

One benchmark output, one diagnostic, one Readiness Report

Sample outputs from the Alt-V Operations pillar

Accompanies the Operations briefing · Alt-V Law

Derek Duffy, APAC Director, Alt-V

Every firm, person and figure in these reports is invented. The Benchmark and Diagnostic that follow are what the two free tools produce when run on the Ashcombe Verity scenario. The paid AI & Operations Readiness Report sample is a worked example from the engine's development set, on a different invented firm (Calder Rowan LLP) - included to show the shape of the paid deliverable alongside the free self-run tools.

The scenarios

Ashcombe Verity LLP - UK regional firm with 280 fee-earners, six offices (Manchester, Leeds, Birmingham, Bristol, Edinburgh, London), £62m fees, independent (not part of a consolidator). COO 18 months into post. Managing Partner asking the Executive Committee: "we've bought the AI tools - why isn't it in the P&L?" The firm sits in a transition from UK regional to national; the 2025 lateral programme added 11 partners, and the business-services ratio has quietly ticked above 10% of fees without a plan to bring it back down.

Calder Rowan LLP - the engine's synthetic development firm. Used here because the Readiness Report is a 25-40pp paid deliverable with its own scoping, and the Calder Rowan case was built end-to-end during the engine's development, voice-reviewed, and sits as the engine's golden-master worked example.

ARTEFACT	SCENARIO	READ IN THE SAMPLE
1. Operations Benchmark	Ashcombe Verity LLP	Five of the ten business-services functions, set against the benchmark for the £50-200m revenue band. Finance runs at 2.87% of fees against a 1.8% median (Lagging); HR and Risk and Compliance are Behind; BD and Marketing is in line and IT is ahead of its median. Modelled annual saving: £374K conservative, £701K likely, £1.2m stretch. The three headline levers are AI-driven automation, shared services and process improvement.
2. Operations Maturity Diagnostic	Ashcombe Verity LLP	Level 2 (Functional) overall. Ops function shape, Matter-level economics and Standard processes each score Level 3; Data posture and AI posture score Level 2 and are both marked as gating. The first-quarter roadmap warns that reported realisation will dip 2-4 points in the third quarter, once leakage that was invisible becomes visible.
3. AI & Operations Readiness Report	Calder Rowan LLP (synthetic)	The paid report across seven domains, with band-level scores, direct-peer and published benchmark panels, prioritised recommendations and an ROI roll-up. Worked end-to-end.

How to read them

The **Benchmark** is lead-gen collateral: the firm's own data, scored against the benchmark set, with a saving range and the three levers that fit its gap profile. The point is to produce a specific thing for the first scoping conversation to open on. The **Diagnostic** is a pre-scoping conversation starter: an honest stage read, with the engagement shape named but scope and investment set with the firm. The **Readiness Report** is the paid deliverable at the next layer: a scoped engagement, scored against both a published benchmark and (where the direct-peer cell has enough firms) a direct-peer comparator, with a prioritised roadmap and ROI.

The tools and the Readiness Report are management-consulting instruments. They focus a conversation; they do not substitute for a scoped engagement.

NOTICE · FOCUS · CREATE

OPERATIONS BENCHMARK

Operations Benchmark Report

Ashcombe Verity LLP

UK · revenue band £50–200m

JURISDICTION UK	REVENUE £62.0m	REVENUE BAND £50–200m
FEE-EARNERS 280	TOTAL HEADCOUNT 365	EQUITY PARTNERS 65

For the reader. A modelled comparison of your firm against UK + AU sector benchmarks across the ten business-services functions, using deliberately conservative capture rates. Variance from actual outcomes typically ±40%. Comparison based on 5 of 10 functions (you selected 5).

AT A GLANCE

Annual saving opportunity

Three scenarios — never a single false-precision number. The Likely scenario is the planning anchor.

CONSERVATIVE

£374K

0.6% of revenue

12–18 months

Focused programme on the highest-gap functions

MOST LIKELY

LIKELY

£701K

1.1% of revenue

24–36 months

Well-executed transformation across all 10 functions


STRETCH

£1.2m

2.0% of revenue

36+ months

Top-quartile structural change — shared services, offshoring, AI at scale



Function-by-function heatmap

Where your firm sits against benchmarks selected for your size band — larger firms (£50m revenue and above) — anchored on PwC 2025 (UK Top 100), Macquarie 2024 / ALPMA 2025 / TR-MLS 2025 (AU) and cross-industry references. Cost ratios as % of revenue. **The hotter the cell, the further behind median.** The tag on each row states the confidence of that row's benchmark; rows marked Low or Placeholder are Alt-V estimates triangulated from published sources, not published benchmarks.

	FUNCTION	YOUR FIRM	VS MEDIAN & TOP-Q	MEDIAN	TOP-Q	STATUS	SAVING
1	Finance MEDIUM <i>Finance department staff + systems, excluding treasury/interest. Triangulated: Hackett cross-industry 0.54-1.76% ceiling; PwC 2025 UK Top 10 adjacency.</i>	2.87% 1.59×		1.8%	1.4%	Lagging	£498K
2	HR / People LOW <i>Alt-V estimate triangulated from ALPMA 2025 staffing ratios x salary data. No published law-firm HR cost % of revenue exists.</i>	1.69% 1.13×		1.5%	1.1%	Behind	£90K
3	IT / Technology HIGH <i>Fully-loaded IT: staff + opex + depreciation. PwC 2025 publishes IT opex 3.4% (Top 10) to 4.5% (Top 51-100) excluding IT staff; fully-loaded median calibrated accordingly.</i>	2.77% 0.69×		4%	3.4%	Ahead	—
4	BD & Marketing MEDIUM <i>Marketing + BD staff and spend. PwC 2025 marketing line runs 1.0% (Top 10) to 1.9% (Top 51-100); BD staffing added.</i>	1.97% 1.04×		1.9%	1.4%	In line	—
5	Risk / Compliance / GC LOW <i>Alt-V estimate. Risk/compliance/GC staffing excluding PII premiums (PII benchmarked separately via SRA research).</i>	1.44% 1.20×		1.2%	0.9%	Behind	£113K

Bullet bars: **amber** = your firm · **solid line** = median · **dotted** = top quartile. Status: *Behind* = above median · *Lagging* = >1.5× median. Where you opened the depth panel for a function on Screen 3, factual comparators appear as a footnote row under that function's line.

RECOMMENDATION

Your top three improvement levers

Selected from eleven evidenced levers by where the firm sits furthest behind, ranked by gap-to-median across the functions each addresses. Ten work the cost base; the eleventh (pricing, profitability & legal project management) works the margin side of the same equation.

1

AI-driven automation
priority lever for your firm's gap profile

EVIDENCED IMPACT

8–20% reduction in addressed function(s)

TYPICAL ENGAGEMENT

£100K–£400K

ADDRESSABLE GAP-TO-MEDIAN

~£784K

WHAT THIS ADDRESSES HERE

Secretarial & Admin, Finance, HR / People, KM / Library / PSLs.

EVIDENCE ANCHOR

AU fe-to-secretarial leverage 4.05 → 4.46 in single year (ALPMA 2025); ILTA 2025: 80% of firms using/exploring GenAI

2

Shared services

priority lever for your firm's gap profile

EVIDENCED IMPACT

8–18% reduction in addressed function(s)

TYPICAL ENGAGEMENT

£150K–£500K

ADDRESSABLE GAP-TO-MEDIAN

~£784K

WHAT THIS ADDRESSES HERE

Finance, HR / People, IT / Technology.

EVIDENCE ANCHOR

Cross-industry 20–35%; law-firm-adjusted 8–18% (Hackett / APQC + Alt-V practitioner experience)

3

Process improvement (Lean / RPA)

priority lever for your firm's gap profile

EVIDENCED IMPACT

10–25% reduction in addressed function(s)

TYPICAL ENGAGEMENT

£80K–£300K

ADDRESSABLE GAP-TO-MEDIAN

~£784K

WHAT THIS ADDRESSES HERE

Finance, HR / People.

EVIDENCE ANCHOR

Lean / Six Sigma in closed-process high-volume work; well-evidenced impact bands

WHAT THIS MEANS

For your firm

Where you're lagging most. Finance. These functions are running materially above the typical benchmark for the size band — and represent the largest opportunity in the saving estimate above.

Where you're behind median. HR / People, Risk / Compliance / GC. Above median for your size band, but not in distress. Typical recovery is 12–18 months with a focused programme.

Where you're already ahead. IT / Technology. Defend these positions — they are the result of past operational discipline. Treat the saving estimate as net of what you've already earned.

Partnership economics anchor. Your revenue-per-equity-partner sits at £954K. The PwC 2025 benchmark for UK Top 51–100 shows PEP at £600K (YoY +5.5%), against a fee-earner-to-EP leverage of 9:1. Where your actual PEP tracks materially above this anchor, the business-services conversation is typically about capacity to spend; where it tracks below, about capacity to save — the lever priorities above shift accordingly.

What this estimate is — a modelled comparison of your firm's inputs against UK + AU sector benchmarks, with capture rates set deliberately conservative. Variance from actual outcomes typically ±40%. The estimate is the conversation starter; the Diagnostic engagement is where the precise numbers and the implementation roadmap get built.

METHODOLOGY

How this report was built

The tool takes your firm's revenue, headcount and per-function cost inputs, compares them to a curated set of UK + AU benchmarks, and produces three saving scenarios using deliberately conservative capture rates.

CONSERVATIVE

40% × gap-to-median × functions behind

LIKELY

75% × gap-to-median

STRETCH

75% × gap-to-top-quartile

Confidence rubric. High = anchored on a named published survey (PwC 2025, Macquarie 2024). Medium = single survey or ratio-based triangulation (e.g. ALPMA staffing ratios × salary data). Low = Alt-V consultancy estimate triangulated from published sources — read as directional. Placeholder = no defensible comparator exists; flagged for follow-on research. **Size banding.** Medians are selected per firm-size band (smaller = under £50m revenue; larger = £50m and above) — published inter-band spreads exceed median-to-top-quartile gaps, so one median for all sizes would misstate your position. Benchmarks re-baked 2026-07-06.

Variance. Modelled estimate, ±40% typical variance from actual outcomes. Not a substitute for an Alt-V engagement — the tool generates the conversation, the engagement does the work.

Privacy. No cookies. No individual analytics tracking. Inputs and contact details stay in your browser — nothing is transmitted to Alt-V until the hosted CRM endpoint goes live. Marketing consent (if granted) is captured separately per UK GDPR + Australian Privacy Act.

Figures in the other currency convert at the Alt-V planning rate of £1 = A\$1.90, as at 5 October 2026.

Sources cited

PwC Annual Law Firms' Survey 2025 (34th edition) view source	Derived from PwC 2025 published figures view source
Macquarie Bank 2024 Legal Benchmarking Report view source	Thomson Reuters Institute & Melbourne Law School — Australia State of the Legal Market 2025 view source
Derived from TR/MLS 2025 indirect-spend share view source	ALPMA 2025 AU Legal Industry HR Issues & Salary Survey view source
ILTA 2025 Technology Survey (Executive Summary via BMT) view source	CBRE 2025 Law Firm Benchmarking Survey Highlights view source
CBRE 2026 Global Workplace & Occupancy Insights view source	SRA Research Publications — PII Market for Law Firms view source
APQC Service Sector cuts (cited via KPMG BSR 2022 — second-hand; direct APQC verification pending) Practitioner / internal reference	Alt-V Operational Review Toolkit (practitioner estimate) Practitioner / internal reference
Informal industry consensus (no published source) Practitioner / internal reference	UK Crown Commercial Service / G-Cloud indicative ranges (cell-by-cell verification pending) view source

NEXT STEPS

Where this goes from here

Book a 30-minute call

Walk through these findings with Derek Duffy, APAC Director. No commitment, no follow-up pressure.

derek.duffy@altruisticventures.com — "Benchmark follow-up call"

Recommended: Rapid Sprint

Your profile shows material gaps in a handful of functions — a 4–6 week sprint across your 3–5 largest gaps is the proportionate next step.

Contact us for a quote · 4–6 weeks · scope confirmed at discovery call

Ask Alt-V to draft the business case

Send your inputs and we'll draft a board-ready business case in Word from this benchmark — turnaround within two working days, no commitment beyond delivery.

derek.duffy@altruisticventures.com — "Draft the business case"

Fee is scoped at the discovery call, not by the self-serve tool. Complexity, political sensitivity and data quality each shift scope materially, and none is visible without a conversation.

Also from Alt-V: the *AI Readiness & Value Realisation Diagnostic* (free, 11-dimension maturity read with a partner-level report) and the *Legal Tech Architect* (free AI benefit estimate + tool-stack overlap audit against a 247-firm peer set). If technology or AI showed up in your gap profile above, they are the natural next self-serve step — ask Derek for the links.

ALT-V OPS MATURITY DIAGNOSTIC

Legal Operating Maturity — self-assessment

Ashcombe Verity LLP

FULL REPORT

Your operating maturity read

L2

Functional


L2 · Functional

The firm has competent functions in silos, but no one owns the end-to-end conversion of partner effort into realised revenue. Most Australian mid-market firms sit here. Moving to Structured takes 12–18 months and £370k–£1.2m at ~150 fee-earners — and the wrong first move is buying a "legal ops function"; install matter-level economics and the pricing-and-scoping playbook first, then decide whether a function is needed to run them.

Figures in the other currency convert at the Alt-V planning rate of £1 = A\$1.90, as at 5 October 2026.

SUMMARY

Your five dimensions



Chip shading is relative to the Australian mid-market median (Level 2).

Current state and next move, by dimension

OPS FUNCTION SHAPE

L3 Structured L3 · Structured

A COO or Head of Operations sits at or near the partnership table and owns the operating agenda end-to-end. The posture depends heavily on that one credible operator — lose them without a successor installed and the firm regresses to Functional within twelve months.

Next move: Step up practice group leader calibre so group heads genuinely own a P&L, and move the ops function from reporting into portfolio-level decision support.

MATTER-LEVEL ECONOMICS

L3 Structured L3 · Structured

Matter-level P&L is produced monthly and reviewed quarterly by practice group and by partner. Practice group leaders can name their bottom five matters by margin, and outlier partners are coached rather than tolerated.

Next move: Invest in real-time data infrastructure so WIP, budget consumption and margin are visible per matter as of the last time entry.

STANDARD PROCESSES

L3 Structured L3 · Structured

Pricing, scoping, staffing and WIP review are documented and demonstrably run. Deviations require documented justification, and finance can produce, this week, the matters that departed from the playbook.

Next move: Convert the playbook into matter-type playbooks referencing the firm's own realisation, effort and leverage data, with deviation requiring justification before the fact.

DATA POSTURE

L2 Functional L2 · Functional

A monthly reporting pack is produced by finance and read by few outside finance. The gap between what is installed and what is used routinely is at its widest at this level.

Next move: Get the dashboards read — a monthly rhythm where practice group leaders review the data and decisions trace to it, backed by contemporaneous time capture.

AI POSTURE

L2 Functional L2 · Functional

Point tools have been bought, usage is inconsistent, and there is no ROI measurement. This is the level where AI spend runs highest relative to AI return: the firm is buying leverage-shaped tools when it needs visibility.

Next move: Appoint a named AI owner with a measured use-case shortlist targeting the visibility gap: time capture, billing narratives, matter-P&L data quality.

Your first quarter

First quarter, Level 2 → 3: install the operating spine — a COO or Head of Operations hire with a partnership-visible mandate, a finance uplift for weekly matter-P&L reporting, and scoping of the pricing-and-scoping playbook. Indicative Q1 investment £180k–£340k.

Expect the reported realisation number to dip 2–4 points in the third quarter, when leakage that was invisible becomes visible. Plan the partnership conversation for the dip before it lands — firms that get ambushed by it abandon the roadmap.

AI POSTURE

AI note for your stage

At your stage, AI closes the visibility gap Level 3 requires: contemporaneous time capture, billing narrative drafting, matter-P&L data quality. Effect size is material — 2–4 points of realisation, 10–20 days of lockup. What AI does not do at this stage is change partner behaviour or scope discipline.

Invest in AI capability one operating-maturity level ahead of where the firm runs today — not two, not five.

ENGAGEMENT

How Alt-V can help

SHAPE 1

The Diagnostic

Four to six weeks. An extended diagnostic with the leadership team, targeted interviews, a defined data pack review, and a one-day executive workshop. Output: a written maturity report — current level per dimension, the prioritised interventions with investment and duration ranges, and the two or three moves the firm should not make. Fixed fee. No implementation.

SHAPE 2

The Move

Six to nine months to move the firm up one level. Alt-V leads the design and installation of the operating changes and stays through the first quarter of running rhythm. Includes hire specs and interview panels; excludes the hires themselves.

Scope and investment are set with the firm once the Diagnostic has run and the engagement shape is confirmed.

Send Alt-V three data points: firm size in fee-earners, the level you read yourself at, and the one operating question you most want an honest answer to.

[Book a scoping conversation](#)

METHOD

How this was scored

How this was scored: each dimension is the median of your answered questions in that dimension; your overall stage is the lowest dimension, not the average, because operating maturity is gated by the weakest load-bearing dimension. Skipped questions are excluded; dimensions with fewer than three answers are not scored. This self-assessment is a lens, not the instrument — the Alt-V diagnostic is designed to be run as a facilitated conversation, and it degrades when run as a form. If this read surprised you in either direction, that is the signal to test it properly.

Generated by Alt-V Ops Maturity Diagnostic (v1.0) · Model calibrated 11 Jul 2026 · Alt-V Law

NOTICE · FOCUS · CREATE

© 2026 Derek Duffy. Licensed to Alt-V for authorised use.

AI & Operations Readiness Report

Calder Rowan (fictitious)

Seven dimensions, evidence-checked, peer-benchmarked · 07 October 2026

SECTION 1

Executive summary

The Calder Rowan AI & Operations Readiness Report

Calder Rowan assesses as **Developing** - the second of five bands - across the seven dimensions of this assessment. The reading is not held down by a single failure; it is a firm that has bought sensibly and governed on paper, but has not yet built the measurement, pricing and knowledge disciplines that turn tools into a defensible position.

Three findings matter more than the score:

First, the governance rating carries two red flags that the score alone would hide. On paper, governance assesses at the middle band: a policy exists, a risk partner is accountable, vendor controls were checked at procurement. But the firm practises in New South Wales without a maintained view of court disclosure obligations, and output verification is left to individual judgement. Either alone would force a red flag under our rules; Calder Rowan has both. A firm with 45% of its revenue in insurance litigation is precisely the firm that will meet Supreme Court Practice Note SC Gen 23 the hard way.

Second, pricing exposure is the firm's largest structural risk, and it is unexamined. Seventy per cent of revenue comes through panels and institutional clients - the buyers most likely to ask, within the next two panel cycles, what AI has done to the firm's cost base. The firm has no answer to where AI efficiency goes, no pricing capability, and an untouched leverage model. This dimension scored lowest of the seven (Developing, at the low end), and unlike the others it deteriorates on its own: every month of client-side AI expectation moves the negotiating baseline while the firm's position stays still.

Third, the firm knows less about itself than it answered. One in three requested evidence documents could not be produced. A claimed strength in precedent curation had to be marked down because no inventory exists to support it. Nobody could say what the firm spends on systems renewals falling due in the next 18 months. These are visibility findings: the gap between what the firm believes and what it can show is itself a readiness measure, and it is the cheapest one to close.

The asymmetry runs in the firm's favour on governance and against it on value. Calder Rowan is closer to a clean governance position than the red flags suggest - a verification protocol and a two-page court-obligations register are weeks of work, not a programme. It is further from its adoption story than the licence count suggests: Copilot has been live for ten months across 400 seats with no baseline and no measured benefit, which means the firm has spent a year generating anecdotes.

Phase 1 is three actions, roughly £18-28K (A\$34-53K) of combined external cost if supported, less if done internally: stand up the verification protocol and court-obligations register; baseline and measure the two live AI deployments; and put the "where does AI efficiency go" question to the board with the panel-exposure map this report provides.

On the modelled figures in Section 6, the eight use cases assessed carry an indicative central-case benefit of **£440K (A\$833K) a year**, against net-new tooling of £30-90K (A\$57-171K) - but that central case assumes adoption disciplines the firm does not yet have. At today's measured adoption, the honest figure is the low band: £140K (A\$267K). The distance between those two numbers is the business case for Phase 1.

SECTION 2

Readiness profile



Readiness profile

Dimension	Score	Band	RAG
D1 Governance & regulatory	2.65	Operational	Red (override)
D2 Adoption & value	2.50	Operational	Amber
D3 Data & knowledge	2.43	Developing	Red
D4 Pricing & commercial exposure	1.75	Developing	Red
D5 Operational maturity	2.83	Operational	Amber
D6 Talent model	2.17	Developing	Red
D7 Tech stack coherence	2.32	Developing	Red
Overall	2.38	Developing	

The floor rule did not fire: no single dimension drags the composite. The profile is a plateau, not a spike - which is the harder shape to fix, and the more common one at mid-market.

Evidence received: 8 of 12 requested documents. Don't-know answers: 2 (insurer engagement; renewal-cost visibility).

D5 · Operational Maturity

Operational

Score 2.83

AMBER

Don't-know count: 0 · Evidence received: 8 of 12



CAPABILITY	PROCESS / PEOPLE	TECHNOLOGY / PROCESS	GOVERNANCE / PEOPLE
Financial Management	3	3	4
Technology & Innovation	4	1	4
Client Services Ops	3	2	2
Strategy Support	3	2	3

Amber border = lens ≥ 1.5 below its capability mean (spread flag; named in findings).

SECTION 3

D6 • Talent Model Resilience

Developing

Score 2.17

RED

Don't-know count: 0 · Evidence received: 8 of 12



SECTION 3

D7 · Tech Stack Coherence

Developing

Score 2.32


RED

Don't-know count: 1 · Evidence received: 8 of 12



SECTION 4

Peer comparison



■ firm score | peer marker

Published-benchmark mode: markers reference PwC / TR-MLS / Macquarie survey data or Atlas archetype norms (cited per row). As our anonymised benchmark base grows, subscribers receive a re-cut against direct peers.

Where Calder Rowan sits against its peers

Peer figures in this report are drawn from published UK and Australian survey sources (PwC Annual Law Firms’ Survey 2025; Thomson Reuters/MLS State of the Legal Market 2025; Macquarie Legal Benchmarking 2024) and from the adoption and stack norms in Alt-V’s Legal Tech Atlas. As our anonymised benchmark base grows, subscribers receive a re-cut against direct peers - same-band, same-jurisdiction firms assessed on this same instrument.

Measure	Calder Rowan	Reference point	Reading
AI policy in force	Yes	<10% of AU firms (2025 survey round)	Ahead
Adoption at month 10	25–50% of licensed users	50–70% by year two planning norm	Behind the curve it needs
Realisation	89%	92% intervention threshold	Three points of silent margin loss
Business-services cost posture	In range	18–24% of total cost (mid-tier)	Typical
Leverage	5.3	Rising sector trend	Unexamined, not wrong
Governance red flags	2	-	The gap between paper and practice

The overall picture: a firm that is ahead of its market on artefacts and behind on disciplines. That is the most common mid-market profile in our source material - and the most fixable, because disciplines are cheaper than artefacts.

SECTION 5

Prioritised recommendations

What to do, in what order

Ranked using the assessment's prioritisation model (benefit, strategic alignment, effort, time-to-benefit, risk), then phased. The ranking is a structured input to the partnership's judgement, not a substitute for it.

Phase 1: this quarter (three actions, ~£18–28K / A\$34–53K external cost if supported)

1. Stand up output verification and the court-obligations register. Why now: two red flags at an insurance litigation firm practising in NSW; exposure is live on every matter filed. Done looks like: a tiered one-page protocol in force, a per-court permitted-use register owned by the policy owner, and the PI insurer's position in writing. Effort: 4–6 internal days. Owner: risk partner.

2. Baseline and measure the two live AI deployments. Why now: ten months of Copilot with no measured benefit means the year's spend bought anecdotes; every later decision in this report needs this data. Done looks like: one baselined use case per live tool, adoption on the monthly exco pack, the Lexis+ pilot given a charter and a decision date. Effort: 5–8 internal days. Owner: COO.

3. Put the efficiency question to the board. Why now: 70% panel revenue and the firm's buyers are ahead of it on this question. Done looks like: a board minute adopting a position on where AI efficiency goes, informed by the panel-exposure map (which clients, which matter types, when their next review lands). Effort: 3–4 internal days plus one board session. Owner: managing partner + CFO.

Phase 2: months 4–12

4. Build the pricing response. One structured fixed-fee pilot on a suitable insurance product line, margin measured; a named pricing owner. This is the scope of a dedicated pricing review (typical band £15–35K / A\$30–70K) if the firm wants it accelerated. **5. Precedent curation sprint, insurance and disputes first.** Counted, dated, owned – the precondition for grounding any AI workspace in firm knowledge. **6. Price the renewal calendar and collapse the security overlap.** Answers the don't-know that scoring flagged; typically self-funding. **7. Adopt one change-delivery method** and apply it to actions 4–6. Directly addresses the spread flag in the firm's strongest dimension.

Phase 3: months 12–36

8. Redesign one practice group's junior workload around AI-absorbed tasks, training rebuilt to match; expand on evidence. This is where the talent findings become a programme rather than a worry.

SECTION 6

Indicative ROI

ROI table was not emitted in structured form on this run — the narrative above carries the ranges in prose. Re-run the narrative stage to regenerate.

Indicative ROI

Eight use cases were assessed, selected against the firm's practice mix and current stack. Four are implementable with tools already licensed; four would require net-new tooling.

Use case	Status	Low	Central	High
Legal research memo first draft (Lexis+ AI)	Implementable	£12.8K	£40.8K	£107.1K
Pre-bill review (Aderant)	Implementable	£10.6K	£35.7K	£89.2K
Time-entry narratives (Aderant/BigHand)	Implementable	£10.6K	£35.7K	£89.2K
Firm-wide Q&A assistant (Copilot)	Implementable	£12.8K	£40.8K	£107.1K
E-discovery review (net-new)	Gap	£34.0K	£102.0K	£267.8K
Workplace investigation review (net-new)	Gap	£34.0K	£102.0K	£267.8K
Chronology & fact patterns (net-new)	Gap	£12.8K	£40.8K	£107.1K
KM search & Q&A (net-new)	Gap	£12.8K	£40.8K	£107.1K
Annual total		£140K (A\$267K)	£439K (A\$833K)	£1.14m (A\$2.17m)

Net-new tooling for the four gap use cases (three products): **£30–90K (A\$57–171K)** all-in first-year. Modelled payback on the central case: **under six months** – tooling cost is small relative to the modelled benefit at this firm's scale, and six months is the floor we apply because no implementation lands faster in practice.

Assumptions, stated plainly: benefits assume 50–70% active adoption by year two (central case modelled at 60%); total cost of ownership for AI tooling typically runs 2.5–3× licence cost; figures are modelled magnitudes from banded per-implementation values scaled to firm size, not forecasts. **And the caveat that matters most: at the firm's current measured adoption, the central case is not yet available to it.** The low band – £140K (A\$267K) – is today's honest number. The distance to the central case is closed by the adoption disciplines in Phase 1, which cost a fraction of the gap.

SECTION 7

Methodology and limitations

Methodology & limitations

Seven dimensions scored 1–5 against behavioural anchors drawn from Alt-V's assessment instruments (AI readiness, operational maturity, talent economics, technology architecture); levels reported as Foundational / Developing / Operational / Mature / Leading. Claims at the top two levels require documentary evidence; unevidenced claims cap at the middle level – one such cap applied in this assessment (precedent curation). "Don't know" answers depress the relevant dimension and are reported as visibility findings – two in this assessment. Cross-dimension consistency checks fired no caps. Two risk overrides forced the governance flag to Red notwithstanding its numeric score. Peer comparison ran in published-survey mode; this firm's cell of the Alt-V benchmark base had not reached the 12-observation threshold for direct-peer percentiles. ROI figures use the Alt-V Benefit Calculator model (banded per-implementation values, size-scaled, adoption-discounted; A\$1.90/£ fixed). Intake documents are destroyed or returned within 30 days of delivery; only anonymised scores and coarse ratios are retained, for 24 months, in Australian-hosted storage, under the consent clause in the engagement letter. Limitations: answers are evidence-checked, not audited; the operational dimension samples four of ten capability areas (the full grid is the scope of an operational review); benchmark statements rely on the published sources cited inline.

Generated by AI & Operations Readiness Report (engine 1.0) · 07 October 2026 · Data lock 2026-10-07 · Alt-V Law (Altruistic Ventures Law Advisors)
Prepared for the recipient firm; not for onward distribution.