

Treating operations as the economic engine

What I've built for firms where the strategy is sensible, the AI tools are licensed, and the margin still has not moved

General briefing · Alt-V Law

Derek Duffy, APAC Director, Alt-V

In short. Most mid-market firms treat business services as the cost line to defend and the AI programme as a tooling decision. Both are category errors. Operations is the leverage spine of the firm; without it, strategy stalls and AI buys licences instead of margin. I've built a reading of law firm operations across five dimensions (the Legal Operating Maturity Model), a sideways view across the ten business-services functions (the Business Operating Model), a cross-firm catalogue of the eleven improvement levers that actually move the needle, a per-function menu of the interventions that tend to apply where, a process-mapping and evaluation engine for the one or two processes a firm needs to see at depth, and a four-phase engagement shape that runs the whole thing in twelve weeks. A sample pack is included.

The conversation I keep hearing

The question lands in different wording but it is the same question. "We have a decent strategy. The partners can state it. We have bought the AI tools everyone said to buy. Why isn't it in the P&L?" The partners asking it are rarely wrong about any one of those three things. What is wrong is the assumption that strategy and tool-buying are what moves the margin. The thing between the two - the operating model - is what converts choice into result, and most firms have not designed theirs.

Operations is not back-office cost control. It is the leverage spine. The firm's economics live on it; a firm's AI programme compounds on it; a firm's strategy either runs on it or stalls on it. The operating-model conversation is where the honest answer to that P&L question lives.

What a firm's operating maturity actually is: the LOMM

Alt-V's **Legal Operating Maturity Model** is five stages across five dimensions. The stages run Owner-Operator, Functional, Structured, Managed, Adaptive. The dimensions are ops function shape (who owns the conversion of effort into margin), matter economics visibility (whether the firm can see what each matter, client and partner earns), standard processes (whether processes exist and whether anyone would notice if they stopped being followed), data posture (how data reaches decision-makers), and AI posture (what posture the firm can credibly take toward AI).

The model's one load-bearing design choice is that the overall stage is the **minimum** of the five dimensions, not the average. "A firm is as mature as its weakest load-bearing dimension" is a judgement from the field: process maturity does not compensate for data illiteracy; data visibility does not compensate for a matter economics black hole. Averaging flatters and lets a firm invest in the dimension it already feels competent in. The anti-flattery stance of the Diagnostic is deliberate.

Below the five-dimension read, the full LOMM diagnostic unpacks ten operating dimensions - matter intake and scoping, rate and fee setting, resource allocation and staffing, WIP and billing discipline, realisation review cadence, panel and client management, practice-group planning, technology stewardship, data and reporting, and people development and career paths - with per-level behavioural anchors on each. The twenty-five-question self-run tool in the sample pack is the lighter read; the ninety-minute facilitated version is the one that drives engagement scope.

A separate **Boutique Operating Maturity model** carries four levels (Founder-run → Structured practitioner → Systematised boutique → Institutional boutique) for sub-fifty-fee-earner firms. The levers are different (partner-time economics, client concentration, founder dependency), and the top LOMM level is explicitly not the target for a boutique. The Diagnostic flags boutique respondents and routes them to a facilitated conversation rather than scoring them against a mid-market model.

The typical mid-market firm reads at Stage 2 (Functional) or 3 (Structured) overall. The gated dimension is almost always one of data posture or standard processes. The pattern across the client conversations I have is remarkably consistent: the firm has invested in the first two dimensions (ops function shape and matter economics visibility), under-invested in the next two, and skipped straight to the fifth (AI posture) without the compounding base beneath it. Which is why AI shows up as licence spend, not margin.

Four firm archetypes we see most often

The LOMM read against the Alt-V client set collapses into four recurring archetypes, each with its own next move:

- **Sydney commercial boutique, BOM-1 moving to BOM-2.** Founder-dependent, high realisation, no documented leverage. The next twelve months are about documenting what the founder does without turning the firm into a factory.
- **Brisbane general practice, LOMM Level 1 moving to Level 2.** Three to five partners, mixed practice, pays itself out monthly. The next move is installing a Head of Operations and a reporting layer the ExCo actually reads.
- **Melbourne insurance defence firm, LOMM Level 2 moving to Level 3:** the flagship repeatable Alt-V shape. Panel-driven, high leverage, margin pressure increasing. The move is matter-level economics, the pricing-and-scoping playbook, and shared services; the firm needs it done in nine to twelve months, not eighteen.
- **National regional-large firm, LOMM Level 3 moving to Level 4.** The ops function is credible; the next move is live data in front of decision-makers and the AI posture that compounds on it. This is the firm that bought the AI tools and asks why they are not in the P&L.

The archetypes matter because the move from stage N to stage N+1 is not the same move at every stage, and the right engagement shape differs. The Diagnostic names the gating dimension; the archetype read scopes what to do about it.

The business-services side: the BOM

The **Business Operating Model** view is the sideways cut of the same firm. Ten business-services functions - finance, HR, IT, marketing and BD, risk and compliance, facilities and admin, knowledge, innovation and legal tech, procurement, and the COO/PMO layer - each carry a cost-to-serve ratio against the firm's own fees and a sector benchmark. The Benchmark tool runs that read in the browser.

Three readings are worth more than any individual function's number. First, the **total business-services ratio** - if the firm's cost-to-serve across all ten functions is materially above the midtier benchmark (sector median is 8.5–9.5% of fees for a mid-market firm), there is leverage sitting in the function mix, not in any one line. Second, the **shape of the over-run** - a firm whose over-run is concentrated in finance and IT has a very different problem to a firm whose over-run is concentrated in BD and knowledge. Third, the **composition inside IT** - a firm with IT costs at or below the benchmark but concentrated in licences, with little in architecture or data engineering, is a firm whose AI investment case will not land, because the plumbing is not there.

The eleven improvement levers

The Benchmark output names, from a fixed catalogue, the two or three largest-opportunity levers for the firm's shape. There are eleven levers, and the Benchmark's job is to pick the right ones; the engagement that follows is where the chosen levers are actually pulled.

Ten of the eleven work the cost base. **Organisation design** (delaying, span-of-control, consolidation) is the first lever most mid-market firms have under-played; **AI-driven automation** is the fastest-moving one; **shared services** and **outsourcing and managed services** have proven, if sometimes over-sold, patterns in finance and IT; **offshoring and nearshoring** is the highest-saving lever available but asks the most of the firm's change capacity; **process improvement** (Lean and RPA, cautiously); **technology consolidation**; **spend and procurement optimisation**; **demand management** (workforce planning included); and **benchmarking and diagnostics** as the lightest-touch starting lever. The eleventh - **pricing, profitability and legal project management** - is the margin-side lever: a one-to-three-percentage-point realisation uplift on a mid-market firm is the equivalent of a very large cost-side programme, and the discipline required is lower.

Per-function the intervention menu is specific, not abstract. In **finance**, the current flagship interventions are time-entry AI, pre-billing automation, matter profitability programmes, and partnership-reporting modernisation. In **HR**, HR-ops centralisation with AI, HRIS data quality, the lateral integration model, and the EA pool model migration (the "secretarial workforce evolution" that becomes a leverage story at ratios like 1:4 to 1:4.5). In **IT**, the AI governance framework, service-desk Tier-1 conversational AI, application rationalisation, and the PMS upgrade conversation. In **marketing and BD**, pitch automation and AI drafting, CRM data programmes, and the key account governance model. In **risk and compliance**, AML/KYC document AI, conflicts process automation, and risk-register modernisation. The menu is deeper than this; the point is that the Benchmark output does not stop at "the saving is £700k" - it names where, in which function, under which lever, with what change-management shape.

The process layer: the Process Discovery and Evaluation Engine

Underneath the firm-level read sits a process-level one. For the one or two processes a firm needs to see at depth - matter opening and AML, billing and WIP-to-cash, file closure and archiving are the usual candidates - Alt-V runs the **Process Discovery and Evaluation Engine**. It is a consultant-run tool rather than a web form: transcripts, SOPs, Visio or drawio, swimlane spreadsheets and whiteboard photos go in; a structured process model, a deterministic metrics engine (cycle efficiency, handoff density, rework loops, control coverage, variant count, automation candidacy), a six-dimension process profile (flow efficiency, handoff burden, control integrity, data friction, variation, automation headroom) and a review deliverable come out. Where the firm's PMS or DMS event log is available, the engine reconciles the described process against the actual event stream - surfacing phantom steps, dark activities, forbidden transitions and the volume-duration divergences the written process does not say exist.

The engine is for the one or two processes the firm actually needs this read on; running it against everything would be theatre. The Benchmark names the function with the largest-opportunity lever; the Diagnostic names the operating-model dimension that gates the firm; the engine lets a specific process underneath them be studied at the depth a real intervention needs.

Why AI stalls on a Functional operating model

The AI posture dimension sits at the top of the LOMM for a reason. The posture a firm can credibly take toward AI is a function of the four dimensions beneath it: without matter-level economics, you cannot measure what AI has moved; without standard processes, AI accelerates whichever version of the workflow the lawyer happens to be running today; without a data posture that puts live information in front of decision-makers, you cannot run a credible pilot, let alone scale one. A firm installing Copilot licences on top of a Functional operating model is paying for a capability it has no way to compound.

The usual sequence in a firm that has got this right is not glamorous. The firm has moved its matter-opening and scoping onto a disciplined process. It has an operational data layer, however modest, that the ExCo actually reads. It has a known weekly rhythm by which matter economics make it to the partners who own them. Only then does AI investment compound, because there is something for it to compound against. The firm that skipped the four dimensions has installed the tools, bought the licences, and will not see the margin - it sits at Functional in the ops dimensions and asserts Adaptive in the AI one, and the model reads back the honest answer: the overall maturity is the minimum.

How a full engagement runs

The reference Operational Review is twelve weeks, four phases, three facilitated workshops. **Weeks 1-2** are set-up: sponsor scope, governance (steering, communications), the data request pack. **Weeks 3-5** are diagnose: the Diagnostic and the Benchmark run against the firm's own numbers, interviews conducted against ten pre-written function-specific guides, the maturity assessment completed. Workshop 1 at the end of week 5 tests the findings against the leadership team. **Weeks 5-9** are options and future-state design: the service-delivery decision framework, the lever-selection prioritisation matrix, the target operating model. Workshops 2 and 3 land in this phase. **Weeks 9-12** are business case, board decision and handover: the board pack, the implementation methodology, the post-implementation review calendar for six and twelve months out.

Engagement shapes compose with scope and timing. A **Snapshot Diagnostic** is three to four weeks against three to five functions. A full **Operational Review** is twelve weeks across all ten.

Comprehensive OR + Org Design combined is sixteen to twenty weeks when the firm wants the review and the organisation redesign in the same scoped piece of work. Scope and investment are set with the firm once the Diagnostic has run and the engagement shape is confirmed.

Where this tends to be useful

Three use patterns come up most often.

1. **Before an AI investment case.** Running the Diagnostic - and ideally the Benchmark - before approving a firm-wide AI programme replaces the "which vendor" conversation with the "do we have the operating base this spend needs" conversation. It is often the uncomfortable six weeks that saves the firm the expensive eighteen.
2. **After a lateral programme has shifted the firm.** Growth through laterals quietly reshapes operations - matter mix, systems load, business-services ratios. A firm that has added fifteen per cent to its partnership in two years is a firm whose BOM has moved; the Benchmark against the new footprint is the honest read.
3. **Before a strategy day, as the ops counterpart to the Strategy Diagnostic.** Firms running a strategy day without an operating-model read usually produce a strategy their current operating maturity cannot execute. The pair of reads - strategy and operations - is the sensible opening act.

How much of your current business-services spend do you actually believe is buying leverage that compounds - and how much is just paying for the lights to stay on?

Derek Duffy

APAC Director, Alt-V

The sample Benchmark output, Diagnostic report and Readiness Report that accompany this note use invented firms, people and figures throughout. The tools, the Process Discovery and Evaluation Engine, and the Operational Review methodology are management-consulting instruments; they do not substitute for a scoped engagement.