision conditions.

Greenfields builds · Slade, Helia, AFCA

First-build AI governance roles need to name architecture because there is no incumbent blueprint. That produces language resembling System Shaping. The report reads these advertisements as constructing a new function within conditions the organisation has already settled, rather than evidencing placement to alter those conditions.

Framework builds · UNSW, J&J, MUFG, Itoh

These roles establish frameworks, populate stewardship arrangements or own risk governance inside a given model. J&J's regional DPO describes distributing privacy capability through stewards. That is close to the capability test, but the advertisement leaves the surrounding model as the given.

Consultancy · Protiviti, ctrl:cyber

Consultancy advertisements can name the relevant capability while leaving the authority to deploy it with an unidentified client. The advertisement does not establish the mandate or sponsorship needed to change a client organisation's arrangements.

The placement test matters

A scope can resemble shaping without placing the appointee where shaping is possible. The interpretation therefore tests both the work named and the conditions for doing it. This is a single-reader test of advertisement evidence; the capability may exist in organisations without being advertised in these terms.

The original L4 band is not a measure of System Shaping. The nine cases belong to the secondary reading of 34 advertisements. A dissenting reading of Slade can count one shaping case where the report counts none; see the preceding page.

Program stewardship is the default description

The concentration is relative to the report's operating-model proposition, not a finding that compliance work is unnecessary.

Work the market can articulate

Across the 27 Compliance / Privacy Governance roles, advertisements repeatedly name policy maintenance, training, breach response, assessment administration and reporting. These are recognisable deliverables, easy to place in a requisition and easy to evidence through program reporting.

A defensible purchase

A competent compliance program protects people, supports regulatory obligations and supplies methods on which other work depends. The critique concerns the aggregate framing: those deliverables are described more readily than the judgement needed at a consequential business choice.

Senior advice has a different purpose

Specialist counsel can give an organisation a defensible legal position on difficult privacy questions. That is necessary expertise. It does not, by itself, determine how a product team brings the question forward, what the decision owner must answer or how a recurring disagreement changes the organisation's routines.

What the concentration leaves open

An organisation can purchase excellent compliance and legal expertise while leaving operational engagement contingent on the business asking for help. The resulting design may fit a stable, low-volume environment. The corpus cannot show whether it fits any particular employer's actual decision load.

"Over-indexing" is an analytical comparison with the operating-model picture proposed by the review. It is not a population benchmark, a measured optimal staffing ratio or an assessment of individual roles.

Framework references are visible; decision effectiveness is not

Mentions show the language of hiring. They do not measure alignment, certification or operating quality.




Figure 5. Advertisements mentioning each reference, across all 167 raw captures. Categories overlap.
Source: 2026 job-advertisement corpus.

A common language

Frameworks help describe obligations, professional method and program components. They can support inspectable evidence of progress.

Alignment is an incomplete test

A reference in an advertisement does not establish whether privacy enters a vendor, product or AI decision early enough to change it.

Other mentions: NSW AI Assurance Framework 5; ISO/IEC 27701 3; NIST Privacy Framework 1. Counts describe captured language, not current requirements; version context is retained in the complete review.

Advice is named far more often than independent challenge

Task coding records the primary and top secondary verbs, rather than every word used in an advertisement.




Figure 6. Selected top task verbs, n = 101 included advertisements. Multiple coded verbs per advertisement; categories overlap. Source: 2026 job-advertisement corpus.

Wording is an imperfect signal

Design can mean an assessment, a framework or an operating arrangement. A verb alone does not establish the capability required.

Partnership is underspecified

Generic strategic-advisory and stakeholder language can describe both process administration and judgement on contested choices. It does not reliably distinguish the work.

Specify the consequential choice

A useful role description says which decisions the practitioner meets, when they arrive and how unresolved disagreement is handled. “Partner with the business” leaves those questions open.

The taxonomy contains 16 verbs. Zero “assure” occurrences means zero among the coded top task verbs, not proof that the word or assurance activity is absent from every advertisement or organisation.

Three risk roles do not establish a privacy Line 2

The small archetype remains informative because it did not grow across both search modes.

The three partial examples

HSBC's Data Risk and Technology & Cyber Risk Specialist (JD-036) places privacy inside a broader risk remit. Protiviti's Technology Risk and Resilience Manager (JD-016) delivers broader technology-risk work for clients. Plenti's Risk & Compliance Manager (JD-081) combines risk and compliance rather than separating privacy operation from challenge.

The important exception

MUFG's Head of AI & Data Privacy Risk (JD-114) names independent review and challenge of impact assessments. It sits in AI Governance Hybrid because AI risk is its primary identity. Independent challenge can therefore be present outside the Risk / GRC Privacy archetype.

The finding is about three advertisements in the 101 included roles and the scope described in them. It does not establish the absence of internal assurance, wider risk coverage or audit.

Two readings remain possible

Privacy-specific second-line capability may be missing, or it may be supplied through wider technology, operational or conduct risk functions. Advertisement evidence cannot decide between those readings for an employer. Statutory officer work also has distinct authority and should not be treated as equivalent to independent risk challenge.

What a three-lines claim requires

If the operating model relies on independent challenge, identify who supplies it, what privacy competence they need and how they remain separate from the work being challenged. Line 2 complements capable Line 1 practice; it cannot perform operational privacy on the business's behalf.

Privacy engineering is rarely a named hiring category

Five Data / Technology Privacy roles describe technical implementation; only one uses Privacy Engineer as its substantive title.

The visible technical work

BHP's Principal Privacy Engineer (JD-012) is the clearest named example. Avanade's Data Protection Manager (JD-037) describes platform delivery. AustralianSuper's Privacy & Data Protection Technology Project Manager (JD-064) purchases technology-side project management.

Platform and project routes

EPI-USE Labs' Lead Consultant SAP Data Privacy (JD-091) describes platform-specific consulting. ANZ's Privacy Risk Business Analyst (JD-070) combines privacy risk with data-project work. Technical privacy delivery is visible, but its role identities are varied.

Where the category ends

The archetype distinguishes technical controls in data infrastructure from broader business embedding. It includes platform configuration, privacy-enhancing technologies and code-level work where privacy is substantive. Technical roles with smaller privacy scope can fall outside the included threshold.

The implication for sourcing

An organisation looking only for privacy-titled engineers may miss capability advertised through data, security or platform engineering. Five included roles is too small a count to establish the size of the wider technical market. The report evidences limited visibility under this search strategy, rather than absence of the capability.

The same visibility problem applies to privacy assurance career routes through broader risk and audit titles. Neither small archetype count is a workforce estimate.

The 17 operational roles describe a recognisable practice

Their difference is the setting of the work, rather than a distinctly more senior profile.

Work alongside the business

The embedded specialist partners with product, technology, data, customer, marketing and vendor teams. Impact assessments and reviews are inputs to business initiatives. The work is recognisably different from maintaining an assessment workflow for a separate compliance process.

A broad sector spread

The corpus includes retail, wagering, insurance, superannuation, capital-markets infrastructure, legal aid and higher education. Six of the 17 advertisements come through recruiters, compared with 22 of the 101 included roles overall. Client deployment and interim work are part of this route.

Operational Privacy is the original archetype label. The band distribution is retained for reproducibility, with mixed bands separate; it is not a practitioner capability assessment.

Leadership, core and entry work

The original bands divide the archetype into 5 L4, 1 L3–L4, 7 L3, 2 L2–L3 and 2 L2 roles. Those codes describe advertised signals: L3 embedding is not proof of Risk-Disciplined Judgement, and the compound L4 band is not a measure of System Shaping.

AI is unevenly distributed

Within the 17 operational roles, 2 carry Strong AI signal, 5 Some, 1 Mentioned and 9 None. QBE and Sportsbet hold the Strong cases. AI is a defining extension of some senior operational remits, rather than the defining feature of all embedded privacy hiring.

Senior embedded work is often bought as an uplift

The examples show different forms of operational leadership rather than a single permanent career ladder.

QBE · JD-058

The Senior Privacy Lead is a nine-month fixed-term engagement for operational uplift, with responsible AI scope. Its contracting horizon matters: the advertisement purchases a defined period of senior intervention, rather than necessarily a permanent head-of-function arrangement.

UNSW and Victoria Legal Aid

UNSW's Data Privacy Lead (JD-034) is an operational lead rather than principally a statutory officer. Victoria Legal Aid advertises both a Data Protection Lead (JD-032) and a Data Protection Specialist (JD-001), giving a rare view of two tiers in the same operational setting.

Sportsbet · JD-126

The Head of Privacy, Data Enablement and RAI brings three remits together. The design differs from a specialist privacy lead advising a separate AI function. It also leaves the question of how that breadth is resourced and made effective outside what the corpus can measure.

The core and entry routes

ASX (JD-013), Australian Retirement Trust (JD-008) and Endeavour (JD-047) describe embedded specialist practice. Correlate Resources' business analyst (JD-050) purchases data-flow and transparency uplift; Hudson's project manager (JD-045) frames privacy through project delivery.

Project-shaped hiring is consistent with treating uplift as a discrete task. It does not prove an employer has no continuing embedded capability. The secondary audit separately tests whether senior scope evidences System Shaping.

Access to the business leaves timing and obligation unresolved

Every one of the 17 operational advertisements uses advisory or partnership framing.

What the text leaves out

Explicit approve, veto, sign-off, final-decision or decision-making-authority phrasing does not appear in the 17 JD bodies. That absence is not itself a defect: the business owner should hold the substantive decision. The more consequential silence concerns when privacy enters and what the owner must do with the judgement.

Why early involvement matters

A practitioner involved when purpose, vendors and data fields are still open can develop a workable alternative. Once those choices are settled, the same expertise is more likely to add conditions to an existing design. Location near the business helps, but does not guarantee this timing.

What the decision owner must answer

Receiving advice and responding to judgement are different arrangements. A model can require the owner to engage before deciding, record the response to unresolved risk and use an escalation route where needed, while leaving the substantive choice and its consequences with the business.

A boundary on the inference

Advertisements cannot show the organisation's actual routines. Timing and obligation may be settled in policies, delegations or working practice that a role description omits. The finding is that the marketed capability rarely specifies those arrangements, including at the leading edge of embedded hiring.

This is an operating-model interpretation of advertisement language. The review contains no comparator corpus showing that privacy descriptions are less precise than descriptions of other professions.

Three AI governance designs coexist

The corpus shows competing arrangements. It does not establish which produces better outcomes.

01 Privacy-led

Privacy function



Privacy + AI governance

Examples in the corpus
Sportsbet · QBE · MUFG

What it requires

An existing privacy function with embedded capability and capacity to carry AI scope.

Where it can struggle

A compliance-led function can reproduce that approach to AI.

The organisational choice

Can privacy reach AI deployment decisions before the design is settled?

02 Data-led

Data leadership



Data + AI governance

Examples in the corpus
AFCA · Helia · Capricorn

What it requires

Data leadership with regulatory familiarity and judgement capacity for ambiguous AI decisions.

Where it can struggle

A technical function may lack the judgement needed for contested deployment choices.

The organisational choice

Who resolves questions that technical controls cannot answer?

03 Standalone

Dedicated AI function



AI governance

Examples in the corpus
Slade · Profectus

What it requires

An AI portfolio that justifies the function and executive capacity to integrate it.

Where it can struggle

Separate leadership can add an interface without embedding governance in business work.

The organisational choice

Who resolves decisions that cross the privacy, data and AI functions?

Source: 2026 job-advertisement corpus. Conceptual comparison, not measured reporting lines or market shares. Evidence spans the 14 AI Governance Hybrid included roles, other archetypes and near misses. Entry conditions and trade-offs are interpretation, not measured outcomes.

AI scope extends well beyond the hybrid archetype

The 14 AI Governance Hybrid roles are not the denominator for all AI-related privacy work.

AI signal	All raw captures · 167	Included roles · 101	Meaning of the code
Strong	40	28	AI governance is central scope.
Some	19	11	AI is in scope, but not central.
Mentioned	10	4	Boilerplate or a single reference.
None	98	58	No coded AI signal.

Senior scope is a clear concentration

18 of the 34 L4 and L3–L4 roles carry Strong AI signal, compared with 10 of the other 67 included roles. The original senior bands combine scope, judgement, leadership and authority; they do not demonstrate System Shaping.

Several routes carry the work

Strong scope appears in senior counsel, operational privacy and compliance roles as well as the hybrid archetype. It is an extension of different professional identities, rather than one settled AI career category.

Source: 2026 job-advertisement corpus. Table columns have different bases. Raw captures include duplicates and near misses. AI signal records advertised scope; it does not assess practitioner competence or employer outcomes.

The hybrid roles lean five, six and three ways

These are primary analytical leans inside 14 included roles, not exclusive market shares.

Privacy-led · 5 of the 14

Breville (JD-010), University of Sydney (JD-102), MUFG (JD-114) and the two ctrl:cyber roles (JD-065, JD-099) lean toward privacy expertise carrying AI scope. Outside the hybrid archetype, QBE and Sportsbet reinforce the pattern through operational privacy leadership.

Data-led · 6 of the 14

RBA (JD-040), Capricorn (JD-090), Protiviti (JD-060), AFCA (JD-082), Helia (JD-147) and ACIC (JD-111) place AI governance alongside information or data governance. RBA carries Some AI signal; the other examples name stronger scope.

Standalone · 3 of the 14

Profectus (JD-028), Slade (JD-152) and Talent (JD-166) describe AI governance as the primary identity. The near-miss evidence adds IAG's AI governance executive (JD-124); CBA's GenAI product owner (JD-129) and Zip's technology, data and AI director (JD-134) organise AI work through product and operations leadership instead.

No measured winner

The patterns coexist within one collection window and several roles straddle them. The corpus contains no comparative outcomes showing which design works better. Choice therefore turns on the work the organisation needs, the existing functions' capacity and the executive ability to integrate the arrangement.

The near-miss examples are outside the 101-role included denominator. Roles without privacy scope inform the standalone-AI design reading but do not support the privacy absorption count.

43 roles carry privacy inside another function

The other 12 near misses are search false positives, not evidence of privacy responsibility being absorbed.



Figure 7. Near-miss composition: 55 = 43 + 12. Absorption: 43 of all 167 raw captures, 26% rounded. Source: 2026 job-advertisement corpus.

A conservative inclusion boundary

Roles surfaced through privacy searches but lacked enough substantive privacy scope for inclusion. This separates a relatively pure main sample from a broad adjacent sample.

Both samples are needed

The included roles show the work marketed as substantively privacy-focused. The absorption cases show privacy purchased within other professional identities.

The threshold is interpretive

A more inclusive threshold would move some general counsel roles into the main sample. It would change percentages without removing the underlying pattern of privacy in broader remits.

Nine of the 12 false positives contain no “privacy” occurrence; three use it only in candidate-data boilerplate. The captured advertisement bodies were checked individually. Cumulative collection stability is not a time trend in labour-market demand.

The six major categories contain 34 absorption cases

Category size and privacy absorption are different counts. The long tail adds nine absorption cases.

Near-miss category	All near misses	Privacy absorption	No substantive privacy scope
General legal counsel	15	14	1
AI / ML / data leadership	8	5	3
Cyber / security leadership	8	5	3
Senior compliance / risk	6	5	1
FOI / information access	3	3	0
Data governance leadership	3	2	1
Long tail	12	9	3
Total	55	43	12

Figure 8. Consolidated near-miss reasons, base 55. The six major categories total 43 roles, but that is not the same 43 as all absorption cases. The final column is the difference between each category total and its absorption count. Source: 2026 job-advertisement corpus.

Absorption leaves several interfaces to be designed

Privacy within a broader remit can be deliberate and effective. The advertisement rarely resolves how it meets a separate privacy function.

Legal and compliance

General counsel roles distribute privacy across a wider commercial practice. Senior compliance and risk roles place it among other regulatory frameworks. These arrangements can fit a modest privacy workload, but the organisation still needs to determine where specialist judgement comes from when the load becomes consequential.

Cyber and AI

The cyber category contains both named privacy responsibility and complete absence. The AI category similarly mixes absorption with standalone identities that name no privacy work. Integration cannot be inferred from a shared governance label or from the existence of adjacent leaders.

Information access

The included sample shows integrated information-access and privacy officers. The near misses show access functions partnered with separate privacy functions. Both are visible designs; the corpus does not establish which delivers better decisions or coordination.

The long tail

Financial crime, research, records, remediation, legal AI adoption, information governance, research ethics and support roles widen the adjacent picture. These small categories establish breadth, rather than independent quantitative conclusions. Academic research roles do not carry the organisational operating-model argument.

The observation concerns role design, not deficiencies at named employers. Scope below the inclusion threshold can still represent useful privacy work; absence of privacy wording does not establish absence of internal coverage.

Settle how privacy reaches the work

The first three design decisions concern where privacy and AI governance sit, and the decisions they must reach.

01 Where privacy enters business work

Evidence. 49 compliance/legal roles versus 17 Operational Privacy roles in the 101 included roles.

Choice. Whether advice on request is sufficient, or privacy input needs a defined place in product, vendor and data decisions.

02 Where AI governance sits

Evidence. Privacy-led, data-led and standalone designs coexist; the corpus establishes no winning pattern.

Choice. Which function can carry the work, and whether executives can integrate a new function with existing privacy and data governance.

03 How cyber and privacy meet in decisions

Evidence. Senior cyber/security near misses absorb privacy or leave it out of the advertised scope.

Choice. Who has authority on each side of the interface, including during incident response, notification and remediation.

Design implications, not measured organisational outcomes. The corpus surfaces the choices; existing capability, risk exposure and decision volume determine the organisation's answer.

A hire needs an arrangement it can work within

The remaining decisions concern the legal advisory load and the source of independent challenge.

04 How privacy legal work is staffed

Evidence. 22 Legal Privacy Counsel roles; 15 broader legal near misses; 14 name privacy as one practice area.

Choice. Whether the privacy load requires dedicated counsel or can sit within a broader remit. The corpus provides no measured headcount threshold.

05 Where independent challenge comes from

Evidence. 3 Risk / GRC Privacy roles in the 101 included roles; none describes a clean Line 2 function.

Choice. Build privacy-specific Line 2, use broader risk coverage or explicitly accept another model. Challenge complements Line 1 capability.

The compliance-led choice can be right

Low decision volume, stable processes and a competent compliance function can justify the existing model. The corpus cannot prove the market is wrong.

Before assuming a hire will change the operating model, establish who has the mandate or executive sponsorship to change the arrangements the hire will work within.

The compliance-led choice can be the right one

A corpus of advertisements cannot prove that organisations have bought the wrong capability.

The strongest objection

Hiring expresses revealed preference: perhaps organisations have considered embedding and assurance and declined them because they do not need the additional investment. For low privacy-decision volume, stable processes and a competent compliance function, that can be a rational design.

Why preference does not settle the question

The argument is that feedback on privacy operating-model failures is slow, often reaches affected people before the organisation and is rarely attributed to role design. Copied position descriptions and familiar recruitment categories can also perpetuate a pattern without a fresh design decision.

A choice with entry conditions

The five decisions are therefore choices, not universal prescriptions. Sector, data intensity, regulatory exposure, existing capability and AI deployment load matter. The practical test is whether the selected arrangement supports the consequential decisions this organisation actually makes.

Who is able to change the arrangement

The capability to recognise structural patterns may already sit with an experienced practitioner. Acting on it requires mandate or executive sponsorship able to change decision routes, forums, reporting and resources. A senior hire does not automatically supply that placement, and a title alone cannot establish it.

The report's judgement that hiring can reproduce an implicit default is interpretation, not a measured finding about the motives of the employers in this corpus.

Central and embedded models solve different problems

These are interpretive patterns, not mutually exclusive employer classifications.

Central function

Program stewardship and statutory-interface work make methodology, regulatory engagement and escalation visible. A central function can provide consistency and auditable reporting. Its weakness is timing when business workflows move faster than questions reach the central team.

The condition it needs

Stable processes and limited decision volume can make central provision effective. Faster product, data or AI activity needs a reliable way to identify consequential choices before they have been shaped. Additional central headcount alone does not establish that route.

The pattern comparison describes design implications. Advertisements do not validate actual reporting lines, timing, decision rights or the effectiveness of an employer's operating model.

Embedded practice

Practitioners close to product, data, vendor and customer work can meet a choice while alternatives remain viable. The operational advertisements describe that proximity through partnership. Their weakness is the absence of a clear obligation on the owner to respond to the judgement.

The condition it needs

The business retains the substantive choice. The model needs defined engagement points, a response to unresolved risk and an escalation route when judgement is disputed. This strengthens the arrangement already visible in the corpus instead of making privacy the owner of every business decision.

Statutory standing and independent challenge are different

Neither automatically brings privacy into a business choice early enough to change it.

Statutory or designated officer

A statutory officer can bring standing, regulator engagement and defined delegated decisions into one role. Information-access and privacy responsibilities often meet here. The scope depends on the jurisdiction and delegation; a GDPR DPO advises and monitors rather than owning business decisions.

The condition it needs

The officer's formal standing must connect to actual operating decisions. A designation does not supply early access by itself. A private-sector internal delegation is also different from a statutory grant, and those authority bases should not be treated as interchangeable.

Three lines

Line 1 operates the business and controls. Line 2 supplies independent risk challenge. Line 3 supplies periodic independent audit assurance. The corpus scarcely advertises privacy-specific Line 2; broader risk and AI-risk remits offer alternative channels.


The condition it needs

Independent challenge assumes capable operational work to challenge. Building a second line over a weak first line cannot provide the judgement the business lacks. Where the organisation uses another model, it should describe that coverage explicitly rather than rely on a nominal three-lines label.

This interpretation uses three lines as its baseline. Organisations deliberately using a different model require an adjusted comparison. No operating-model documents were examined in the corpus.

Make the moment of engagement explicit

A conceptual decision path translates the report's findings into questions for an operating-model discussion.



Test the route in practice

Trace a recent vendor, product or data-use choice. Establish when privacy entered, which alternatives remained available and what the owner did with the advice. Reporting lines alone cannot answer those questions.

Change recurring conditions

Where the same difficulty repeats, System Shaping asks whether routines, interfaces, capability or escalation need to change. That work requires a mandate or direct executive sponsorship; the business continues to own the substantive choice.

Figure 9. Conceptual synthesis of the operating-model argument, informed by the 2026 job-advertisement corpus. This is not an observed process map, a tested intervention or a new research classification.

The advertised routes offer different kinds of work

The corpus describes demand, not salaries, placement probabilities or the value of one form of practice over another.

Compliance and legal routes

Compliance is the largest archetype, with recognisable program-stewardship work. Legal privacy is a well-defined specialist route, anchored by senior in-house counsel at banks, telecommunications and software employers. Legal qualifications and professional experience limit direct entry to that track.

Statutory officer route

Regulator engagement, delegated decisions and FOI, RTI, GIPA or PRIS experience form a distinct capability profile. Jurisdiction-specific knowledge matters. This route is not interchangeable with corporate embedded practice simply because both use privacy titles.

Embedded and project routes

The 17 operational roles show a genuine path through business-facing privacy work. The senior examples are often fixed-term, interim or uplift-shaped. Entry through data-flow analysis and project delivery starts with a different orientation from entry through register and program administration.

Engineering and assurance routes

The small dedicated archetypes suggest searching wider professional identities. Data, platform and security engineering can carry technical privacy work; technology risk and audit can carry privacy challenge. The corpus cannot establish the size or reward profile of those wider careers.

Australian job-board evidence from one 2026 collection window does not support remuneration rankings or promises about future hiring. All routes represent substantive work; their different visibility is not a ranking of professional value.

Capability development and advertised seniority are different

The framework's progression is cumulative; an appointment title is an incomplete guide to what a practitioner can do.

AI widens several routes

Senior privacy roles often name responsible AI, but the work is carried through legal, operational, compliance and hybrid designs. A practitioner needs to understand the decisions and the organisation's chosen pattern, rather than assume that familiarity with one framework covers all AI governance work.

Exposure supports judgement

Embedded work can provide access to consequential choices. Risk-Disciplined Judgement requires more than proximity: locate the choice, reason from effects on people, test necessity and proportionality, develop alternatives and identify who owns unresolved risk.

From recurring cases to System Shaping

A practitioner who notices the same difficult decision arriving for the same structural reason can begin examining the arrangements that produce it. Changing those arrangements is a different capability from personally handling each case and depends on mandate or executive sponsorship.

Keep the response proportionate

Not every hard privacy question needs an operating-model intervention. A clear rule-based answer or established professional process can be sufficient. Capability lies partly in recognising when the problem is a difficult case and when it is a recurring condition the organisation can change.

These are capability implications, not individual career advice. The canonical levels are Rule Knowledge, Operational Fluency, Risk-Disciplined Judgement and System Shaping; they are not equivalent to the original advertised L1–L4 bands.

The design needs to make room for judgement

The evidence is consistent with the proposition. It cannot establish how decisions work inside organisations.

What the corpus makes visible

Compliance and legal advisory demand dominate the included sample. Embedded work is recognisable but secondary. Privacy-specific independent challenge is rarely named. Adjacent roles distribute privacy into wider remits, while AI governance has several simultaneous organisational homes.

What remains outside the evidence

Internal capability, actual decision timing, the response to privacy advice and comparative operating outcomes require evidence from inside organisations. The thesis about how privacy fails goes beyond what advertisements alone can test. Hiring language supplies one checkable contribution.

What the audit adds

The secondary reading distinguishes governance construction and difficult judgement from System Shaping. Under its interpretive test, the advertisements do not evidence a mandate to reshape the organisation's recurring privacy decision conditions. The single-reader limitation and the Slade dissent remain part of that result.

What a later corpus should watch

A future edition can examine whether advertisements begin specifying early engagement, an owner's obligation to respond, independent challenge and mandate to change recurring arrangements. Stable coding and explicit crosswalks matter for comparison. Changes in advertisement language would show that design has begun making room for judgement, not prove the judgement changes outcomes.

A fuller evidence base would combine this corpus with practitioner surveys, program assessment, enforcement and breach evidence, and direct examination of operating decisions.

The query strategy shapes the corpus

Collection and downstream assistance were separate stages, with a manual-only boundary established before capture.

Broad sweep, then focused queries

Sprints 01–04 used privacy and adjacent terms, including privacy risk, GRC privacy, data governance privacy, responsible AI privacy and information governance privacy. They produced 141 captures at a reported aggregate 57% inclusion rate. Sprint 05 used privacy-titled queries: 21 captures, 76% included. Sprint 06 added 5 captures, 80% included.

Why both modes matter

Focused-only searching would suppress much of the adjacent-role evidence. Broad-only searching would supply fewer direct privacy hits in smaller archetypes. The blended 60% inclusion rate is therefore specific to this two-mode strategy rather than a general probability that a privacy search returns a privacy role.

A four-stage process

Manual intake retained metadata and full JD text. AI-assisted classification applied archetypes, bands, signal strengths, task verbs and evidence phrases. A separate QA checklist tested over-reading and unsupported classifications. Manual resolution then fixed final inclusion status and near-miss reasons.

QA and the historical record

140 captures were accepted without a review flag, 16 flagged cases were resolved or revised, and 11 duplicates were excluded. Historical sprint rows sum to 102 included and 54 near misses because a later reclassification was not restated into the originating row. The final corpus contains 101 included roles and 55 near misses.

The historical sprint table is retained in the complete review. Collection-sprint totals measure corpus assembly, not labour-market growth over time.

Every classification needs an evidence trail

The instrument records what advertisements describe, using conservative judgement and verbatim evidence phrases.

Archetypes and signals

One primary archetype captures substantive role scope; secondary archetypes record genuine overlap. AI signal uses Strong, Some, Mentioned and None. Operational signal uses Strong, Moderate-Strong, Moderate and Weak. These are classifications of role descriptions, not assessments of people appointed into them.

Task orientation

The 16 verbs are advise, interpret, review, draft, maintain, monitor, report, assess, train, implement, embed, challenge, assure, design, govern and lead. Primary and top secondary verbs are recorded. Train is presented with routine work in the final report, although grouped differently at coding; no coded value changed.

Original bands remain frozen

L1 is rules awareness, L2 artefact production, L3 operational embedding and L4 operating model and contested judgement. Mixed bands are separate. L3 does not demonstrate Risk-Disciplined Judgement; L4 also captures leadership and authority and does not directly measure System Shaping.

The secondary reading

All 28 L4 and 6 L3–L4 advertisements were re-read, and all 50 L3 roles scanned. The test looked for shaping recurring decision conditions and sufficient placement to do so. The audit adds interpretation without changing the original codes. It inherits the single-reader limitation more sharply because its boundaries are interpretive.

No independent second coder or inter-rater reliability statistic is available. Conservative thresholds, the separate QA checklist and captured evidence phrases mitigate judgement risk; they do not eliminate it.

The limits govern how far the findings travel

The review's strongest claims concern the captured text. Inferences about organisational operation need other evidence.

Coverage and window

Three job boards and a short window omit internal hiring, private search and some government channels. Platform mix and query choice affect the distribution. The search strategy is likely to under-represent statutory and technical work; those expectations are not statistical corrections to the published counts.

Description versus operation

A role can operate differently from its advertisement. Employers may hold capability in existing staff, consultants or wider functions. Missing wording does not prove missing practice, mandate or assurance. The report contains no outcomes or employer-maturity assessment.

The private corpus is not redistributed with this edition. Reapplying the methodology to a fresh corpus can produce different counts; preservation of method does not guarantee the same findings.

Judgement and thresholds

Archetype, band and inclusion calls use a single analyst with AI assistance. The include/near-miss boundary is a major reproducibility risk. A more inclusive threshold could move roughly 8–12 roles, but that is a sensitivity judgement rather than a second coded dataset.

Scope of the implications

Salary was not systematically captured. Career implications apply to this Australian search context, not other countries. Three-lines comparisons depend on that model being relevant. AI arrangements may move quickly; forecasts and design trade-offs are interpretation, not measured future outcomes.

Sources and references

Research sources, framework references and supporting methodology.

Complete review

The Privacy Capability Labour Market Review · 16 October 2026. The complete document contains expanded role inventories, historical collection tables, coding definitions, the audit reasoning and additional qualifications.

Framework and thesis

[The Privacy Capability Framework](#) sets out the canonical practitioner levels. [Privacy as a Decision Capability](#) supplies the thesis against which the labour-market proposition is checked. The band crosswalk is provided on page 12.

This report summarises the findings of The Privacy Capability Labour Market Review, 16 October 2026. The complete research record contains the expanded evidence and methodology.

Regulatory and statutory references

The review references the Privacy Act 1988, OAIC guidance and breach reporting, WA PRIS Act 2024, Victorian Privacy and Data Protection Act 2014, Queensland Information Privacy Act 2009 and NSW Privacy and Personal Information Protection Act 1998. Regulatory context is reported as at 16 October 2026.

Standards, evidence and authorship

Framework and credential names on page 17 are references appearing in captured advertisements, not verification of compliance. Employer examples trace to JD identifiers retained in the private corpus. Author: Fiona Chan. Collection: 11 May–2 June 2026. Publication date: 16 October 2026.