

The Legal Leader's Guide to Scaling AI

Governance, sequencing, and the business case

When managed well, governance is not a brake on legal AI. It is what makes scaling it possible, and it is the capability that the rest of the business will come to Legal to provide.

The caution is the advantage

The caution that senior lawyers bring to AI is usually treated as a problem to be managed. It is actually an advantage that is often overlooked. The same instinct that makes a good lawyer distrust an unproven tool is exactly what lets a legal team adopt new ones safely. The same instinct also equips a legal team to write the guardrails, policies, and governance standards that a safe rollout needs. Put to work, that caution is what lets Legal adopt AI at scale without giving up the obligations that justify the function.

The pressure is real and recent. Across corporate legal, AI has already moved past experimentation. According to FTI Consulting and Relativity, general counsels' use of generative AI rose from 44% to 87% of departments in a single year.¹ Yet only about half of these departments have a formal plan in place. Adoption of AI has surpassed their stated intentions. While most legal teams are using AI, significantly fewer have clearly defined what specific purposes they are using it for. This gap is the focus of this guide.

44% → 87%

General counsels' use of generative AI, in a single year.
FTI Consulting & Relativity.

Only about half have a formal plan for it



Adoption has outrun intention. That gap is the subject of this guide.

This guide is written for the chief legal officer or general counsel working to close that gap, and for the legal operations leaders who turn those decisions into practice.

Draws only on published research, publicly disclosed production deployments, and the observed practices of legal teams already operating at scale.

Most of the legal week is not judgment work

Some argue that Legal should be the last function delegated to AI, because it involves judgment, and judgment cannot be automated. The first part is correct. The second part assumes that a lawyer's week mainly consists of judgment, which is not true.

Thomson Reuters estimates that approximately 60% of a legal team's time is spent on routine and repeatable work.² The standard contracts, recurring research questions, compliance checks, and first drafts follow a pattern the whole department already knows. That is precisely the work today's AI does well. The time spent on complex judgment and decision-making is actually less than many outside the legal field might think.

THE LEGAL WEEK, SPLIT HONESTLY

60%

ROUTINE AND REPEATABLE WORK

Standard contracts · recurring research · compliance checks · patterned first drafts

40%

JUDGMENT

The work that defines the profession, and stays with the lawyer

Most of the legal week is the exact work AI is good at.

The prize, at scale, is substantial: the global legal services market runs to roughly USD 1 trillion.³

The shift is in who is now asking for AI: the demand has moved from vendors and individual associates to the CEO. It is less about the technology itself and more about who is now enquiring about it. According to the 2026 Chief Legal Officers Survey from the Association of Corporate Counsel (ACC), sponsored by FTI Consulting, 47% of chief legal officers said their CEO now expects them to develop technology and AI proficiency.⁴ That expectation lands on the legal leader personally, not on a vendor's pitch or an ambitious associate's side project. The demand is coming from the top, and it lacks patience, because boards can see peers already deploying AI and expect legal to keep pace.

So the question inside most departments has quietly flipped. The real question is no longer whether to use AI, since almost everyone already does. It is how to scale it without sacrificing the obligations that give Legal a reason to exist. That is the harder question, and the better one, because unlike the technology, the solution is entirely within a legal leader's control.

The complaints reduce to three causes

Every legal department describes its problems a little differently: one calls it a contract backlog, another slow deal cycles, a third “legal is the bottleneck”. But the underlying list barely changes from one company to the next. That consistency points to something structural rather than particular to any one team. The complaints collapse into three root causes, each with a different fix: work volume has outgrown manual methods, legal's own information is not in a state any AI can use, and board-level ambition has outrun the department's actual readiness.

The first is that the volume of legal work has grown, but the methods of doing it have remained unchanged. Companies review thousands of contracts manually, often resulting in contract backlogs and commercial teams waiting days for sign-off that should otherwise take hours. The less obvious cost is financial. According to the 2025 Corporate Legal Operations Consortium (CLOC) State of the Industry Report, outside counsel absorbs somewhere between 45% and 55% of the typical legal budget,⁵ and much of that is routine drafting and diligence billed at premium rates, with little visibility into where the money actually went.

The second is that Legal's own information is not in a state that any AI can use. Part of this challenge is fragmentation: contract and matter data are spread across email, shared drives, signature platforms, and legacy systems, with no single place that holds the truth. But the bigger issue is that even when the data is found, most of it is not captured in a machine-readable format. It lives in scanned PDFs, inconsistent templates, and free-text emails, without the structure, metadata, or clean version history needed for reliable machine reading. This is the problem that rarely reaches the board report and hinders more AI projects than model limitations do, because a tool can only work with data it can access and understand.

The third challenge is the gap between strategic ambition and operational maturity (readiness). While the strategy is ambitious and approved at board level, the governance framework and guardrails often lag by a decade. When the technology is adopted at a faster pace without the necessary governance and sequencing, this discrepancy becomes costly, especially in compliance. Regulatory risk then accumulates invisibly until audit or enforcement. The financial impact is significant, as DLA Piper reported EUR 1.78 billion in GDPR fines in just one year.⁶

THE FAMILIAR COMPLAINTS COLLAPSE INTO THREE CAUSES



Legal has a unique trait not shared by any other function, and it shapes everything that follows. It is not just about adopting AI; Legal is also the department that the rest of the company relies on for rules governing AI usage. Any policies or frameworks Legal develops must withstand scrutiny from the board, and not solely function at a desk. While this presents a short-term constraint, it also offers a long-term opportunity. A legal team that masters AI governance will possess the crucial capability the entire organisation will soon require.

Privilege and regulation, on conditions you can write down

Every serious conversation about legal AI revolves around a key question. Can we use this without giving up privilege, confidentiality, or our professional responsibilities? It is the right question to ask. And the answer is straightforward: yes, provided certain clear conditions are satisfied. The advantage is that these conditions are specific enough to be written down and handed to a team.

The first condition is timing. Run the confidentiality and privilege review with stakeholders before the tool is connected. Examine where the data flows, who can access it, and if the vendor trains on what you feed it. Then, limit access to only what the task requires. The tool should see what it needs and nothing more. Often, the mistake is not malicious. It occurs when someone gives a capable system complete access to the entire document management system (DMS) simply because it was the quickest way to get it operational.

The second condition is familiar to the profession because it aligns with existing rules on delegation, competence, confidentiality, and supervision. Delegating a task to an AI assistant is subject to the same expectations as delegating to a junior: the supervising lawyer stays responsible for competence in the work, for keeping client information confidential, and for reviewing the output. The standards apply to work done with an AI assistant exactly as they apply to work done with a junior lawyer. The key test is comparing the review process: if a lawyer reviews the AI output the way they would review an unsupervised first-year's draft, and discloses its use where expected, they meet those same professional standards of competence, confidentiality, and supervision. This standard is not new; only the AI assistant is.

Accuracy raises the same kind of concern as privilege and professional duty: something to be managed with a defined process, not a reason to avoid the tools. Here, the evidence asks for careful handling rather than avoidance. A 2025 Stanford study found that even the best legal research tools still return a wrong or unsupported answer between 17% and 33% of the time.⁷ These errors range from citing cases that do not exist to misdescribing what a real case actually held. That shapes how a department uses these tools, not whether it should. Successful teams typically adopt two practices: first, they base the system on verified source documents, ensuring responses come from the firm's own materials rather than a model's memory. Second, they assign a human to review every output, logging corrections. When these practices are followed, accuracy shifts from being a matter of trust to a matter of record.

Regulation is less restrictive than its reputation implies. Under the EU AI Act, AI used in the administration of justice can be classified as high risk, while most routine internal tasks, such as contract analysis, document review, and research assistance, are generally not classified as high risk; they are instead assessed on a case-by-case basis. The Monetary Authority of Singapore enforces similar standards for regulated entities, with various jurisdictions implementing their own.

Meeting the standard is an ongoing task rather than a one-time check, and it rests on two things: someone is responsible for keeping the policy current as rules change, and the regulatory watch covers every market the organisation operates in, not only the largest.

FOUR FEARS, FOUR ANSWERS

THE FEAR	THE CONDITION THAT ANSWERS IT
Privilege	Run the confidentiality and privilege review before the tool is connected, not after: where the data goes, who can reach it, whether the vendor trains on what you feed it. Then scope access to the task. The tool sees what it needs and nothing more.
Professional duty	Apply the rules that already cover delegation: competence, confidentiality, supervision. Review the output as you would an unsupervised first-year's draft, and disclose its use where a client or regulator expects it.
Accuracy	Ground the system in verified source documents, keep a human on every output, and log the corrections. Accuracy becomes a matter of record, not faith. 17–33% error rate in the best tools · Stanford, 2025⁷
Regulation	Most routine in-house work sits outside the EU AI Act's high-risk tier and is judged case by case; MAS sets comparable expectations. Someone owns keeping the policy current, and the watch covers every market you operate in.

Start where the privilege risk is lowest.
Automate what is well defined. Never
automate the judgment call.

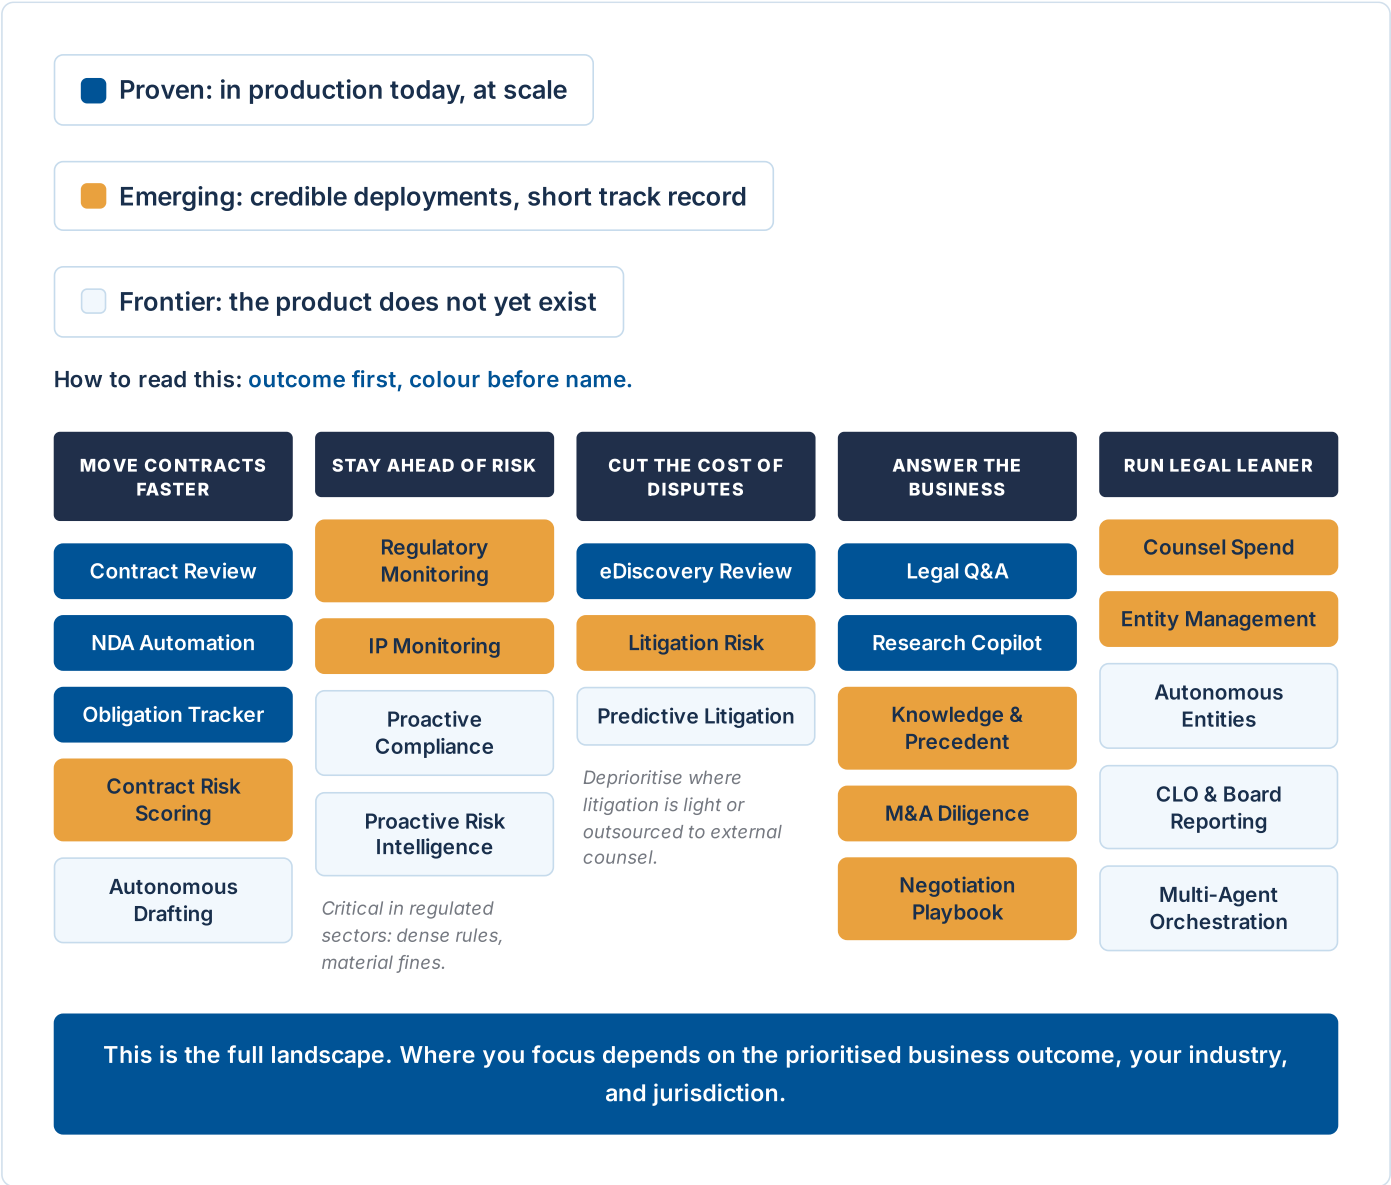
These conditions are not brakes. They are what turns a cautious yes into a defensible one, and they keep the legal team in charge of its AI rather than governed by it.

What is proven, and what is only promised

The landscape shown below represents what legal AI can do or claims to do. As a shopping list, it is not very useful. However, as a filter, it does real work, and it is primarily driven by two habits. The first is to read across the top rather than down the side, since the columns represent outcomes, not products. Market offerings typically present categories first and ask which you prefer, but it is more effective to focus on the outcome a department must achieve this year and then select the tools from that specific column. The other options can be deferred.

The second habit is to look at the colour code before reading the name of the use case. Every use case here is classified as proven, emerging, or frontier, with the colour mark telling you more than the name does. Proven means it runs in production today, at a real organisation, and at scale. Emerging suggests credible deployments but a short track record. Frontier means the idea is sound, and the product has not yet been developed. Technologies like predictive litigation, autonomous entity management, and multi-agent orchestration are all in the frontier stage because there is nothing to benchmark them against, and the landscape keeps them separate rather than folding them in with the proven work. When a vendor markets a frontier capability as ready to buy, they are making a claim ahead of the evidence, and you are within your remit to ask for the disclosed deployment that supports it.

THE USE-CASE LANDSCAPE · FIVE BUSINESS OUTCOMES AI IS ALREADY DRIVING IN LEGAL



Emerging and frontier use cases build on the proven layer they reuse.

The gap is particularly noticeable with agentic AI, a system that performs multi-step tasks independently rather than answering one question at a time. There is real interest: Deloitte's 2026 survey of senior legal leaders found roughly 61% of legal teams are experimenting or piloting these systems, while about 9% are running them in production.⁸ Thomson Reuters estimates real agentic use across professional services near 15%.⁹ Together, these figures show a capability that many teams are piloting, but few are running in production.

A good first use case has three properties: high volume, a repeatable structure, and clear decision criteria. Contract review, legal research, and compliance monitoring all exhibit these traits, which explains their recurring presence in the work that has actually reached production. A candidate lacking one of the three traits is not a poor choice. It just belongs later, after initial wins have generated data and confidence to support it.

The gap is the whole story

LEGALFLY's 2025 *AI Governance Gap* survey, covering 154 general counsels across the UK, France, and Germany, puts the problem in two figures: 77% of legal and compliance functions now use AI in some form, yet only 18% have a formal governance framework around it.¹⁰ This adoption-versus-governance gap highlights the core problem: most legal functions are already running capable AI systems on live client matters, with little documentation on processes. This creates significant exposure, and the professional risk of that sits with the lawyer, not the tool. Conversely, it presents an opportunity that governance remains the least developed area. Teams that tackle the gap early are not only more secure but also lead the industry, upholding a discipline that many others neglect.

ADOPTION WITHOUT A FRAMEWORK


77%

of legal and compliance functions now use AI in some form



18%

have a real governance framework around it¹⁰



When legal technology fails to deliver what was expected, the cause is rarely the technology itself. Gartner predicts that nearly half of first-time Contract Lifecycle Management (CLM) rollouts will fall short of expectations,¹¹ and the recurring reasons are not technical. They are unclear ownership, no agreed-upon rules, and no plan to encourage actual usage. AI is no exception, and the stakes are higher because an AI tool is live on client matters from its first week.

This is where the effort belongs. Model quality is not something a legal team manages. It sits with the vendor and improves with each release, regardless of what the team does. Governance and change management are different. Who owns the tool, what rules it runs under, and whether people actually use it are all decisions the department makes, and these are the factors that determine whether a rollout succeeds.

So what does a framework actually look like in the teams that have one? It consists of three simple layers.

Ownership: set up a designated group, call it a Legal AI steering group, that sets priorities, approves use cases, and manages ethical and compliance risks. Having a single accountable body keeps decisions from being made impulsively by whoever happens to be in the room.

Operating model: this addresses who builds and maintains workflows. There are a few established structures and no one-size-fits-all answer. The goal is to intentionally choose a model suited to your team's specific needs, rather than inadvertently adopting a mixed approach.

THREE OPERATING MODELS · WHO BUILDS AND OWNS THE WORKFLOWS

MODEL	WHO BUILDS AND OWNS	TRADE-OFF	BEST FIT
IT-led / centralised	A central IT team builds and runs the tools; Legal specifies needs and signs off on output	Most control, at the cost of throughput	Functions that prioritise tight control and have IT capacity
Centre of Excellence	A small group from legal ops and IT owns the use cases, sets standards, maintains a governed workflow library	Balances agility with governance	The common default for large legal functions
Citizen development	Lawyers build their own workflows on governed low-code platforms; IT provides the guardrails	Fastest path from idea to value, needs discipline	Teams with high data maturity

Data rules: these are established before the rollout, not after it. In practice, that means a plain map of which documents go to which tool, whether any tool trains on the material it is given, and where outputs live, all condensed to a single page that the team can actually refer to. The rules can be straightforward; unredacted M&A financials never touch a public chatbot, while a private CLM instance is cleared for all contract text. Human-oversight thresholds should be incorporated into the training process rather than being part of a static policy that no one revisits. These thresholds specify which outputs need senior review, the actions to take if issues arise, and who is responsible for quality control.

None of this is a process for its own sake. It is the whole of what separates the 18%¹⁰ from the rest, and on current numbers, the rest is most of the market. The teams that build it early are not the ones slowing down. They are the ones equipped to scale.

A pilot shows what is possible. Production shows what is true.

Legal AI solutions get talked about far more than they get deployed, and the distinction matters. A pilot demonstrates potential but a production deployment shows what happens when people actually use it at scale. Three publicly reported deployments give us a useful cross-section: one shows the scale of adoption, one shows a measurable improvement in accuracy, and one shows a significant reduction in time spent on routine work. Together, they offer evidence of what legal AI looks like in practice and not just what vendors say it could do.

THREE PUBLICLY REPORTED DEPLOYMENTS, RUNNING AT SCALE

ALLEN & OVERY · 2023

3,500+ lawyers, 43 offices

In 2023, Allen & Overy put Harvey in front of more than 3,500 of its lawyers across 43 offices.¹² That scale matters because the firm moved beyond testing AI with a small group and made it available across its global practice.

HSBC · 2023

>60% fewer alerts

HSBC pointed AI at a narrower, harder target: screening transactions for financial crime. It reported cutting alert volumes in its anti-money-laundering checks by more than 60% in 2023, while surfacing two to four times more true positives.¹³ Compliance staff stop chasing alerts that lead nowhere.

KLARNA · 2024

~1 hour →

~10 minutes

Klarna's in-house lawyers use an enterprise AI assistant to draft their routine contracts. One of them told City A.M. in 2024 that a task which used to take about an hour now takes about 10 minutes.¹⁴ The lawyer still reviews and adapts the draft, but AI cuts the time spent getting to a usable first version.

Harvey is a legal AI assistant built on OpenAI models for drafting, research, and document review. Allen & Overy's rollout was one of the first firm-wide legal AI deployments anywhere.

The visible results are only the endpoint. Behind them are quieter operating habits that determine whether an AI programme can actually take hold and scale.

When those disciplines are neglected, they produce predictable failure modes: essentially the same habits in reverse.

FIVE OPERATING HABITS BEHIND SUCCESSFUL LEGAL AI PROGRAMMES

The same habits that help an AI programme scale can also become predictable failure points when they are skipped.

THE HABIT THAT HELPS A PROGRAMME SCALE	WHEN THE HABIT IS SKIPPED
✓ Find out how lawyers are already using AI. Survey who already uses which tools, including personal ChatGPT and Claude accounts.	✗ Assume lawyers already know what AI can do. A plan built on a false baseline is likely to miss what people actually need.
✓ Pilot with sceptics, not just enthusiasts. If a sceptical lawyer finds the tool useful, that credibility can help bring the wider department along.	✗ Pilot with enthusiasts only. Results can collapse when the tool is deployed to the wider department.
✓ Track AI outputs and whether they were correct. Measure accuracy over time and identify recurring errors while they are still easy to fix.	✗ Automate before review standards are clear. No one knows which outputs need a second look or how to judge them.
✓ Add AI in where lawyers already work. Embed these tools inside Microsoft Word, emails, and contract management systems, not as a separate tool they have to remember to open.	✗ Make lawyers leave their usual tools. A tool that pulls them out of Microsoft Word or their contract management system will get used less, regardless of how good it is.
✓ Pilot with clean data, fix the rest in parallel. Use a controlled, high-quality dataset for the first use case while fixing broader data issues.	✗ Wait for perfect data or ignore data quality entirely. Both extremes can stall a programme or create avoidable risk.

As the legal AI market becomes more competitive, buyers have more options than ever. The gap between general-purpose AI and specialist legal AI tools is also narrowing: specialist tools increasingly differentiate themselves through workflow integration and legal-specific features, rather than fundamentally better AI. That makes flexibility more valuable. Before committing to a platform, buyers should ask four critical questions.

FOUR QUESTIONS FOR ANY VENDOR

1. Do our inputs train your model?
2. Is our data held in a private instance, a dedicated environment not shared with other customers?
3. How is legally privileged data protected, both contractually and technically?
4. How secure is your access to the underlying AI models, and what happens if that access changes?

TWO HABITS TO GO WITH THEM

- Test the integrations, not just the demo.
- Schedule an external review six months after launch.

Sequence by readiness, not ambition

A roadmap like this is most effective when it is sequenced by readiness rather than ambition. Boards are often impressed by use cases that are not yet prepared to succeed. Teams that do well begin with use cases they can realistically deploy, let the first win build confidence and capability, and then move to harder problems.

The roadmap below assesses every use case on two dimensions: business value and ease of deployment. Value runs vertically, deployment difficulty runs horizontally. Colour shows the maturity of the underlying technology: Proven, Emerging, or Frontier.


Readiness is about how prepared you are to deploy a use case successfully. That depends on factors such as data quality and access, required integration, and the permissions or confidentiality involved in the workflow.

The result is four practical categories:

- **Quick Wins:** high value, relatively easy to deploy. **Do now.**
- **Big Bets:** high value, but harder to deploy. **Plan and fund.**
- **Fill-ins:** lower value and relatively easy to deploy. **Use when capacity is available.**
- **Long Games:** lower immediate value and harder to deploy. **Watch rather than buy.**

The point is not to avoid ambitious use cases. It is to earn the right to tackle them by starting where readiness is highest, learning from deployment, and using those gains to move towards the harder opportunities.

EVERY USE CASE, PLACED BY VALUE AND EASE OF DEPLOYMENT



These four quadrants represent an evaluation of each use case. Sequenced by phase rather than plotted by value, the same elements serve as a roadmap.

The same use cases, sequenced by phase

Turn readiness into timing and the quadrant becomes a roadmap. The columns represent the outcome a department must achieve under pressure, while the rows indicate when each use case is prepared for deployment. The colour continues to reflect maturity.

ILLUSTRATIVE ROADMAP · PRIORITISE BY OUTCOME, SEQUENCED BY PHASE



Read that way, the sequence mostly writes itself. The early wins are where proven capability meets contained privilege risk, with contract review as the clear example, and these can be implemented within 6 months. The middle tier seeks more integrated solutions and organisational confidence, so it lands between 6 and 18 months. The farthest tier is frontier territory, 18 months and out, with some items that still have nothing to benchmark against, which the chart honestly indicates. All other tasks are supplementary, requiring minimal effort at any stage and can be added as capacity allows.

The most important part of the chart is the section easiest to overlook because it sits beneath everything else: the foundation row. This includes data infrastructure, AI capability and literacy, organisational confidence, governance and bar compliance, and clear vendor and build-versus-buy rules. These five are not just another phase of work; they are the support for every other column above. Ignoring the foundation might still allow quick wins, but it can be costly. Early successes can mask the missing foundational elements. The major initiatives meant to justify the programme often fail because they lack this solid base. This is again the governance chapter's main idea, now presented as architecture. The foundation is the advantage, while everything visible depends on the part that remains unseen.

Two key figures keep expectations realistic, based on experience from teams already running programmes like this. Typically, these initiatives take 18 to 36 months from start to finish, with substantial scaling usually happening 6 to 12 months after a successful pilot and foundational governance are established. And the most common way these programmes fail is overcommitting early, expanding department-wide before the support and governance can hold it. The chart represents a continuous loop rather than a final destination. It is meant to track one outcome from the short term to the long term, then begin anew, since the resources gained from the first result can alter what qualifies as a quick win the next time.

A model the CFO can interrogate

To illustrate the method concretely, meet Serene Tan, General Counsel at a mid-sized company, and Arjun Pillai, her legal operations manager. They serve as an example based on real situations faced by legal leaders, rather than represent a single company's experience.

Serene has been managing the legal department for 6 years. She has the AI programme mapped, governance established, and a successful pilot. What is missing is an essential meeting with her CFO to determine if it can scale. She anticipates how the meeting might go: a slide of vendor statistics, a confident headline number with no context, and a seasoned finance chief who, after seeing countless tech business cases, gains little insight from this one.

Instead of creating a pitch, Serene and Arjun decided to build a model from the CFO's perspective. They start by measuring a basic departmental workflow, the standard NDA, which accounts for 2,000 documents annually and roughly an hour of manual work per document. This modest figure is crucial because savings that are not measured against anything are merely hopes, not actual savings.

The calculations then become transparent, focusing on three key areas, each one represented as a row in the model that the CFO can easily review.

THE VALUE MODEL, AS THE CFO SEES IT

VECTOR 1 · ROWS 15–16

Capacity released

Once a draft lands in under 2 minutes and review takes 15, roughly 45 minutes come back on every NDA the system handles. Applied to half the volume in year one and rising from there, that is around 750 hours returned to the team, valued at the department's own loaded rate.

$\text{volume} \times \text{time saved per matter} \times \text{loaded rate} \times \text{touchless rate (50\% / 65\% / 75\% over Y1 / Y3 / Y5)}$

VECTOR 2 · ROW 17

External counsel cost reduced

A conservative share of the NDAs that used to go to outside counsel now stays in-house.

$\text{routine matter volume} \times \text{external rate} \times \text{share insourced (benchmark 40–60\%)}$

VECTOR 3 · ROW 18

Risk & compliance protected

Renewal alerts and obligation monitoring switch on at signature, so the missed deadline that used to surface as a bad surprise is tracked instead.

$\text{breaches avoided} + \text{missed-obligation cost} + \text{fine-risk reduction}$

Net of platform and change effort, the first year clears its own cost with room to spare, and the model shows exactly how.

Every assumption open to challenge · every input swappable for a real one

Use Case ROI Model (illustrative).xlsx			
Home Insert Formulas Data Review View			
B19 fx = B16 + B17 + B18 → capacity released + counsel cost ...			
	A	B	C
1	Use case ROI model · NDA workflow (illustrative)		
2	BASELINE & ASSUMPTIONS		
3	Measured before any tool...	the baseline	not a hope
4	Annual NDA volume	2,000	from the matter log
5	Hands-on time per NDA · ...	60 → 17 min	draft <2 + review 15
6	Time returned per NDA	~45 min	= before – after
7	Touchless rate · Y1 / Y3 / Y5	50% / 65% / 75%	
8	Loaded hourly rate	your actual	from finance
9	NDAs sent to outside cou...	your actual	from e-billing
10	External cost per NDA	your actual	from e-billing
11	Share insourced · bench...	40–60%	
12	Missed obligations per ye...	your actual	from the matter log
13	Cost per missed obligation	your actual	finance + risk
14	THREE VECTORS · YEAR ONE		
15	Capacity released · hours	~750 hrs	= B4 × B6 × B7 ÷ ...
16	Capacity released · value	= B15 × B8	hours at your own...
17	External counsel cost red...	= B9 × B10 × B11	share that stays in...
18	Risk & compliance protec...	= B12 × B13 × cover	tracked, not surpr...
19	Gross benefit · year one	= B16 + B17 + B18	the three vectors, ...
20	COST TO RUN		
21	Platform + change effort	your quote	licence, integratio...
22	Net · B19 – B21	Year one clears its own cost, with room to sp...	
	ROI Model	Assumptions	Actuals Ready

The number is not the point Serene is making, though. The model is. She hands the CFO a sheet where every assumption can be challenged, and every input replaced with actual data. Nothing is accepted without scrutiny, which is the only way a business case withstands two rounds of finance review.

The important part comes next: the CFO does not just approve the NDA case; they request that every future use case follow the same format. From then on, the model is not a one-time pitch but becomes the department's ROI standard, owned by the steering group. No workflow gets funded without passing through the same evaluation. Over time, the forecasts are replaced by actual results directly in the model, so the case never has to be rebuilt from scratch. This exemplifies how governance can shift from being a constraint to a driver of growth. A single disciplined model, applied once, then applied to everything.

Under it all runs one rule about what to measure. Track cycle time, capacity, and counsel spend, not just tool usage. Usage indicates that the tool is opened, but these numbers show it is effective. When budgets tighten, it is the second kind of evidence a programme is asked to produce.

Not a launch. A loop.

Everything in this guide comes down to a sequence that a legal leader can start this quarter. And it does not begin with a tool.


Instead, it starts with the department examining itself, including its operating and governance model, team structure, decision-making processes, and existing systems such as the CLM, the document store, and the matter database. Integration influences adoption more reliably than model quality. The same first step establishes where the data sits, its quality and sensitivity, and sets the baseline, with a commitment to measurement from day one, avoiding late-stage reconstruction.

Next, review the evidence. Before mandating anything, effective teams examine what similar functions in their sector and jurisdiction have actually implemented and disclosed, and compare it against their own regulatory and legal requirements. There is a subtle reason to do this early. A leadership team that has seen the tools in action can focus on the order of implementation rather than if it should happen. A shared understanding among leadership clarifies the decision-making process beyond what any written memo can achieve.

The third step is to let the lawyers articulate their pain points through structured sessions. These discussions surface genuine use cases, which are evaluated for effort, privilege risk, and value. A practical use case from actual experience tends to have a more lasting impact than one that came from a slide, with the difference usually becoming apparent within 6 months. This step is also the least likely to be skipped, as it is where resistance is either brought in or created.

Sequencing typically begins only after initial steps, and at this stage, it often becomes straightforward. Focus on quick wins initially, then scaling, and finally frontier work, all within a realistic 18- to 36-month timeline. Each use case should have its own business case based on three key factors, with governance and build-versus-buy decisions made before the first deployment rather than during it.

FOUR STEPS THAT CURVE BACK INTO A LOOP



STEP 1

The department looks at itself

Operating and governance model, matter flows, decision rights, the systems anything new must live inside. Where the data sits, how clean, how privilege-sensitive. Set the baseline from day one.

STEP 2

Study the evidence

What comparable functions have put into production and disclosed, mapped against your regulatory and bar obligations. Shared AI literacy at the top: a leadership team that has seen the tools work argues about sequence, not whether.

STEP 3

Let the lawyers name the pain

Structured sessions with the people doing the work surface the real use cases, scored on effort, privilege risk, and value. A use case from the desk outlasts one from a slide.

STEP 4

Sequence, then go again

Quick wins first, scaling next, frontier later, inside an honest 18-to-36-month arc. Each use case carries its own business case on the three vectors. The measured result from the first cycle makes the case for the second.

Change management, not technology, is what decides these programmes.

It is a continuous cycle, not a single launch. Each workflow undergoes assessment, a pilot involving sceptics, and a contained rollout, always keeping a human involved. The result from that first cycle then justifies the next step. The approach eliminates the need to repeatedly argue for AI, because each successful workflow provides evidence for the next. Support built on tangible, proven results is the most sustainable foundation for a programme.

Every thorough examination of legal technology reaches the same conclusion, and it is the appropriate point to conclude. Success depends on change management, not technology itself. The technical models are already sufficiently advanced; the human and organisational aspects, whether people, governance, and sequence of implementation, remain uncertain. This is the more challenging task, but it justifies the effort. A legal team that masters AI governance not only implements it more quickly than other departments but also becomes the go-to resource for others to learn how to do it.

References

Sources for the figures and production deployments cited above, in MLA style. Each entry is numbered to match its superscript marker in the text. Vendor-reported and illustrative figures are labelled where they appear.

1. FTI Consulting and Relativity. "AI Adoption in Corporate Legal Departments Doubles According to The General Counsel Report." *FTI Consulting*, 11 Mar. 2026, www.fticonsulting.com/about/newsroom/press-releases/ai-adoption-in-corporate-legal-departments-doubles-according-to-the-general-counsel-report.
2. "From Strategy to Execution in Legal Operations with AI." *Thomson Reuters*, 14 Apr. 2026, legal.thomsonreuters.com/blog/from-strategy-to-execution-how-modern-legal-teams-bridge-the-knowing-doing-gap/.
3. Mordor Intelligence. "Legal Services Market Size, Share and Trends Analysis." *Mordor Intelligence*, 2025, www.mordorintelligence.com/industry-reports/legal-services-market.
4. Association of Corporate Counsel. "The Changing C-Suite: CLOs Grow in Strategic Leadership as AI and Global Uncertainty Influence Corporate Decision-Making." *Association of Corporate Counsel*, 29 Jan. 2026, www.acc.com/about/newsroom/news/changing-c-suite-clos-grow-strategic-leadership-ai-and-global-uncertainty. Reporting findings of the 2026 ACC Chief Legal Officers Survey, sponsored by FTI Consulting.
5. CLOC. *2025 State of the Industry Report*. Corporate Legal Operations Consortium, 2025, cloc.org/wp-content/uploads/2025/02/2025-CLOC-2025-SOTI-Report.pdf.
6. DLA Piper. "GDPR Fines and Data Breach Survey: January 2024." *DLA Piper*, Jan. 2024, www.dlapiper.com/en/insights/publications/2024/01/dla-piper-gdpr-fines-and-data-breach-survey-january-2024.
7. Magesh, Varun, et al. "Hallucination-Free? Assessing the Reliability of Leading AI Legal Research Tools." *Journal of Empirical Legal Studies*, vol. 22, no. 2, 2025, pp. 216–42, <https://doi.org/10.1111/jels.12413>.
8. Deloitte. *The AI Imperative: Reshaping of the Legal Industry*. Deloitte, 2026, www.deloitte.com/nl/en/services/legal/research/ai-imperative-reshaping-of-the-legal-industry.html.
9. Thomson Reuters Institute. *2026 AI in Professional Services Report*. Thomson Reuters, 2026, www.thomsonreuters.com/en/reports/2026-ai-in-professional-services-report.
10. LEGALFLY. *The AI Governance Gap*. LEGALFLY, 2025, www.legalfly.com/report-overviews/ai-governance-gap-key-findings.
11. Cimplifi. "Why Many First-Time CLM Implementations Fail." *Cimplifi*, 2024, www.cimplifi.com/resources/why-many-first-time-clm-implementations-fail/. Citing Gartner's prediction that nearly 50% of initial CLM implementations will fall short of expectations.
12. Allen & Overy. "A&O Announces Exclusive Launch Partnership with Harvey." *Allen & Overy*, 15 Feb. 2023, www.aoshearman.com/en/news/ao-announces-exclusive-launch-partnership-with-harvey.
13. "Google Cloud Launches AI-Powered Anti Money Laundering Product for Financial Institutions." *PR Newswire*, 21 June 2023, www.prnewswire.com/news-releases/google-cloud-launches-ai-powered-anti-money-laundering-product-for-financial-institutions-301856403.html.
14. Jones, Jess. "Klarna Lawyers Drafting Contracts with ChatGPT and 9 in 10 Staff Using AI Everyday." *City A.M.*, 14 May 2024, www.cityam.com/klarna-lawyers-drafting-contracts-with-chatgpt-and-9-in-10-staff-using-ai-everyday/.

Disclaimer

This guide is general information, not legal advice. It reflects the views of the Aicadium team and does not represent the position of any client, portfolio company, or the individual practitioners consulted in its preparation. Aicadium is not a law firm and does not practise law.

The guide takes an outside-in view, drawing only on published research, publicly disclosed production deployments, and the observed practices of legal teams operating at scale; illustrative figures are labelled as such.

Nothing here creates a lawyer-client relationship or should be relied on for any specific matter. Regulatory positions, bar rules, and tool capabilities vary by jurisdiction and change quickly, so read anything time-sensitive against current sources and your own qualified counsel before acting on it.

Governance first was never the cautious path. It was the one that ends up in front.