

When legal AI starts to act

Why corporate legal orchestration will define the next generation of legal leadership

Thomson Reuters™

Corporate legal departments have spent the past two years investigating and testing how to adopt AI. A more consequential question is now here: what happens when AI moves from helping professionals with individual tasks to acting across the workflows, systems, and decisions through which the enterprise operates?

That shift is already exposing an uncomfortable gap. Ask a general counsel (GCs) whether Legal makes a significant contribution to organisational objectives and 86% will say yes. Ask the C-suite and only 17% agree; 42% say Legal contributes little or nothing at all. Agentic AI arrives at precisely the moment Legal must both transform more quickly and to prove its value to the business more clearly.

General Counsel assessment of Legal's strategic contribution

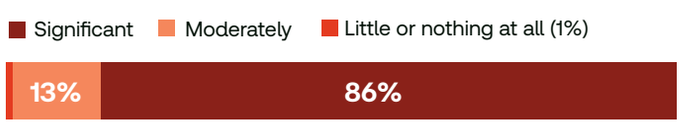

C-suite assessment of Legal's strategic contribution

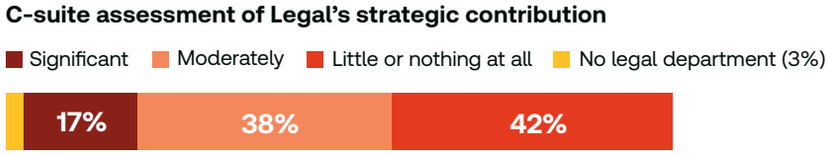

Data source is Thomson Reuters Institute, 2026 State of the Corporate Law Department.

The first wave of AI put powerful tools into professionals' hands. The next wave introduces agents that can plan, retrieve, coordinate, and act across systems. They will not stay inside Legal, and the questions they raise — who authorised this action, what information did it rely on, who answers for it — cannot be settled by Legal alone.

This creates a short window for legal leaders to get ahead of the operating-model questions before technology choices harden into fragmented workflows, duplicated platforms, and unclear accountability.

The next legal AI challenge is orchestration: coordinating trusted legal intelligence, enterprise context, AI agents, workflows, and human judgement around business outcomes, with governance built into the way work is performed. The legal teams that address this now will guide and manage the risks their businesses are introducing with agentic AI. Those that wait may find that the business has picked an operating model without their input or visibility.

Highlights

Disconnected tools are not a mere inconvenience. They are a structural bottleneck, and they make it more difficult for you and your staff to meet your goals.
Consider this:

- AI is moving faster than most legal operating models. Decisions made now will determine where agents can act, what outcomes are tracked, and whether Legal can defend the answers given.
- More AI will not solve an operating-model problem. Workflow, data, skills, decision rights, and change management have to move with the technology.
- The governance question is changing. The issue is no longer only who can use AI, but what AI itself may be authorised to do.
- Legal has an opportunity to lead. Orchestration can keep judgement defensible, connect AI to enterprise outcomes, and give the wider C-suite a reason to engage now.

The risk of waiting: AI is moving faster than Legal's operating model

The Thomson Reuters® Institute describes the work of a corporate legal department as four semi-independent goals: to be effective, efficient, enabling, and protective. Priorities shift over time, with technology now the fastest-rising area of concern. The share of General Counsel (GCs) citing it as a strategic priority doubled from 14% to 28% in a year, and 86% of those technology comments explicitly mention AI. The Institute describes AI as moving from implementation to strategic integration. But the research also shows an uneven pattern. Most departments describe unlocking capacity, but fewer describe deploying it through strategic partnering, business integration, or revenue-linked work. Without closer legal ties to the business, the C-suite is not likely to change its opinions about Legal's role in supporting business goals.

Hours saved on repetitive, low-value work is a number Legal understands. But that number can be uncomfortably close to "potential headcount cuts" when a CFO sees Legal as purely a cost centre. Capacity is clearly a major pain point: nearly half of GCs (48%) name staffing and resource constraints as the top barrier to delivering more value, while only 8% name technology or process limitations. Capacity pressure is what makes agents attractive, but capacity alone will not close the visibility gap with top management.

Top GC capacity pain points



Source: Thomson Reuters Institute, 2026 AI in Professional Services Report

The danger is not that Legal moves too slowly to buy AI; it is that the enterprise moves quickly enough to embed agents while roles, workflows, data, and decision rights remain designed for a pre-agentic world.

Agentic AI raises the stakes in both directions. It offers the potential to both speed up and streamline repetitive, low-value workflows, while also enforcing consistency. But it could also become a black box that produces legal work that can't be audited or traced to a line of reasoning, and that leads senior business leaders to believe everything is fine because "the AI said so". As the State of the Corporate Legal Department report put it, focusing solely on efficiency "could also be a double-edged sword," inviting further resource cuts when departments show only what AI-enabled tasks have accomplished.

Instead, a legal team must deliberately implement AI. A useful framework is the GCO 2030 vision, first explained in an article by Thomson Reuters Chief Legal Officer Norie Campbell and Thomson Reuters Institute Senior Director Elizabeth Duffy. GCO 2030 provides five archetypes for AI-empowered legal teams, each model with a different focus to match teams' varying goals. Departments must choose an operating model and work to implement it across a team, rather than grasping at random solutions to handle disconnected tasks.

Can buying more technology close the gap?

There is a familiar pattern in major technology shifts: companies initially use the new technology to simply speed up existing ways of working. The breakthrough comes later, when the operating model is redesigned around what the technology makes possible. Agentic AI creates that same choice for corporate legal. The danger is not that Legal moves too slowly to buy AI; it is that the enterprise moves quickly enough to embed agents while roles, workflows, data, and decision rights remain designed for a pre-agentic world.

Too much of legal AI is still at the speed-up stage. The Future of Professionals Report found that among professionals whose organisation has a named AI strategy, 35% say day-to-day practice does not match it; among corporate respondents, that rises to 44%. The reported causes — missing tools, insufficient training, and a lack of shared understanding — are classic change-management failures, and the problem is one of leadership and organisational change that no software licence alone will solve.

The governance perimeter is also porous. Some 34% of professionals overall, and 36% of corporate professionals, report using AI tools their employer has not sanctioned. People route around a strategy that does not help them get work done.

The wider market is converging on the same diagnosis. Gartner has predicted that over 40% of agentic AI projects will be cancelled by the end of 2027, citing escalating costs, unclear business value, and inadequate risk controls. McKinsey argues that value comes only when the operating model is redesigned, and BCG that agentic AI success is predominantly a people and change-management challenge.

Unsanctioned tools, stalled projects, and untethered strategy can look like separate problems. They point to the same underlying issue: AI being deployed faster than the operating model around it. That is why the next 12 to 24 months matter. Decisions about architecture, permissions, workflow ownership, skills, and measures of value made now will be expensive to unwind later.

The strategic question shifts from 'Where can we deploy AI?' to 'How should legal work, responsibility, and decision-making be designed when AI becomes an active participant in the workflow?'

The inflection point: From assistance to delegated action

An agent does not simply answer a prompt. Depending on its design and permissions, it can break a goal into steps, retrieve information, interact with other systems, initiate actions, and keep working towards an outcome.

Consider the difference. A generative AI (GenAI) assistant that summarises a contract is a productivity tool; a professional uses it, reviews the output, and decides. A system that retrieves the relevant agreement, applies approved legal and company knowledge, identifies deviations, routes routine issues, escalates high-risk exceptions, starts an approval workflow, and records the decision is taking part in the operating process itself. In the first case, a person uses a tool. In the second, the tool is driving the process and routing to humans for input.

That distinction is why adopting AI tools is not the same as transforming the function. The GCO 2030 model is direct about it: to move beyond incremental gains, legal departments must redesign workflows, improve data infrastructure, invest in training, and hire for adaptability and technical literacy. Greater autonomy adds questions of accountability, oversight, and human responsibility on top. The strategic question shifts from "Where can we deploy AI?" to "How should legal work, responsibility, and decision-making be designed when AI becomes an active participant in the workflow?"

The 2030 legal department is being designed now

As we noted above, the GCO 2030 Project examines the pressure from AI and provides five archetypes for an AI-transformed legal department. Each archetype is a model for a particular kind of legal team and work: some large organisations may implement multiple models for different teams, but most will want to pick one for focus.

Your department's best-fit model will depend on where these fall on your priority list:

- High-volume automation
- Better strategic advising
- Stronger business partnership
- Business self-service within legal guardrails
- Smarter collaboration with outside counsel
- Global leverage across teams, languages, and time zones

Every GC faces constraints of time, budget, and clarity. That's why it's essential to pick the right path for their organisation. But all can draw on the experience of our General Counsel Office as it has moved from experimentation towards real transformation.

Corporate legal orchestration is the connective capability that makes a legal department's chosen operating model deliverable.

Read through an orchestration lens, the five archetypes share a dependency. Automation at volume needs defined boundaries. Strategic advising needs authoritative intelligence and enterprise context. Business partnering that lets other functions move without waiting for Legal needs delegated authority those functions can trust. Collaboration with outside counsel needs visibility into how AI is being used. Global leverage needs the same standards to hold across jurisdictions, languages, and systems. The archetypes describe where Legal wants to go. Orchestration describes what must be in place to get there safely, and in a way the rest of the enterprise can see. Choosing an archetype is a leadership decision. Orchestration makes it deliverable.

Corporate legal orchestration: Five tests of readiness

Corporate legal orchestration is the connective capability that makes those choices deliverable. It is the coordination of trusted legal intelligence, enterprise knowledge and data, AI and agents, workflows and business systems, and human expertise around end-to-end legal and business outcomes.

The future enterprise will not run on one AI system — and that future is already taking shape. Specialist legal AI will coexist with enterprise AI. Business functions will deploy their own agents. Outside counsel are already choosing their own technology tools, and GCs have limited visibility there: two-thirds (67%) of corporate counsel say they do not know how their outside firms approach AI.

Legal is not the only function moving either. The AI in Professional Services Report found that 47% of corporate legal departments have enterprise-wide GenAI access, against almost two-thirds (62%) of corporate risk and compliance teams. The longer these choices evolve independently, the harder it becomes to reconnect intelligence, automation, and accountability later.

AI adoption

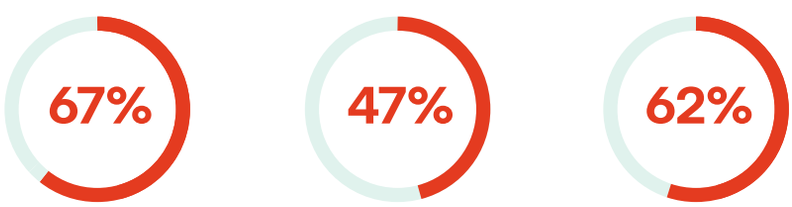

of **corporate counsel** say they do not know how their outside firms approach AI.

of **corporate legal departments** have enterprise-wide GenAI access.

of **corporate risk and compliance teams** have enterprise-wide GenAI access.

Source: Thomson Reuters Institute, 2026 AI in Professional Services Report

Orchestration asks whether the whole system of legal work has become better, rather than whether an isolated task has become faster.

It can be tested in five questions.

1. **Authority. Is the intelligence behind every action authoritative?** Is it clearly grounded in primary law, practice guidance, and the organisation's own standards, rather than a model's unverified web search? As AI gains the ability to act, the provenance of what it reasons on becomes more consequential, not less.
2. **Context. Does enterprise context shape the answer?** Playbooks, precedent, contract data, risk appetite, and business priorities should inform what the system does, not just what it says.
3. **Boundaries. Is delegated action bounded?** Routine activity should proceed safely, while exceptions halt processes until the right expert takes action. What the system may do, recommend, or initiate is defined in advance.
4. **Proof. Can every decision be traced and defended?** Sources, steps, approvals, and human interventions should be recorded in a form an auditor, regulator, or court will recognise.
5. **Value. Does the outcome appear in terms the business understands?** Think about reporting on business velocity, quality, risk, resilience, and strategic capacity, not just hours saved.

What does this look like in practice? The legal team loads a newly arrived agreement from a vendor. The system matches it to the approved template and checks against a playbook. With that authoritative guidance, it outlines the company's fallback positions and drafts a low-risk deviation within a pre-approved position. That reviewed agreement goes to the GC for approval, with notes to flag a data-transfer clause to Privacy. The playbook ensures compliance with the needs of business partners including Procurement, Finance, and Security. Each step of the agent's work is recorded. The GC only needs to spend time on the decisions that need it. Six months later, Legal has clear metrics to report to the business: cycle time down, exceptions down, external spend avoided, with an audit trail to prove it. None of this requires a single vendor or model. It requires that authority, context, boundaries, proof, and value have been designed in.

Taken together, these five questions provide a practical way to begin discussing readiness across the executive team without reducing maturity to the number of AI tools deployed.

Few companies will be equally ready across all five dimensions. A legal department may have authoritative AI-enabled intelligence but limited integration with enterprise workflows; another may have sophisticated automation but weak governance over delegated action; a third may have strong controls but struggle to demonstrate business value. This creates an emerging agentic readiness gap: the distance between what AI is technically capable of doing and what the legal team is operationally ready to let it do. The next stage of legal AI maturity will be defined less by how much AI a department uses than by how coherently authority, context, boundaries, proof, and value work together.

Each member of that group brings a different question, and a credible orchestration programme has to answer all of them.

How does Legal make the business case for investment with the executive buying committee?

Role	The question they will ask	What orchestration must provide
Lawyers and legal operations	Can we scale repeatable work without losing judgement, and prove it?	Legal quality and risk control; process architecture, data, metrics, and adoption
Finance	Does this investment produce measurable enterprise value?	A measurement logic linking AI activity to avoided spend, throughput, leakage, and redeployed capacity
Operational and business leaders	Is the end-to-end process faster and more reliable?	Fewer queues, clearer ownership, faster exception handling across hand-offs
IT and digital transformation	Will this create another AI silo?	Interoperability with enterprise AI and systems of record; agent lifecycle ownership
Information security	What can an agent access, and what can it do at machine speed?	Identity, permissions, monitoring, and lifecycle controls that define the boundaries of orchestration
Compliance, risk, and privacy	Under what authority did the agent act, and can we show it?	Provenance, traceability, risk classification, retention, and escalation built into workflows
Procurement	Does this vendor fit our enterprise requirements, not only our feature list?	Evaluation of security, interoperability, data use, model governance, resilience, exit terms, and measurable value

The cost of late alignment is predictable: Security discovers agents with permissions it never reviewed; Procurement finds duplicate purchases; Finance struggles to connect costs with benefits; Legal is presented with workflows it is expected to govern but did not design. The practical advantage of getting the group together early is not more governance for its own sake. It is the ability to move faster later because authority, architecture, controls, and value measures have been agreed before implementation.

These questions are intentionally cross-functional, and no single member of the group can answer them alone. That is the point. Agentic AI turns legal transformation from a departmental technology programme into an enterprise design problem.

Traditional technology governance asks whether a system is approved and who may access it. Agentic systems add a harder question: what may the system itself be authorised to do?

Governance has to get ahead of delegated action

Traditional technology governance asks whether a system is approved and who may access it. Agentic systems add a harder question: what may the system itself be authorised to do?

Agentic AI in legal settings can result in silent failures, diffusion of accountability, and confidentiality risk, all of which argue for testing, governance, and rigorous human oversight. Current Microsoft® security guidance stresses identity, policy, monitoring, and lifecycle governance for agents, and the World Economic Forum's work on trusted agent adoption emphasises authorisation, auditability, and enforceable delegation. For legal work involving privileged, confidential, or commercially sensitive information, these are not controls to add after deployment. They define the space in which orchestration operates. GCs are also increasingly framing AI around business implications, with a number pointing to AI regulatory preparedness, across global jurisdictions, as an area of focus.

A practical starting point is risk-tiered autonomy, built on the difference between reversible tasks and consequential actions.

Tier	Typical activity	Control
Automate	Summarising for review, classifying, routing, first-pass drafting against approved playbooks	Sampling, monitoring, and clear escalation triggers
Recommend	Proposing mark-ups, positions, or risk ratings	Human review and explicit approval before anything leaves the system
Reserve	Approving a contractual position, sending an external communication, changing a record of authority	Direct human decision, supported by traceable evidence

The aim is appropriate autonomy and managed risk. Nor should governance be positioned as a brake on innovation. Security, Compliance, Legal, and IT teams that agree common standards in advance do not have to renegotiate acceptable behaviour every time a new agent or workflow appears. That is what allows innovation to scale.

The case for change must answer a human question as well as a financial one: what does freed capacity let people do that is more valuable or more fulfilling?

The workforce model cannot be an afterthought

The hardest questions may concern people rather than technology. If routine legal activity is increasingly performed by AI, what becomes of the lawyer's role? What does a manager supervise when part of the team is digital? Which capabilities become more valuable?

The likely answer is not the disappearance of legal expertise but a change in where it is applied: less time retrieving, formatting, and transferring information; more time on exceptions, negotiation, strategic interpretation, and advice where context and judgement matter most. Legal operations moves further into process architecture, data, technology governance, and benefit realisation. New hybrid roles will emerge around legal engineering, AI governance, and agent management.

There is also an apprenticeship problem. If agents take on first-pass analysis and drafting, legal teams must decide how junior professionals learn and how experienced ones keep developing judgement. Short-term gains must not weaken the pipeline of trusted expertise.

The appetite is there. According to the AI in Professional Services Report, 72% of corporate legal professionals believe GenAI should be applied to legal work, and 60% feel excited or hopeful about its prospects. What they need is a credible path. The AI Success Pyramid (individual users, operations, leadership, strategy) describes how to build adoption inside the department. Orchestration extends the same logic past Legal's walls. Leaders must explain the destination, involve professionals in redesigning work, create incentives that reinforce new behaviours, provide credible training, and hire for adaptability and technical literacy. The case for change must answer a human question as well as a financial one: what does freed capacity let people do that is more valuable or more fulfilling?

Six questions to answer before you scale

A first test of readiness. Legal teams do not need a perfect future-state architecture before they begin, but they do need to make several decisions before agentic AI scales around them. A useful test is whether Legal, Finance, Operations, Technology, Security, Compliance, and Procurement can answer the following six questions together:

1. Which end-to-end workflows (contracting, regulatory compliance, investigations, disputes) matter most, rather than which isolated tasks?
2. What trusted intelligence and enterprise context does each workflow require?
3. What may AI do, recommend, or initiate at each level of risk?
4. How will specialist legal AI work with enterprise platforms, and who owns the lifecycle of each agent?
5. What changes for roles, skills, and incentives, and how will adoption be supported?
6. Which business outcomes will define success, and who will report them?

Teams that can answer these six questions credibly have the foundations for orchestration. If the answers change depending on whether the question is put to Legal, IT, Security, Finance, or Procurement, that is itself useful information: the organisation may have an orchestration gap before it has an AI capability gap. Understanding and benchmarking that readiness is the logical next step.

Legal can treat AI as a route to faster drafting and review — or it can help the organisation decide how trusted intelligence, automated action, and human judgement should work together.

The opportunity: Make Legal part of the enterprise AI operating model

There is a strategic choice here. Legal can treat AI as a route to faster drafting and review and remain principally a consumer of enterprise technology. Or it can help the organisation decide how trusted intelligence, automated action, and human judgement should work together — becoming an architect of the enterprise's AI-enabled operating model. The window to make that choice is now, while roles, controls, and workflows are still being formed.

That does not mean Legal should own enterprise AI. It means the function has something distinctive to contribute: a professional culture built around authority, evidence, risk, judgement, and accountability. Those characteristics become more valuable when machines can not only generate information but also act on it.

The prize is a governed operating environment in which legal intelligence flows into business decisions, routine work proceeds with less friction, exceptions surface earlier, AI acts within boundaries everyone understands, and Legal's contribution becomes visible to the C-suite. The 86% and the 17% can converge, but only if the work is designed so the business can see it.

Efficiency will arrive as a by-product. What each GC and their legal team must decide is what kind of legal function the enterprise will need in 2030 to start building towards it now. Merely having the most agents will not win the next phase of legal AI. Success will rely on legal teams that coordinate AI agents, people, data, and systems, around outcomes the enterprise values and can trust. Because agentic AI changes not just the tools people use but the organisation of work itself, orchestration and change management are two sides of the same transformation.

The question to take back to your executive team

If AI starts to act across Legal tomorrow, have we already agreed what it may do, what it must know, when a human must decide, and how we will prove the outcome?

How ready is your organisation to transform Legal through adopting an orchestrated approach to technology and AI?

As AI moves from assisting professionals to acting across enterprise workflows, the challenge is no longer simply AI technology adoption. It is ensuring that trusted intelligence, enterprise context, human judgement, and appropriate governance all work together to deliver measurable business outcomes.

Thomson Reuters is exploring these challenges with corporate legal leaders as organisations prepare for the next generation of AI-enabled corporate legal operations.

Contact us today to learn more.

Thomson Reuters™

© 2026 Thomson Reuters TR8991038/11-26