


IA before AI

Identify the gap undermining your AI strategy

Inside: the case for why Information Architecture gates AI performance, a five-level diagnostic to score your organization's standing today, and a use-case map showing what's safe to attempt at each level.



Contents

Introduction.....

3

Part 1

6

Why IA sets the ceiling on what
AI can do

Part 2

9

Building outcome-focused
IA momentum

Part 3

13

Building better IA leads to
better AI outcomes

Conclusion.....

17

Getting your house in order

“You have to tidy the house before you have a party.”

Reena SenGupta, Founder, RSGi

AI performance is bounded by IA quality.

Information Architecture (IA) determines **what information AI can see, its context, how it understands relationships, and how reliably it can retrieve and apply knowledge.**

Strong IA turns AI from a probabilistic text generator into a dependable decision support system. Weak IA turns even the best models into confident guessers.

While successes may also be celebrated, every legal AI failure this year shares the same origin story: adoption moved faster than the information architecture underneath it could support. A firm piloted a tool, got a confident wrong answer, and blamed the model. But if the firm fails to clean the house before they start the party, the tool is not at fault.

The gap between how quickly organizations deploy AI and how ready their information is for AI exacts a cost in undesirable results, such as fabricated citations and “quoted” judgment text or employment clauses drafted under the wrong state’s law. AI pilots highlight the data work needed before going mainstream. None of that is a “model” problem. It is an information architecture problem, and it is entirely solvable.

That’s not rhetorical flourish. In 2024, [Stanford RegLab](#) found that general-purpose language models hallucinated on verifiable legal questions 69 – 88 percent of the time.

A follow-on Stanford study, [Hallucination-Free? Assessing the Reliability of Leading AI Legal Research Tools, March 2025](#), benchmarked several AI-driven legal research tools against GPT-4, which hallucinated 43 percent of the time when used on its own. The report found that when retrieval-grounded, curated legal content sat between the model and the answer, hallucinations were reduced but not eliminated. Both Lexis+ AI and Westlaw Ask Practical Law AI systems still hallucinated at 17 percent, while Westlaw’s AI-Assisted Research was nearly double that, at 33 percent. Same model class, different information layer, a materially different result.

What’s more, [Thomson Reuters’ 2026 Future of Professionals Report](#) suggests the same pattern holds well beyond legal research tools: 41 percent of professionals say they lack access to AI tools built on verified, well-governed content in the first place, while 91 percent say their organization is falling short of what AI could deliver as a result.

Information architecture is both the problem and the solution to the largely unrealized gains promised by artificial intelligence.

Part 1

Explains why IA sets the ceiling on what AI can do.

Part 2

provides a five-level framework to score where your organization sits today.


Part 3

Lays out which legal use cases are safe — or reckless — at each level.

Let’s get to it.

Part 1

Why IA sets the ceiling on what AI can do



1

IA determines what AI can “know”

AI systems do not reason over reality; they reason over **representations of information**.

Good IA:

- Defines authoritative sources
- Separates drafts from final content
- Clarifies ownership, jurisdiction, versioning, and applicability
- Encodes relationships (precedent → issue → outcome)

Without this, AI:

- Treats everything as equally valid
- Blends obsolete, conflicting, or irrelevant material
- Cannot distinguish signal from noise

Result:

AI accuracy improves not because the model is smarter, but because **the informational universe it operates in is coherent**.

2

IA reduces “hallucinations” (the real root cause)

Most AI hallucinations are **retrieval failures**, not reasoning failures.

Poor IA causes:

- Missing context
- Broken connections between concepts
- Undocumented assumptions
- Implicit knowledge trapped in people’s heads or free text

Strong IA:

- Makes implicit knowledge explicit
- Establishes clear retrieval paths
- Narrows the search space
- Anchors answers to traceable sources

Result:

AI **stops inventing answers** because it can find the right ones.

3

IA enables precision instead of probability

LLMs work probabilistically: they predict what sounds right, given the inputs.

IA introduces constraints, which shift AI from:

- “Most likely response” → to
- “Correct response within defined boundaries”

Examples of IA-driven constraints:

- Institutional knowledge type
- Guidance
- Templates
- Checklist
- Examples

Result:

AI output becomes **situationally accurate**, not just linguistically fluent.

4

IA turns AI from generic to context aware

Without IA, AI answers like a smart generalist. With IA, AI answers like **your organization**.

IA supplies:

- Taxonomies (what things are)
- Ontologies (how things relate)
- Metadata (why things matter)
- Workflow states (where things are in a process)

This allows AI to:

- Tailor outputs to role, task, and stage
- Respect internal standards
- Apply organizational judgment patterns

Result:

AI aligns with how **you** work, not how the **internet** works.

5

IA enables explainability and trust

In regulated domains, an answer is useless if it can't be explained.

Strong IA allows AI to:

- Cite sources
- Show reasoning paths
- Surface underlying assumptions
- Enable human review and override

Weak IA forces:

- Black box outputs
- Post hoc rationalization
- Manual reverification

Result:

IA makes AI **auditable, defensible, and adoptable**.

6

IA scales human judgment instead of replacing it

The most effective AI systems **augment judgment**, they don't replace it.

IA captures:

- Decision frameworks
- Risk heuristics
- Pattern recognition
- Exceptions and edge cases

AI then:

- Applies these consistently
- Flags anomalies
- Surfaces trade offs
- Accelerates human decision-making

Result:

AI scales how experts **think**, not just what they **produce**.

7

IA lowers cost and increases the ROI of AI

Organizations often overinvest in models and underinvest in structure.

Strong IA:

- Reduces prompt complexity
- Improves retrieval efficiency
- Lowers error correction costs
- Enables reuse across tools and use cases

Weak IA:

- Requires constant prompt tuning
- Generates rework
- Fails to generalize across workflows


Result:

IA is the **highest leverage AI investment** you can make.

The core principle

AI amplifies the quality of your information architecture.

If IA is weak, AI accelerates **failure**.



If IA is strong, AI accelerates **value**.

Next step


Knowing that IA sets the ceiling doesn't tell you where your own organization stands relative to it, or what's safe to attempt today. Part 2 turns that argument into a practical measurement.

Part 2

Building outcome-focused IA momentum

What follows is a practical information architecture (IA) development framework designed for AI-enabled, knowledge-intensive organizations (legal, professional services, regulated industries).

It is deliberately **outcome-focused**, rather than tool-centric, and shows how IA advancement unlocks better AI performance.



Where does your IA development stand today?

Move from information chaos to AI-ready intelligence

Level 1

Ad hoc information

“We have data, but no structure.”

Characteristics

- Information scattered across email, shared drives, DMS, personal folders
- No consistent naming, tagging, or versioning
- Knowledge exists primarily in people’s heads
- Search is unreliable and trust is low

AI impact

- High hallucination risk
- Inconsistent outputs
- Heavy human reverification required
- AI used only for generic tasks

Signals you’re here

- “I don’t know where the latest version is.”
- Different answers depending on who you ask
- AI demos look impressive but fail in practice

Level 2

Organized but fragmented

“We tidied up, but nothing connects.”

Characteristics

- Basic folder structures and taxonomies
- Some metadata, inconsistently applied
- Knowledge stored by system, not by use case
- Manual governance, low adoption

AI impact

- Improved retrieval in narrow contexts
- Still context-blind across matters, clients, or domains
- Prompt engineering compensates for weak structure

Signals you’re here

- Search works “sometimes”
- AI pilots succeed only with curated datasets
- Heavy reliance on power users

Level 3**Standardized information architecture**

“We agree on how information works.”

Characteristics

- Enterprise wide taxonomy and naming conventions
- Clear content types (draft vs final, advisory vs precedent)
- Defined ownership and lifecycle states
- Consistent metadata across systems

AI impact

- Reliable retrieval augmented generation (RAG)
- Fewer hallucinations
- Outputs are traceable and explainable
- AI usable in core workflows

Signals you're here

- People trust search results
- AI answers cite authoritative sources
- Reduced rework and duplication

Level 4**Contextual & judgment-aware IA**

“Information reflects how experts think.”

Characteristics

- Ontologies map relationships (issue → risk → outcome)
- Embedded business and legal context
- Decision frameworks and playbooks encoded
- Information architecture aligned to workflows

AI impact

- Context aware, situationally correct outputs
- AI supports decision-making, not just drafting
- Risk-sensitive recommendations
- Consistent application of expert judgment

Signals you're here

- AI adapts responses to role, matter, and stage
- Outputs feel “on-brand” and organization specific
- Experts trust AI as a thinking partner

Level 5**Adaptive, AI native information architecture**

“IA and AI continuously reinforce each other.”

Characteristics

- IA dynamically evolves based on usage patterns
- Feedback loops refine structure and relevance
- Knowledge flows across systems seamlessly
- Governance is embedded, not bolted on

AI impact

- Predictive and proactive AI
- Agentic workflows with guardrails
- Continuous learning without loss of trust
- Scalable, defensible AI at enterprise level

Signals you're here

- AI improves over time without manual reengineering
- Human oversight focuses on strategy, not correction
- IA is seen as core infrastructure, not support

Practical at-a-glance guidance

Level 1

Ad Hoc → AI guesses

Level 2

Fragmented → AI assists

Level 3

Standardized → AI delivers

Level 4

Contextual → AI advises

Level 5


Adaptive → AI augments judgment

How to use this framework:

- **Assess your current** level honestly (not aspirationally)
- **Target the next level**, not the final one
- **Invest first in structure**, not tooling
- **Treat IA as strategic infrastructure**, not content hygiene

Key principle

AI success is capped by
IA momentum.



Most organizations try to jump from Level 1 or 2 directly to Level 4 or 5 by adding more powerful models. It almost always fails. Why? Because it's not the AI model. It's the architecture.


Next step

Knowing where you stand now only matters if you know what it makes possible. Part 3 maps each of the five levels to the legal use cases that are genuinely safe today — and the ones that aren't "yet."

Part 3

Building better IA leads to better AI outcomes

Now we map IA levels to concrete legal use cases to show what becomes possible (and safe) at each level and to demonstrate where AI adds genuine value.



IA levels mapped to legal use cases

Level 1

Ad hoc information

Legal reality: Information chaos

AI role: Experimental, non-reliant

Viable legal use cases

- Generic legal research (public law only)
- First draft brainstorming (emails, outlines)
- Plain English explanations of known concepts
- Personal productivity (summaries, notes)

Not viable (high risk)

- Matter-specific advice
- Contract analysis
- Litigation strategy
- Client-facing outputs

Why? AI has no way to distinguish:

- Draft vs final advice
- Jurisdictional relevance
- Client context
- Authority or currency

Outcome: AI saves time only if humans redo the work.

Level 2

Organized but fragmented

Legal reality: Islands of order

AI role: Assisted execution

Viable legal use cases

- Internal research on curated datasets
- Clause lookup from limited libraries
- Document comparison and summarization
- Due diligence on scoped repositories

Conditionally viable

- Contract review (narrow playbooks)
- Precedent search (practice group specific)

Why? Some structure exists, but:

- Context doesn't travel across systems
- Metadata is inconsistent
- AI outputs depend on careful prompt control

Outcome: AI works in pilots and silos, not at scale.

Level 3

Standardized information architecture

Legal reality: Trustworthy information layer

AI role: Reliable delivery engine

Viable legal use cases

- Legal research with authoritative citations
- Contract drafting from approved templates
- Consistent clause risk flagging
- Matter summaries and chronologies
- Litigation document analysis (eDiscovery + synthesis)

Newly unlocked

- Retrieval-augmented generation (RAG)
- Explainable AI outputs
- AI embedded into core workflows

Why? AI can now:

- Reliably retrieve the **right** information
- Cite sources
- Respect versioning and jurisdiction
- Operate within governance boundaries

Outcome: AI reduces workload, not just effort.

Level 4

Contextual & judgment aware IA

Legal reality: Codified legal thinking

AI role: Decision support partner

Viable legal use cases

- Risk-graded contract review
- Playbook-driven negotiation support
- Litigation strategy analysis
- Regulatory impact assessment
- Advisory memos tailored to client profile

High value capabilities

- Issue → risk → outcome reasoning
- Context-sensitive drafting
- Exception handling and escalation
- Consistent application of firm judgment

Why? IA now encodes:

- Legal reasoning patterns
- Risk tolerances
- Practice-specific heuristics
- Client and matter context

Outcome: AI scales expert judgment, not just output.

Level 5

Adaptive, AI native information architecture

Legal reality: Intelligence infrastructure

AI role: Supervised autonomy

Viable legal use cases

- Agentic contract lifecycle management
- Proactive regulatory monitoring
- Continuous compliance validation
- Automated matter triage
- Predictive litigation and risk insights

Advanced capabilities

- Multistep legal workflows
- AI-initiated actions with guardrails
- Continuous learning from outcomes
- Real-time knowledge evolution

Why? IA and AI form a feedback loop:

- Usage refines structure
- Outcomes improve relevance
- Governance is embedded by design

Outcome: Lawyers focus on strategy, ethics, and judgment — not correction.

At-a-glance guidance revisited

Level 1

Ad Hoc → AI guesses

Level 2

Fragmented → AI assists

Level 3

Standardized → AI delivers

Level 4


Contextual → AI advises

Level 5

Adaptive → AI augments judgment

Tailor your strategy

IA development and AI strategy should reflect your unique needs.




Most legal AI failures stem from attempting Level 4 or 5 use cases with a Level 1 or 2 information structure. That means upgrading your AI model isn't the highest leverage move you can make: it's doing an IA momentum assessment and acting on what you learn.

The context we've provided can be a frame of reference to identify gaps that can help solve your AI investment issues. Score your organization honestly against the five levels, then commit to closing just one level, not all five. Once you know where you stand, decide what outcomes you want to achieve, and work incrementally toward the level that will enable you to attain that goal.

Conclusion

Getting your house in order



Critical insight for legal leaders

It bears repeating ... most legal AI failures happen because organizations attempt Level 4 or Level 5 use cases while operating at Level 1 or 2 IA maturity.

This creates:

- Hallucination risk
- Ethical exposure
- Reputational damage

Across every level of this model, one pattern holds: AI performance, not model choice, reflects IA maturity. A better model layered on chaotic information still produces guesses; a modest model layered on standardized, contextual information produces trustworthy output. That is the entire argument in one sentence, and it is why swapping in a better model is not a fix for information that isn't AI-ready.

The Reena SenGupta quote we started with is the same argument, in metaphor. Getting your house in order is what the five levels above describe: authoritative sources, clear ownership and versioning, encoded relationships, judgment made explicit rather than left in someone's head. And to keep the metaphor going, we've woven iManage Context Fabric™ through our platform to underpin the AI equivalent of a tidy house. It is a governed intelligence layer that reads and connects an organization's matters, documents, people, and in-flight activity by translating its taxonomies, ontologies, metadata and workflow states into something AI can reason over, not just store.

Context Fabric is the difference between a lawyer reconstructing an unfamiliar matter from a random email and stepping into it fully briefed. It's a draft clause drifting quietly away from the firm's playbook, versus being flagged the moment it begins to drift so someone can decide whether to accept the deviation or bring it back in line. None of that replaces the momentum work we lay out in this guide. Rather, it's what that work is for.

Shall we get this party started?

See the [next-generation iManage platform](#) in action: reach out to your iManage representative or book a demo now.

Book a demo

iManage is dedicated to Making Knowledge Work™. Built with more than 30 years of experience serving the world's leading law firms, corporate legal teams, and professional services organizations, iManage gives knowledge workers a governed foundation to capture, manage, and build on the work that drives better decisions. As your strategic partner, we employ our award-winning AI-enabled technology, an extensive partner ecosystem, and a customer-centric approach to provide support and guidance you can trust to make knowledge work for you. iManage is relied on by more than one million professionals at 4,500 organizations around the world.

Visit www.imanage.com to learn more.

