

The Legal AI Readiness Checklist

Seven Questions to Answer Before Your Next Deployment

For general counsel and legal operations leaders. Complete this before the tool evaluation, not after it.

How to use this. Readiness has very little to do with the technology. AI fails in legal environments because the data, workflows, and governance underneath it are not prepared, and it amplifies whatever is already there. Answer all seven honestly. A **no** is not automatically a blocker, but it is a line item that needs a named owner and a date before you scale. Score yourself at the foot of the page.

1	Do we have approved use cases, or are people experimenting randomly? GOOD LOOKS LIKE A written tier list of prohibited, pre-approved, and tracked uses, covering employees, contractors, and third party providers.	☐ YES ☐ PARTIAL ☐ NO OWNER AND DATE
2	Do we have a data and confidentiality policy covering prompts, outputs, and connected systems? GOOD LOOKS LIKE One policy that reaches every connected system, not only the chat window where the prompt is typed.	☐ YES ☐ PARTIAL ☐ NO OWNER AND DATE
3	Do we know when human review is mandatory, and who owns that review? GOOD LOOKS LIKE A written trigger for mandatory review and a named reviewer per workflow. Accountability for outputs does not transfer to the vendor.	☐ YES ☐ PARTIAL ☐ NO OWNER AND DATE
4	Do we have a vendor diligence and contracting standard for AI tools? GOOD LOOKS LIKE A diligence standard applied before signature, covering data handling, model change, and what happens on termination.	☐ YES ☐ PARTIAL ☐ NO OWNER AND DATE
5	Do we have training by role, not just a generic all hands policy? GOOD LOOKS LIKE Training differentiated by role, so lawyers, paralegals, and legal operations each get what they actually use.	☐ YES ☐ PARTIAL ☐ NO OWNER AND DATE
6	Do we know how to measure value beyond anecdotes? GOOD LOOKS LIKE A baseline in hours, spend, or risk outcomes, captured before deployment rather than reconstructed after it.	☐ YES ☐ PARTIAL ☐ NO OWNER AND DATE
7	Do we have a named owner for governance after the launch, not only during it? GOOD LOOKS LIKE One named owner with time allocated past go live, usually the general counsel or a designated deputy.	☐ YES ☐ PARTIAL ☐ NO OWNER AND DATE

Seven yes

Proceed to tool evaluation. Scope the first pilot to one workflow with a measured baseline.

Five or six

Close the named gaps first. Give every no an owner and a date, then start narrow.

Four or fewer

Run a readiness program before the spend. The work is usually data and process, not tool selection.

The readiness gap

The distance between having access to AI tools and having the operational foundation to use them well. It shows up as fragmented data, undefined workflows, limited governance, and low adoption capacity. **AI amplifies what already exists,** so an inconsistent process becomes inconsistent output at scale.

Why the checklist decides the outcome

Most pilots do not fail on model quality. They fail on integration, data, and workflow, which is exactly what those seven questions surface before the spend happens.

~95% of enterprise GenAI pilots show no measurable P&L impact ~5% reach real operational return MIT NANDA, GenAI Divide: State of AI in Business, 2025

The five dimensions behind the questions

Experienced general counsel assess readiness across five dimensions. Every question on page one maps to at least one of them.

ONE Data Structured, standardized, accessible legal data with minimal gaps in the record.	TWO Process Workflows standardized, documented, repeatable, and measured before automation.	THREE Governance Usage policy, data controls, approval framework, audit, and named accountability.	FOUR Technology Real integration with the ELM, CLM, eBilling, and matter systems already in place.	FIVE Change Leadership alignment, role based training, user trust, and workflow fit.
---	---	--	--	--

What to do with your score

YOUR SCORE	WHAT TO RUN NEXT	WHAT IT STILL LEAVES OPEN
Seven yes	A disciplined tool evaluation. Pilot one workflow against a baseline you captured first.	Readiness is not permanent. Re-run these seven at each new use case, model change, or vendor switch.
Five or six yes	Gap closure, then a narrow pilot. Every no gets an owner and a date before scope grows.	Partial governance still leaves the general counsel accountable for the outputs the policy does not cover.
Four or fewer	A readiness program. Usually data remediation and process mapping ahead of any tool decision.	This is a budgeted program, not a policy memo. Aging systems may need remediation or replacement first.

Readiness is the work. The tool is the easy part.

Swiftwater's AI Lab helps legal departments find where AI creates real leverage and builds the governance to use it responsibly. When the data underneath is the blocker, we migrate it off the legacy stack, clean it, and hand back a platform that is actually ready.

- Legal AI readiness assessment
- Legal AI governance frameworks
- AI use policy and tiering
- Vendor diligence and contracting
- Data remediation and migration
- Process mapping and standards
- Tool evaluation and selection
- Role based training and adoption

400+

years of combined practitioner experience across Americas, EMEA, and APAC

Legal AI Solutions

Legal AI Hub

Book a discovery call