

The AI & Technology Pillar of the Legal Operating System

What Alt-V can help a law firm see, decide and run differently about AI - and the three tools that read it on your estate

Autumn 2026 briefing · Alt-V Law

Derek Duffy, APAC Director, Alt-V

The AI pillar sits in the LOS as the control layer that cuts across every other pillar. Pricing is where the AI story shows up first on the invoice; Operations is where it shows up in process maps; Talent is where it shows up in leverage; Strategy is where it shows up in the firm's choice of where to compete. The AI pillar's job is to turn piecemeal procurement into a programme - one architecture, one register, one governance frame, one measurement discipline. The three diagnostic tools below read what a firm has today across those four surfaces, and name the one or two moves that would convert spend into capability.

Why most AI programmes stall at the “we said yes to everything” line

Mid-tier firms land in one of five patterns, and the pattern decides the pillar's first move. **Pattern 1: unregulated enthusiasm.** Individual partners procure point tools; the firm has eight to twelve licences, no overlap view, no measurement, no written policy. The partnership feels “ahead” because licence numbers look big, but there is no visibility of what runs on client work. **Pattern 2: pilot-forever.** A sixty-lawyer Copilot pilot runs for eighteen months without a go-decision; adoption plateaus at 23 per cent of invited seats; the pilot owner cannot say whether the firm should scale, pivot or stop. **Pattern 3: governance before adoption.** A committee drafts a policy that bans everything useful; adoption goes underground; the firm now has shadow-IT risk on top of the original risk. **Pattern 4: deferred.** The firm takes a “wait and see” position; two years later, insurer renewal, panel RFPs and client AI questionnaires are each asking questions the firm cannot answer. **Pattern 5: scrap-heap integration.** Historic tools (DMS, PMS, email security, CLM) are not integrated with the AI stack; the newest Copilot or Harvey deployment cannot read where the work actually lives.

The five patterns are not sequential. A firm can be on three of them at once - unregulated enthusiasm in corporate, pilot-forever in disputes, scrap-heap integration everywhere. The pillar's first job is to read which patterns are live, which practice groups each applies to, and whether the governance, measurement and procurement frames are strong enough to hold the resulting plan.

The AI Operating Model: five stages

The AI Operating Model (T05 of the AI Toolkit) is Alt-V's anchor for what “mature” looks like without pretending the maturity ladder matters more than the firm's commercial position. The five stages are **Reactive · Experimenting · Deploying · Governed · Compounding.**

Reactive firms respond to individual requests from partners or clients. No strategy; no budget line; no policy; no measurement. The default position for most UK firms outside the LB50 in early 2026.

Experimenting firms have one or two pilots running without a go-decision frame; the pilots produce opinions rather than data. **Deploying** firms have procured firm-wide tools (Copilot, Harvey, Lexis+AI or similar) and are rolling them out to practice groups; adoption is uneven; the first measurement data is coming in, usually inconclusive. **Governed** firms have the four control artefacts in place: a partnership-endorsed AI strategy, a published AI policy with teeth, a vendor consolidation plan with retirement

decisions landed, and a measurement discipline that reads productivity impact at matter-type level.

Compounding firms have moved past governance into commercial realignment: pricing has moved to AFA on the matter types where AI changes the economics; leverage has been redesigned where AI replaces first-year work; partner compensation reflects AI-augmented delivery. The jump between stages is not time; it is the specific capability move each stage asks for.

The AI Operating Model Position one-pager in the sample pack reads the firm against this five-stage anchor, names the gating dimension, and sets up the three capability moves Phase 2 of an engagement would run.

The four surfaces an AI programme has to control

Surface 1: strategy and sponsorship. Which use cases the firm is pursuing, why, and who owns the outcome. The question is not “which AI tool should we buy” but “which practice group’s economics do we need to change, and over what period”. The surface fails when the strategy is drafted by COO / CIO without partnership endorsement - the strategy exists, but it cannot authorise the leverage or pricing changes that make it real. **Surface 2: architecture and integration.** Which tools the firm runs; how they integrate with the DMS, PMS and client systems; which overlaps exist; which retirements are planned. The surface fails when procurement is practice-group-led without a firm-level register. **Surface 3: governance, risk and policy.** The AI policy, its enforcement route, the training discipline that stands behind it, the client disclosure and incident response position. The surface fails when policy exists on paper but the training records would not survive an insurer’s proposal form or a panel RFP. **Surface 4: value realisation and measurement.** The KPI discipline that reads whether the AI is actually changing matter economics; the governance reporting cadence; the leading indicators that would warn if adoption stalls. The surface fails when the firm cannot say - twelve months after Copilot licence procurement - whether the licence has produced any measurable productivity shift.

The three tools in the pack read against these four surfaces: the **Legal Tech Architect** reads surface 2, the **AI Readiness Diagnostic** reads surfaces 1, 3 and 4, the **AI Governance Scorecard** reads surface 3 at depth.

The twelve-dimension AI readiness diagnostic

The AI Readiness Diagnostic sits on twelve dimensions, built from the AI Readiness Research Synthesis (June 2026) and composed with the NIST Generative AI Risk Management Framework and ISO 42001. The twelve dimensions, with their owning role:

- **Strategy and Leadership** (Managing Partner) - firm-level AI strategy, executive sponsorship, board / exec-committee AI literacy.
- **Partnership Governance and Sponsor Coalition** (Managing Partner / COO) - decision rights, sponsor coalition breadth, practice-group autonomy versus firm-wide standards.
- **Data and Knowledge Management** (Head of KM) - DMS hygiene, taxonomy depth, precedent quality, AI-grounding readiness.
- **Technology and Integration** (CIO) - cloud posture, PMS / DMS integration, AI tool stack, security architecture.
- **Talent and AI Literacy** (HR / KM Lead) - role-based literacy, training depth, champions, recruiting.
- **Culture and Change Readiness** (COO) - prior change track record, change appetite, ADKAR readiness.
- **Operating Model and Process Maturity** (COO) - workflow redesign, 10-20-70 spend split, organisational model.
- **Governance, Risk and Responsible AI** (COLP / Risk Partner) - AI policy, NIST GMMM coverage, ISO 42001 alignment.

- **Regulatory and Professional-Conduct Readiness** (COLP / Risk Partner) - jurisdictional matrix, per-court permitted use, disclosure obligations.
- **Partner Economics and Pricing Realignment** (CFO / Managing Partner) - billable-hour versus AI-efficiency tension, AFA pilots, leverage redesign, partner-comp alignment.
- **Value Realisation and Measurement** (CFO / COO) - KPI discipline, baseline-and-measure protocol, governance reporting cadence.
- **Pilot Discipline** (COO) - Mode-A alternative to the Value Realisation dimension; charter, success criteria, go / no-go decision discipline before procurement. Used when the firm has no live AI deployment yet.

The diagnostic grades each dimension on a five-level scale and uses a floor rule - a firm cannot read as mature on a dimension when the lowest-graded dimension is more than one band below the mean. The floor rule prevents the "strong on strategy, weak on everything else" firm from recording a false-positive reading. The sample pack shows what the output reads like when the floor rule fires on governance and measurement, which is the Enderby Finch position.

The five-cluster AI governance scorecard

The AI Governance Scorecard reads a specific slice of surface 3: whether the firm's AI governance is survivable under external questioning. The scorecard sits on five clusters and twenty questions:

- **Policy and tool control:** whether the firm controls which tools are in use and what may go into them.
- **Training and competence:** whether the people using AI have been taught its failure modes, and whether the records would show it.
- **Output verification:** whether AI output is checked before it is relied on, and whether anyone could prove it.
- **Supervision and accountability:** who actually owns AI oversight; management through structures, or each fee-earner alone.
- **Client disclosure and incident response:** what clients and courts are told, and what happens when something goes wrong.

The overall band is the lowest cluster band; a firm cannot read strong when any cluster is weak. The headline gap is the lowest cluster broken by priority (verification first, then supervision, policy, disclosure, training). The scorecard is deliberately short - twenty questions, roughly fifteen minutes - because the real engagement is not the diagnostic, it is the policy stack and training programme that close the gaps the diagnostic names.

The three tools in the pack

Legal Tech Architect ([calculator/gateway.html](#)) is a five-sub-tool suite that reads the firm's current legal-tech estate and names where consolidation would release spend, where overlaps are producing hidden cost, and where integration gaps are blocking adoption. The five sub-tools are the Benefit calculator (ROI model per use case), the Audit view (vendor overlap map), the Run-cost exposure (what the current stack is costing), the Peer benchmark (how the firm reads against bracket peers), and the Selector (which AI tools would fit the firm's live workflows). The sample pack runs the Audit view for Enderby Finch - the vendor overlap read is where the Legal Tech Atlas does its visible work.

AI Readiness Diagnostic ([AI Readiness Diagnostic/builds/diagnostic/](#)) is the twelve-dimension capability read. Express path produces an indicative report in roughly ten minutes; Full path produces a comprehensive report with evidence capture at roughly forty minutes. The sample runs the Express path on a Mode-B scenario (AI deployed) to show the floor-rule firing on the governance and measurement dimensions.

AI Governance Scorecard (ai_governance_scorecard) is the five-cluster, twenty-question governance read. Jurisdiction-sensitive - the UK edition surfaces the SRA / COLP position; the AU edition surfaces the equivalent regulator and PI insurer frame. The sample runs the UK edition for Enderby Finch; the headline gap is output verification.

The eight-to-sixteen-week engagement shape

Weeks 1 to 4: Strategy and architecture. Running the three diagnostics and the Legal Tech Architect against the firm's actual estate. Partnership consultation on strategic ambition (which practice groups' economics change, over what horizon). Draft AI strategy and vendor consolidation plan. **Weeks 5 to 8: Governance and policy.** Standing up the AI governance committee, drafting the AI policy, building the training programme, writing the client disclosure position. The Policy Stack Configurator (T07 of the AI Toolkit) and the AI Governance Pack sit here. **Weeks 9 to 12: Measurement and realisation.** Baseline KPIs by matter type, measurement protocol, governance reporting cadence, leading indicators dashboard. The AI Program Health Dashboard (T08) sits here. **Weeks 13 to 16: Partner-economics alignment.** Working with the CFO on which matter types move to AFA, which practice groups' leverage redesigns, which partner-comp dimensions reflect AI-augmented delivery. Composes with the Pricing and Partner Pay pillars.

A shorter engagement (eight weeks) covers Phase 1 and Phase 2 only. A full engagement covers all four phases and typically runs twelve to sixteen weeks across two or three partnership touchpoints.

What the pack does not do

The three tools name the firm's current position and the shape of the first intervention. They do not stand up the governance committee; the AI Governance Pack does that. They do not price a specific matter type under AI efficiency; the Pricing Engine does that, driven by the AI readiness read. They do not redesign practice-group leverage; that is a partner-pay conversation driven by the Talent Economics Model. The tools are the first readings - the diagnostic, the architect, the scorecard - on top of which the engagement runs the specific capability builds.

Written up for law-firm partnerships, boards and executives considering their AI programme shape. The AI Operating Model, the twelve-dimension readiness diagnostic and the five-cluster governance scorecard are the Alt-V templates documented in the AI Toolkit (July 2026), the Legal AI Industry Briefing (July 2026) and the Legal AI Refresh (October 2026).

Confidential briefing. Figures and firm names used in the companion sample pack are invented and the firm referenced is fictitious. The methodology is Alt-V's IP, developed on the base of ten years' operational-executive experience in the UK and global legal market.

NOTICE · FOCUS · CREATE