

AI excellence: Choosing strategic, data-driven investments

Contents

Introduction	3	Seven principles	10
• What will define legal services for the foreseeable future?	3	• What separates organizations winning with AI from those in a piloting loop.	11
• State of the market	4		
• What separates high-maturity organizations	5	The path to AI excellence	13
		• This is an AI architecture puzzle – to which the research holds the key	14
The AI excellence playbook	6		
• AI excellence: An architecture problem with a layered solution	7		

What will define legal services for the foreseeable future?


Spoiler: The bottleneck isn't the model or the tool; it's the foundation beneath it.

In every major technology cycle, we see a point where experiments begin to coalesce into tangible results. This is when the early movers start to dominate the narrative, leaving the less adventurous scrambling to close an ever-widening gap.

AI has already crossed that threshold. And the organizations that will define legal services for the foreseeable future aren't the ones with the flashiest tools or the biggest marketing budget. They're the ones that put infrastructure before ambition — quietly, methodically building the architecture that makes AI work.


That's what this playbook is about.

Grounded in data gathered from **3,000+ decision-makers from 26 countries**, it's a strategic guide for technical leaders whose organizations are striving toward AI excellence.




The state of the market


Everyone is doing something with AI. Very few have operationalized. Governance is lagging across the industry, while clients are raising the bar.




piloting, implementing,
or using AI



have AI fully integrated
into workflows




have experienced a policy
violation from AI



allow employees to use public
AI with little oversight



per day lost to
information search



say client needs directly
shape AI usage

What separates high-maturity organizations

The research findings around AI adoption rates are striking — but even more remarkable is the influence that knowledge work maturity has on every other business metric.

These outcomes aren't specific to AI. They are the exact kind of business outcomes that show up in board presentations, investor conversations, and competitive positioning statements. The data is clear: Mature knowledge work infrastructure doesn't just make AI perform better; it makes the whole organization perform better.

**4x**

more likely to be in the top financial performance quartile

28% of most mature vs 7% of least mature

**80%**

of the most mature firms are operate at a profit

vs 54% of the least mature

**77%**

of the most mature firms grew revenue last year

vs 39% of least mature organizations

The AI excellence playbook

AI excellence: An architecture problem with a layered solution

“It is all about IA, before AI.
If AI is the train, information
architecture is the tracks.”

Source: Reena SenGupta, Executive Director, RSGI Limited,
iManage Knowledge Work Benchmark Report 2026

The most common AI investment mistake is racing to adopt AI before getting the foundation in place. The research shows that when organizations attempt to deploy advanced AI capabilities on top of fragile, fragmented, or ungoverned knowledge infrastructure, they often end up stalled — or abandon the projects altogether. Unreliable outputs, governance breaking down, and disconnected workflows result in an unrealized ROI.

AI excellence requires a layered architecture. But unlike a rigid stack where every layer must be fully built before advancing, this framework is better understood as a progression of increasing capabilities — with one hard rule at the base,

and meaningful flexibility above it. Not every organization needs every capability within every layer. What matters is understanding which layers are load-bearing for your specific goals and which advanced capabilities will amplify performance rather than add complexity.

The following framework maps the **six interdependent layers of AI-ready infrastructure** — from the non-negotiable foundation to the human and organizational capabilities that allow the benefits of using AI to compound over time.

Data foundation layer

This essential layer is non-negotiable. Its centralized, governed storage, consistent metadata, version control, and lifecycle management yield clean, accessible data. It is the prerequisite to all other layers. AI working with fragmented, outdated, or siloed content produces fragmented, unreliable outputs regardless of the chosen model's quality. Every other layer in the framework is contingent on this critical layer being in place.

If it isn't, stop here and fix that first. No need to boil the ocean — just pick one workflow in a single practice area. Build success with a single workflow that is predictable, well understood, and low risk, and later apply that success to other areas.

Common pitfall

Investing in AI tools while content remains fragmented across SharePoint, email, and local drives. No AI model can compensate for a broken foundation.

Knowledge layer

The intelligence layer on top of your data foundation. Semantic search, automated tagging, template and clause recommendation, and expert identification. Natural language query — the single most adopted AI use case globally, at 55 percent — is unlocked here. Not every organization needs the most advanced capabilities in this layer, but organizations that invest here consistently outperform those that don't. The gap in natural language query adoption between the most mature organizations (58%) and the least mature (42%) is a signal that this layer compounds results. The more intelligently your content is organized and enriched, the more value any AI model can extract from it.

Common pitfall

Without a knowledge layer, AI tools must navigate through noise from redundant, outdated, or trivial content (ROT).

Governance layer

This is the permissions layer that determines what AI can touch, when, and with what accountability. Ethical walls, matter-level confidentiality, audit trails, policy enforcement, and “shadow AI” controls. This is not a future concern. Thirty-six percent of organizations have already experienced a measurable policy violation. Governance is a present-tense operational requirement, and this is the layer that earns client trust, regulatory resilience, and the organizational latitude to deploy AI more broadly.

Common pitfall

Not identifying a leadership role to steward governance. The most effective approach is to have a leader AND bake AI guardrails into workflows.

AI infrastructure layer

This integration backbone connects AI tools and agents to your knowledge and workflows without duplicating data or creating silos. APIs, model hosting flexibility, orchestration for multi-agent workflows, and monitoring. Critically, not every workflow capability requires this layer. Collaboration tools, external sharing, and many workflow automations operate independently of AI infrastructure. This layer becomes load-bearing specifically when organizations want to deploy AI at scale across systems: when the goal shifts from one AI tool to a coordinated AI ecosystem. Thirty-two percent of organizations cite integration with existing systems as their top barrier to AI adoption. This is the layer where you solve for that issue.

Common pitfall

Locking into a single AI vendor's closed ecosystem before interoperability standards are established.

Workflow layer

This is where AI stops being a separate tool and starts being a native capability embedded in how work happens. AI-assisted drafting, summarization, and risk flagging inside and alongside the document management platform. Automated routing and classification. Client-facing AI tools. But this layer also includes foundational workflow capabilities such as external collaboration, document sharing, and process automations that deliver value. Mature knowledge work organizations are nearly twice as likely to integrate AI into client-facing tools as less mature ones (46% vs. 24%). Integration is the multiplier. But the baseline workflow investments in this layer pay off regardless of AI maturity.

Common pitfall

AI in innovation pilots that never reach production due to misaligned workflows.

Human experience layer

Human infrastructure determines whether AI scales or stalls. AI literacy that extends across the organization, not just IT, includes change management and adoption programs, and new role creation. The most mature organizations are already making space for trust architects, digital ethics officers, and knowledge strategists. Thirty percent of organizations cite inadequate training as the number one barrier to effective collaboration, and AI amplifies that gap rather than filling it, with most AI failures linked to low adoption. This layer, and an intentional strategy to focus on human adoption, helps AI become a cultural capability. and where the ROI gets stalled as it's not just a technical challenge.

Common pitfall

Fifty-seven percent of the most mature say AI is creating entirely new roles; only 29 percent of the least mature say the same.

Seven principles

What separates organizations winning with AI from those in a piloting loop

Seven operating principles supported by the research distinguish AI leaders from AI aspirants. These are not theoretical; they are based on observable behaviors mapped to measurable business outcomes.

1

Data foundation before acceleration: this one is not optional.

The only non-negotiable layer in the framework, a centralized, governed knowledge infrastructure must precede using AI at scale. No exceptions. The principal question may shift from “have you built it?” to “how much do you need?” and the answer depends on your ambitions. But no amount of sophistication in the upper layers compensates for a broken or missing foundation.

2

Governance is a competitive function.

Organizations with mature governance suffer fewer violations, earn more client trust, and earn broader organizational latitude to deploy AI capabilities. Governance is not a constraint on AI ambition; it is the precondition for it.

3

Knowledge architecture is AI infrastructure.

Regardless of how sophisticated the model is, your information architecture determines how much value AI can extract. The data foundation dictates the height of the ceiling. The quality of the outputs is bounded by the quality of the inputs. Advanced investment in the knowledge layer raises the ceiling further. While not every organization needs to max out every layer, it does need to understand which layers are constraining its results.

4

Integration is the multiplier — but the baseline pays off regardless.

Eighty-five percent of organizations have adopted AI. Only 17 percent have integrated it. The gap is not about adoption; it is about architecture. That said, many of the workflow capabilities that drive daily productivity — such as collaboration, external sharing, and document automation — deliver real value independent of AI maturity. The goal is not to withhold workflow investment until AI infrastructure is perfect. The goal is to build the integration layer before attempting to scale AI tools across it.

5

Client expectations reshape your AI mandate.

Among the most mature organizations, 74 percent say their clients directly shape their AI strategy. This is not a future signal; it is a present one. The firms whose clients are demanding AI governance transparency are already competing on expectations that less mature organizations haven't yet begun to quantify.

6

The maturity dividend is measurable.

Twenty-eight percent versus 7 percent in the top financial performance quartile. Eighty percent versus 54 percent profitable. The return on sequenced, disciplined knowledge work maturity is not theoretical. It appears in the bottom line.

7

Sequencing determines ROI, not budget size.

The organizations that compound their AI advantage over time are not necessarily the ones spending the most. They are the ones spending in the right order. Transformation is a deliberate progression. Counterintuitively, racing ahead may slow progress and — ultimately — undermine success.

The path to AI excellence



This is an AI architecture puzzle – to which the research holds the key

The organizations that will win using AI over the next decade are not waiting for a better model. They are not waiting for AI regulation to stabilize. They are not waiting for competitors to show them what's possible. They are building the knowledge infrastructure that makes any model, any regulation, any competitive pressure something to leverage rather than survive.

Get started today to understand where you are, where you need to go, and how to get there.

Where is your foundation? Take the 15-question [AI readiness diagnostic](#). Identify your weakest layer. Don't commit new AI budget before you know what it's landing on.

What phase are you in? Map your current investments against the [four-phase maturity roadmap](#). Know what you are trying to build. Know what it requires before you try to build it.

Understand your shadow AI risk. Name an owner for AI policy, compliance, and risk today. Understand what AI tools are in use today. Make sure you understand and follow your AI guardrails. Take the [shadow AI risk audit](#) and identify your most urgent gap. Not because of liability, although that's important, but because strong governance is what allows everything else to scale.

Organizations that ask these questions – and answer them with objectivity – are on the path to AI excellence. They're building a foundation that continues to return value even when the current AI cycle is replaced by the next.

Get the data. See where investment flows — and why
AI adoption and governance define the winners.

iManage Knowledge Work Benchmark Report 2026

Download now



www.imanage.com



[Blog](#)

© Copyright 2026 iManage