

The AI & Technology Pillar of the Legal Operating System

Two-page summary of the pillar - the control-layer that cuts across every other pillar of the LOS

Autumn 2026 briefing · Alt-V Law

Derek Duffy, APAC Director, Alt-V

The AI pillar is the LOS control layer. It cuts across Pricing (AI is where the discount pressure comes from), Operations (AI is where the process maps compress), Talent (AI is where leverage redesigns), and Strategy (AI is where the firm's choice of where to compete gets real). The pillar's job is to turn piecemeal procurement into a programme - one architecture, one register, one governance frame, one measurement discipline. Three diagnostic tools read what a firm has today across those four surfaces.

The five failure patterns

Mid-tier firms land in one of five. **Unregulated enthusiasm:** point tools acquired piecemeal, no overlap view, no measurement. **Pilot-forever:** eighteen-month pilots with no go-decision. **Governance before adoption:** a policy that bans everything useful; adoption goes underground. **Deferred:** "wait and see" leaves the firm unable to answer insurer, panel and client AI questionnaires. **Scrap-heap integration:** new AI tools cannot read where the work actually lives (DMS, PMS, CLM).

A firm can be on three patterns at once, in different practice groups. The pillar's first job is to read which patterns are live and which practice groups each applies to.

The AI Operating Model: five stages

Reactive · Experimenting · Deploying · Governed · Compounding. Reactive firms respond to individual requests. Experimenting firms have pilots producing opinions rather than data. Deploying firms have procured firm-wide tools and are rolling them out, adoption uneven. Governed firms have the four control artefacts - endorsed strategy, enforced policy, consolidation plan, measurement discipline. Compounding firms have moved past governance into commercial realignment - pricing, leverage and partner comp reflecting AI-augmented delivery.

The jump between stages is not time; it is the specific capability move each stage asks for.

The four surfaces to control

Strategy and sponsorship: which practice groups' economics the firm is changing, over what horizon, with partnership endorsement.

Architecture and integration: which tools, how they integrate with DMS / PMS / client systems, which overlaps and retirements exist.

Governance, risk and policy: AI policy with teeth, training that would survive an insurer's proposal form, incident response position.

Value realisation and measurement: KPI discipline reading matter-economics shift, governance reporting cadence, leading indicators that warn if adoption stalls.

A firm cannot be strong on one surface and weak on the others - the weakest surface sets the ceiling.

The three diagnostic tools

Legal Tech Architect reads surface 2 - five sub-tools (Benefit, Audit, Run-cost, Peer, Selector) over the firm's live estate. The Audit view is the vendor-overlap map that surfaces consolidation targets. **AI Readiness**

Diagnostic reads surfaces 1, 3 and 4 - twelve dimensions on a five-level scale with a floor rule that prevents false-positive maturity readings. **AI Governance Scorecard** reads surface 3 at depth - five clusters, twenty questions, jurisdiction-sensitive (UK / AU), fifteen-minute completion.

The three compose: Architect reads the estate; Diagnostic reads the capability; Scorecard reads whether the firm would survive external questioning.

The twelve-dimension readiness diagnostic and the five-cluster scorecard

AI Readiness Diagnostic sits on twelve dimensions: Strategy and Leadership · Partnership Governance and Sponsor Coalition · Data and Knowledge Management · Technology and Integration · Talent and AI Literacy · Culture and Change Readiness · Operating Model and Process Maturity · Governance, Risk and Responsible AI · Regulatory and Professional-Conduct Readiness · Partner Economics and Pricing Realignment · Value Realisation and Measurement · Pilot Discipline (Mode-A alternative). Each dimension grades on a five-level scale. The floor rule caps the overall band when the lowest dimension is more than one band below the mean - this is where the "strong on strategy, weak on everything else" firm reads accurately rather than aspirationally.

AI Governance Scorecard sits on five clusters: Policy and tool control · Training and competence · Output verification · Supervision and accountability · Client disclosure and incident response. Twenty questions in total, four per cluster. 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 eight-to-sixteen-week engagement

Weeks 1 to 4 land strategy and architecture. Weeks 5 to 8 land governance and policy, with the Policy Stack Configurator (T07) and the AI Governance Pack. Weeks 9 to 12 land measurement and realisation, with the AI Program Health Dashboard (T08). Weeks 13 to 16 land partner-economics alignment, composing with the Pricing and Partner Pay pillars. A shorter eight-week engagement covers Phase 1 and Phase 2 only.

Written up for law-firm partnerships, boards and executives considering their AI programme shape. Full briefing with worked example at [briefing_full](#) in the pack; sample reads of the three tools follow the Operating Model Position companion.

Confidential briefing. Figures and firm names used in the companion sample pack are invented and the firm referenced is fictitious.