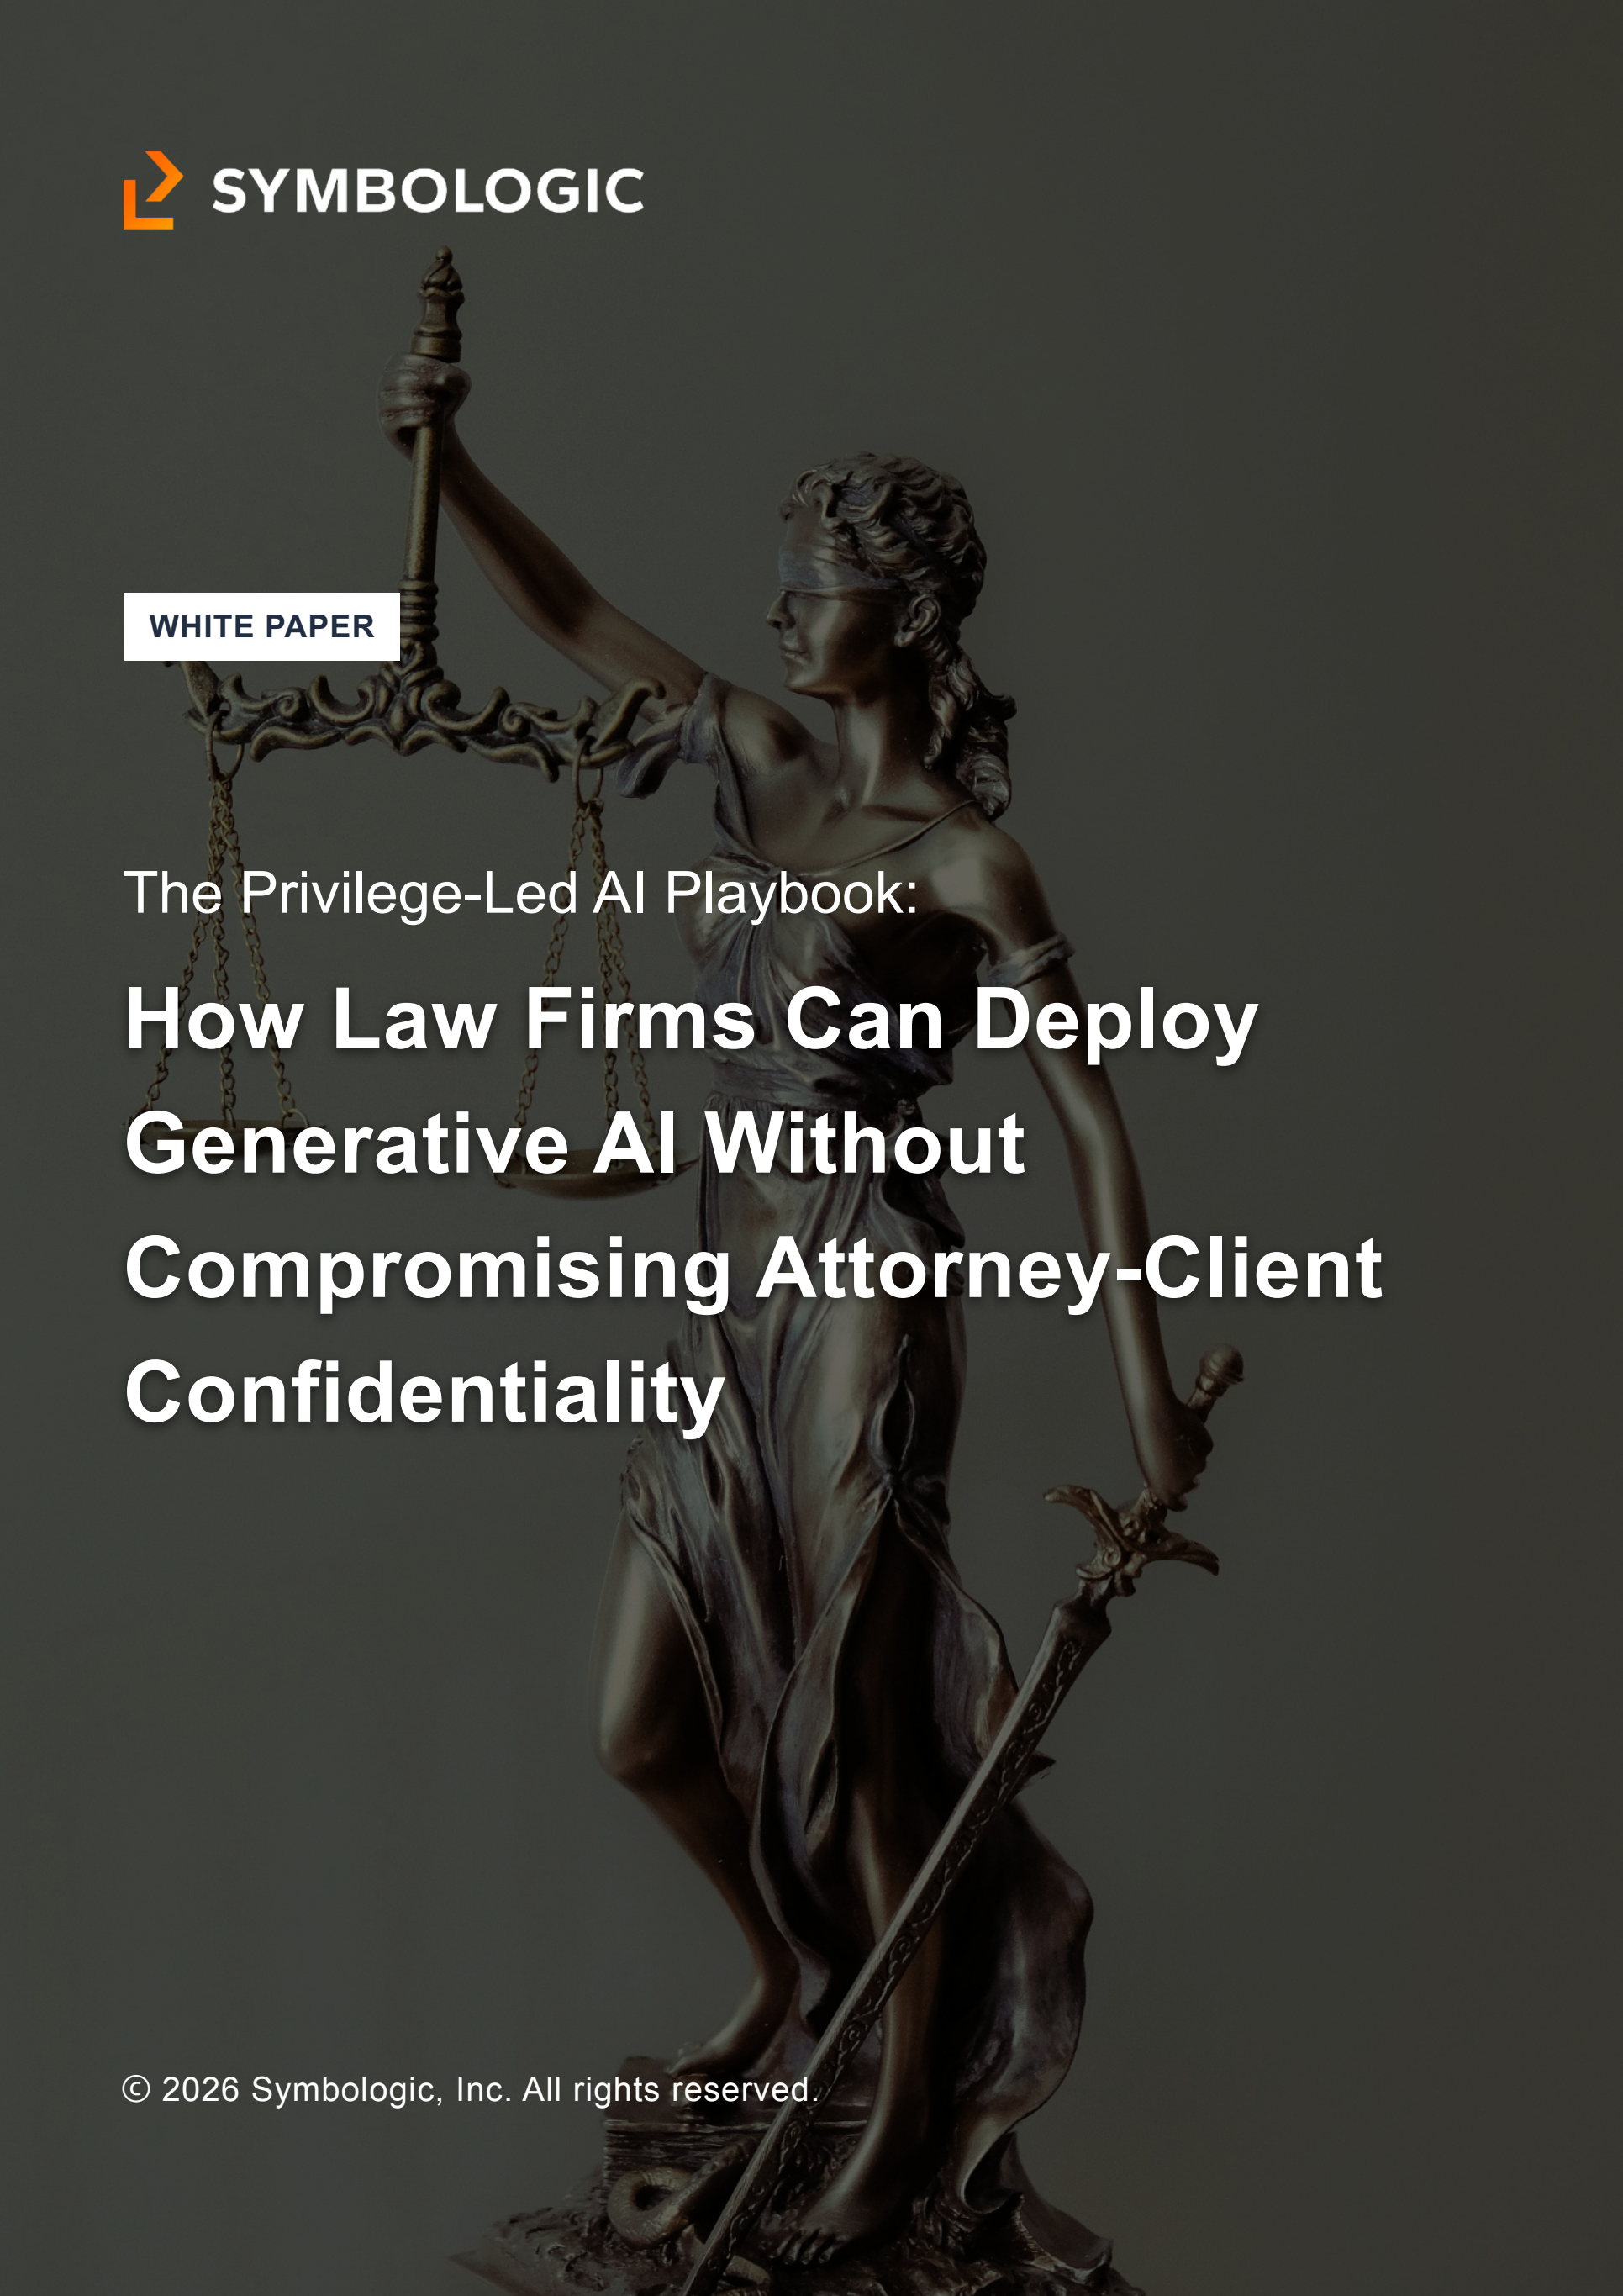
SYMBOLOGIC



WHITE PAPER

The Privilege-Led AI Playbook:

How Law Firms Can Deploy Generative AI Without Compromising Attorney-Client Confidentiality

© 2026 Symbologic, Inc. All rights reserved.

Executive Summary

Every Am Law 200 firm is under pressure to deploy AI. Many are stuck.[1]

The reason is not budget. It is not skepticism about AI's potential. It is a structural problem that no amount of vendor negotiation resolves: associates use consumer AI tools without IT's knowledge. According to the 2025 Thomson Reuters Institute, “The Future of Professionals” report[2], a third of law firm professionals use AI tools their firm has not authorized for work.

The report highlights what it calls "the risk of introducing unvetted technology into an environment of sensitive, privileged information" and concludes that "the firm carries that exposure whether or not it has chosen to".

The result is a familiar pattern. Pilots stall at the security review. Partners keep asking when AI will actually be available for real client work. Meanwhile, years of institutional knowledge can sit underused inside content repositories such as iManage Work or NetDocuments, accessible to the teams who know exactly where to look, but effectively out of reach for everyone else.

This paper argues that the problem is not AI. It is architecture. A different architectural approach materially reduces the privilege risk: deploy AI that goes to where the data lives, not the other way around. No architecture can substitute for sound firm policy and lawyer judgment, but it can remove the key risk that policy alone cannot fix: disclosure to a third party at the moment of ingestion.

With this approach, a firm can query its full library, whether housed in iManage or another document repository, in natural language, receive sourced answers in minutes rather than hours, and accomplish all of this without a single document leaving the firm's controlled environment. In addition, the results of those queries then become part of the firm's institutional knowledge and easily available to anyone else with the same (or a similar) query, scoped, like every query, to the requesting user's own iManage permissions. Deployment takes one business day. Because no new data environment is created, the sequential legal, information-security, and procurement sign-off that normally gates enterprise AI adoption shrinks to a review of self-hosted software rather than a data-governance negotiation.



© 2026 Symbologic, Inc. All rights reserved.

1. The AI Imperative in Legal Services

1.1 The Competitive Pressure Is Real and Accelerating

Legal AI has crossed the chasm from early adopter curiosity to mainstream procurement. According to an August 2026 story in the ABA Journal[3], Am Law 100 law firms have said they're spending more on artificial intelligence and innovation investments than on any other expense category. Boutique practices are winning business by advertising AI-enabled turnaround times. Corporate legal departments are actively scoring outside counsel on technology efficiency.

The economics are straightforward. When associates billing at $600 an hour spend a significant share of their time on information retrieval rather than legal judgment, clients notice. And they push back. Firms that can demonstrate AI-driven efficiency without compromising quality are taking mandates from firms that cannot.

This is no longer a future-state conversation. The gap between early movers and laggards is widening every quarter.

1.2 The Knowledge Access Gap

Long before AI entered the picture, law firms quietly accepted a structural productivity tax: decades of institutional knowledge sitting inside iManage or other document repositories: searchable, but not synthesizable.

The pattern is universal. Consider a scenario where a senior partner handled a complex indemnification clause dispute in 2021. The analysis, the negotiating history, and the outcome are all in iManage or a similar repository, buried in a workspace that three associates on a related case cannot access. In 2026, a junior associate on a different matter faces the same issue. The research starts from scratch.

This happens thousands of times a year across every large firm. The knowledge exists. It simply is not accessible at the moment it is needed, to the people who need it.

© 2026 Symbologic, Inc. All rights reserved.

1.3 Why Adoption Has Stalled

The opportunity is obvious. The business case is clear. So why is there an “execution gap” between law firms adopting AI and using it effectively?[4]

Because every mainstream enterprise AI platform carries a structural requirement that collides directly with attorney-client privilege. Understanding why requires looking at how these platforms actually work.

2. The Attorney-Client Privilege Barrier

2.1 The Fundamental Constraint

Attorney-client privilege is not a compliance checkbox. It is a load-bearing pillar of the legal profession. Communications between attorney and client made in the course of seeking legal advice are protected from compelled disclosure. That protection can be waived, and it can be waived without anyone intending it.

The question every law firm faces when evaluating AI platforms is this: does sending client documents to an AI platform's servers constitute a disclosure to a third party sufficient to trigger waiver?

Courts have now answered that question directly.

2.2 What the Courts Are Saying

In February 2026, Judge Jed S. Rakoff of the Southern District of New York ruled in United States v. Heppner[5] that documents generated using a public generative AI platform were not protected by attorney-client privilege. The holding rested on three findings:

- **No attorney relationship.** The privilege attaches to communications between a client and counsel. A consumer AI platform qualifies as neither.

© 2026 Symbologic, Inc. All rights reserved.

- **No confidentiality.** Using a public AI platform eliminated the confidentiality element required for privilege to attach, because the consumer tool's own terms reserved the right to use inputs for model training and to share data with third parties, including government authorities. On those terms, there was no reasonable expectation of confidentiality.
- **No work product protection.** The work product doctrine did not protect the materials because they were not prepared at the direction of counsel and did not reflect defense strategy.

Judge Rakoff's opinion explicitly addressed a “question of first impression nationwide.” It is already being cited across bar association ethics opinions, practice group guidance, and law school curricula. It is not an outlier. It is the opening chapter of a judicial framework that will govern how firms approach AI for the foreseeable future.

Key Principle from Heppner: *Courts will apply traditional third-party disclosure analysis to AI, and the confidentiality of the environment that receives client information is decisive. Whether enterprise contractual protections preserve privilege once documents are copied into a vendor's AI environment is a question no court has yet answered. No firm wants its client's matter to be the test case.*

2.3 Why the Standard Vendor Response Falls Short

AI vendors are aware of these concerns. The typical response is a data processing agreement, a zero-retention policy, and an enterprise pricing tier marketed as secure. These protections are genuine and worth reviewing carefully, but they do not solve the problem.

These protections may well matter: Heppner turned on consumer terms that expressly undermined confidentiality, and enterprise agreements are written to do the opposite. But whether they are sufficient to preserve privilege once client documents have been copied into a vendor's environment is a question no court has decided. A firm that uploads privileged documents to a third-party AI platform is betting the privilege on an untested waiver defense, and it will learn whether the bet paid off only in litigation, with a client's matter as the vehicle.

Contractual protections address data security, and they strengthen the argument that confidentiality was maintained. What they cannot do is prevent the threshold event from occurring: a protected communication entering a third party's environment. An architecture that never creates that event has nothing left to argue.

© 2026 Symbologic, Inc. All rights reserved.

2.4 Ethical Obligations Add Independent Risk

Beyond the privilege question, ABA Model Rule 1.6 imposes a separate obligation[6]: lawyers must make reasonable efforts to prevent the inadvertent or unauthorized disclosure of client information. Comments to the rule address cloud-based services specifically, requiring lawyers to investigate the service's nature, the vendor's data access, and the security measures in place before using it for client matters. The ABA has since applied this obligation directly to generative AI: Formal Opinion 512[7] confirms that the same confidentiality duty governs a lawyer's use of AI tools in client work.

State bar associations in New York, California, and Texas have each issued formal ethics opinions applying this framework to generative AI tools. Texas's Opinion 705[8], issued in February 2025, grounds these obligations in existing duties of competence, confidentiality, and candor rather than creating new ones. The analysis is consistent across jurisdictions: using a third-party AI platform for client work without adequate due diligence and, in many circumstances, client consent creates Rule 1.6 exposure entirely independent of whether privilege is ultimately waived.

The cumulative effect is a compliance posture that most firms are not comfortable accepting. And so AI deployments stall.

3. The Hidden Cost of Staying Stuck

3.1 The Cost of Asking Twice

Most large organizations already store their work in iManage or a similar repository, and most of those repositories now ship with AI built-in to help users find and work with what is already there. That is real progress, and it is worth acknowledging directly. According to iManage's own March 2026 growth report[9], seventy-nine percent of the Am Law 100 (the top half of the Am Law 200 by revenue) anchor their AI strategy on its platform. It enforces access control at the workspace and document level, maintains version history, and supports compliance workflows. It is exactly the right tool for document management.

But being able to ask a question is not the same as reliably getting the right answer on the first try. A partner asking “How have we handled arbitration clauses in buy-side M&A agreements over the past five years?” does not want a plausible-sounding summary that might be wrong. They want an answer they can act on without checking it by hand, or the work effectively happens twice: once by the AI, and once by the lawyer confirming it.

© 2026 Symbologic, Inc. All rights reserved.

That second pass is where the real cost hides, and it hides in three places. It shows up as token spend, when the first answer was not good enough and the same question gets rephrased and resubmitted until it is. It shows up again, multiplied by headcount, when different associates across the firm ask a version of the same question in the same month and each one pays the full cost of synthesis from scratch, because nothing in the system remembers that the question was already answered. And it shows up as associate hours, spent double checking an AI-generated summary before it goes in front of a partner or a client. None of these costs appears on a vendor's pricing page. All of them show up on the matter's bill.

iManage and other similar solutions have been investing in native AI capabilities of its own to help close this gap, a good-faith effort worth acknowledging here. However, Section 5.3 explains what it takes to get a synthesized answer right on the first attempt, and Section 7 explains why the cost of an already-answered question should not have to be paid again.

3.2 The Compounding Cost of Duplicated Work

Every matter pays the cost of siloed institutional knowledge. Associates duplicate research that partners already completed on prior matters. Clients pay for that duplication. Partners supervise research that did not need to take as long. The firm's accumulated legal judgment, embedded in years of closed-matter documents, can sit inaccessible to the teams who most need it.

Knowledge management teams have tried to close this gap with practice group wikis, know-how databases, and matter summaries. These efforts are valuable and perpetually under-resourced. They depend on voluntary lawyer contributions at the end of already long workdays, and they degrade as personnel changes and practice areas evolve.

The opportunity cost compounds silently, and it compounds twice. Every matter that starts from scratch (instead of building on the firm's own precedent) costs more than it needed to: in associate hours, if the research is done by hand, and in token spend and re-verification, if it is rushed through an AI that does not get it right the first time. That compounding cost, search on one side and unreliable synthesis on the other, translates directly into competitive positioning, associate productivity, and client perception of value.

© 2026 Symbologic, Inc. All rights reserved.

4. Why Conventional AI Architectures Do Not Solve This

4.1 The Vacuum Cleaner Model

Every mainstream enterprise AI search and knowledge platform operates on the same underlying architecture. To answer questions about a firm's documents, those documents must be copied to the platform's infrastructure and processed into a searchable index. The AI model and the data need to be in the same environment, which means the data has to move.

This is not a design flaw. It is the only way retrieval-augmented generation works when the model and the data start in different places. The documents have to travel to the AI. Every conventional platform is built this way.

For law firms, the consequences cascade across four dimensions:

- **Privilege exposure at ingestion.** Privileged documents enter a third party's environment at the moment they are copied to the platform. Under the disclosure-based analysis Heppner applied, the firm's privilege position then rests entirely on contractual defenses that no court has yet tested. Heppner is explicit: no subsequent vendor commitment undoes this disclosure.
- **Dramatically expanded breach surface.** A centralized index of sensitive client documents across every matter and every practice group becomes a single high-value breach target. One platform compromise exposes everything.
- **Regulatory scope expansion.** Moving data to a third-party environment can trigger GDPR, state privacy law, and sector-specific compliance obligations that never applied when data remained within the firm's infrastructure.
- **Cost that scales with headcount, not with need.** A conventional platform treats every query the same way: a fresh model inference, run from scratch, whether the question needs reasoning or has one correct answer sitting in a database, and whether or not someone else at the firm already asked something similar. Token spend rises in a straight line with usage, with no structural way to recognize a repeat question or route a simple one around the most expensive path.

© 2026 Symbologic, Inc. All rights reserved.

4.2 The Binary Choice No Firm Wants to Make

The practical result is a forced choice: accept the privilege and ethics risk, or limit AI to generic tasks that do not touch client matters.

Firms that choose the latter are effectively using AI as a slightly better search engine. Generic research. Writing assistance. Internal memos. They are capturing perhaps 20 percent of the productivity gain AI can deliver in legal practice, while their most aggressive competitors are running faster.

This is not a sustainable position. And it is not a necessary one.

So what is the alternative? If contractual promises cannot cure a structural disclosure problem, and limiting AI to generic, non-client tasks leaves most of its value on the table, the only way out is an architecture that never creates the disclosure in the first place. That is the premise the rest of this paper takes up.

5. A Different Architecture

Symbologic Librarian is Symbologic's answer to the architecture problem described above: an AI layer that runs inside the firm's own environment and works directly against content already stored in iManage or another document repository, without copying, indexing, or relocating any of it. The sections below explain how it works (5.1), how it extends the firm's existing access controls (5.2), how it handles two fundamentally different kinds of legal questions accurately (5.3), and, now that the architecture has been introduced, how it compares structurally to conventional AI platforms (5.4).

5.1 The Lens Approach: AI Goes to the Data

Symbologic Librarian starts from a different premise. Rather than pulling data to where the AI lives, Librarian discovers and works with the data where it lives, like viewing the firm's existing repository through a different lens, without ever moving what it sees. Users ask questions, and Librarian answers by analyzing all relevant information. The data never moves, but when the data changes, Librarian refreshes the indices it maintains, and those indices, like Librarian itself, reside inside the firm's own environment.

© 2026 Symbologic, Inc. All rights reserved.

Before walking through a specific query, it helps to define one term. The intent router is the component that classifies each incoming query, by subject matter and by whether it calls for synthesis or a precise, computed answer, and directs it to the appropriate processing path (see Section 5.3).

When a lawyer asks “What positions have we taken on limitation of liability clauses in enterprise SaaS agreements?”, the following occurs:

1. **Classify.** The query is classified by Librarian's intent router.
2. **Query.** A structured question is constructed and processed by Librarian's reasoning engine, which ensures the latest versions of data are pulled from iManage's native API.
3. **Retrieve.** iManage returns the relevant data within the boundaries of the user's existing permissions.
4. **Respond.** Librarian synthesizes a response and returns it with citations to specific iManage documents.

At no point does a document leave iManage. No copy is created on Symbologic's infrastructure, because there is no Symbologic infrastructure in the loop: Librarian, its indices, and its caches all run inside the firm's environment. Where model inference runs is the one remaining variable, and Section 5.3 places that choice, including a fully in-boundary option, in the firm's hands.

The documents never entered a third-party environment, so the third-party disclosure question that Heppner turns on, and that conventional platforms must answer with untested contractual defenses, never has to be answered at all. On the self-hosted posture described in Section 5.3, the same is true of the model itself.

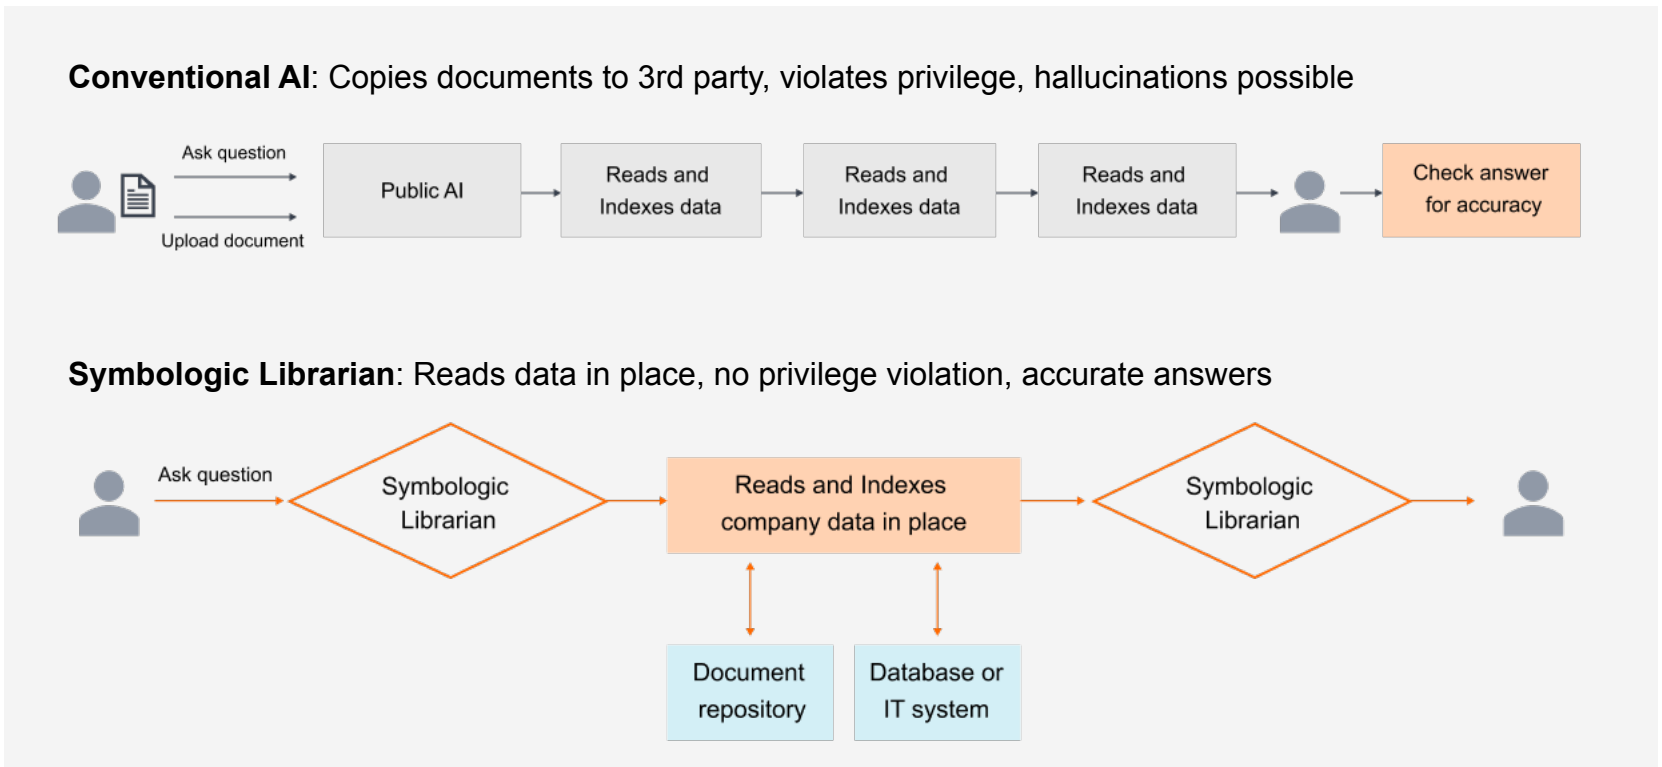


Figure 1. Conventional AI architecture vs. Symbologic Librarian's Lens Approach.

© 2026 Symbologic, Inc. All rights reserved.

5.2 Access Control That Extends What the Firm Already Built

Librarian does not introduce a parallel permission system. It embraces the access controls the firm has already built in iManage and extends them automatically to every AI query, rather than asking IT to define a new rulebook for AI on top of the one that already governs human access.

The practical consequence matters enormously in legal practice. If an associate does not have access to a particular matter workspace in iManage, Librarian cannot surface documents from that workspace in response to the associate's queries. Ethical walls, conflicts screens, and matter-level access restrictions all carry forward automatically. No new permissions need to be defined. No additional configuration is required. Because Librarian enforces exactly the same boundaries iManage already enforces, it cannot surface documents a user could not otherwise reach through iManage itself.

For firms that have invested years in carefully configuring iManage access controls, this is significant. That investment extends to AI without any additional work.

5.3 Deterministic Accuracy for Questions That Cannot Afford to Be Wrong

Legal work involves two fundamentally different types of questions. Librarian's intent routing engine classifies every query at submission and routes it to the appropriate processing path:

Question Type	Example	How Librarian Handles It
Synthesizing	“What damages theories have we advanced in trade secret misappropriation cases?”	LLM working against iManage content with source citations
Deterministic	“Has the exclusivity period in this agreement expired?”	SQL direct execution, deterministic computation over extracted, source-cited terms a lawyer can verify

A lawyer asking “Has the payment cure period in Section 12(b) passed?” receives a precise answer derived directly from the contract, not a probabilistic estimate that might be right 94 percent of the time. In legal practice, the other 6 percent matters. Librarian routes those questions around the uncertainty entirely.

© 2026 Symbologic, Inc. All rights reserved.

Librarian is model-agnostic by design, and where inference actually runs is a deployment decision the firm makes and can verify, not a default the vendor bakes into the product. Three postures are available. A managed cloud endpoint is fastest to stand up, but the query and its retrieved evidence leave the firm's network on every call, under contractual no-training terms. A private, in-region endpoint dedicated to the firm's own tenancy keeps everything inside the firm's cloud account and jurisdiction. A fully self-hosted, open-weight model running on the firm's own hardware keeps inference entirely inside the firm's boundary; no external party is involved at all. A firm can deploy the environment handling its most sensitive privileged matters, such as client contracts, negotiation history, and litigation strategy, on the self-hosted posture specifically, so that queries touching that work are never transmitted to a third-party model provider in the first place.

The same three postures are worth asking about for any AI feature that touches privileged content, including features built into tools a firm already licenses, not only standalone platforms.

5.4 The Architecture Gap: Conventional Platforms vs. Symbologic Librarian

The table below illustrates the structural difference across six dimensions that matter most to legal IT and risk leadership.

Criteria	Conventional AI Platforms	Symbologic Librarian
Data Movement	Documents copied to vendor infrastructure	No documents leave iManage
Privilege Risk	Documents enter vendor environment at ingestion; privilege rests on untested contractual defenses (per Heppner)	No ingestion-time disclosure; third-party model exposure is a deployment choice, not a default
CISO Review	Sequential legal/InfoSec/procurement sign-off required	No new environment to audit
Deployment Timeline	Months	One business day
Access Control	Separate permission layer required	Extends iManage ACL automatically
Breach Surface	Centralized index is a high-value target	No copy exists outside the firm

© 2026 Symbologic, Inc. All rights reserved.

6. What Deployment Actually Looks Like

6.1 One Day. No Migration. No Infrastructure Outside Your Walls.

A firm with iManage can have Symbologic Librarian operational in one business on existing infrastructure. Self-hosted open-weight inference is assessed during onboarding.

This is possible because no data moves and nothing runs outside the firm's own environment. Librarian is not a multi-tenant SaaS product that firms send data to. It is self-hosted software, deployed on the firm's own infrastructure, whether that is Kubernetes on AWS, Azure, or Google Cloud VPC, on-premises, or purpose-built AI hardware such as NVIDIA DGX. The firm's data never leaves its own tenant, because Librarian never leaves the firm's tenant either.

There is no migration to plan, no new repository to build, and no access control system to reconfigure. Deployment consists of connecting Librarian to iManage through the iManage native API. For firms that have completed their iManage implementation, the technical work is straightforward.

This is a meaningfully different deployment profile than any conventional enterprise AI platform, which requires data ingestion, index building, security review, and often a months-long change management process before a single lawyer can run a query.

6.2 CISO Review: Not a Bottleneck

The security review that gates every conventional AI deployment exists for a specific reason: the platform needs access to sensitive data in a new environment, so legal, information security, and procurement each have to sign off in sequence, reviewing encryption, vendor certifications, breach notification procedures, and contractual liability for data leaving the firm.

Few of those questions arise with Symbologic Librarian, because no new data environment is created: the software runs inside the firm's own tenant, and from a data-flow perspective it behaves like a sophisticated user running queries against iManage. The firm's existing controls remain fully operative. What remains for the CISO is the review any self-hosted software holding repository credentials deserves, covering code provenance, the update mechanism, credential handling, and the logging and retention of queries and answers, which themselves contain privileged strategy. That is a scope measured in weeks, not the months a new data environment typically demands.

© 2026 Symbologic, Inc. All rights reserved.

The review that typically requires that full sequential sign-off becomes a standard vendor due-diligence conversation, because the question at the center of it (where is our data going?) has a one-word answer: nowhere. That is the difference between an AI practice that is live this quarter and one that is perpetually in security review.

6.3 Practice Group Use Cases

Litigation. Associates query iManage across all accessible matters using natural language. “What damages arguments have we advanced in ERISA class action cases?” returns a synthesized response with citations to specific documents. Research that took hours takes minutes. The work is traceable, auditable, and grounded in the firm's own precedent rather than a general web search.

Corporate and Transactional. Contract review combines natural language synthesis for contextual analysis with deterministic routing for precise condition checks. Librarian surfaces non-standard clause language by comparing against the firm's standard position documents in iManage, with none of those documents leaving the firm's environment.

Intellectual Property. Prior art research, prosecution history review, and claim construction work draw on the firm's accumulated matter history. Associates can query the firm's experience with specific claim elements across the full iManage library, scoped to their permissions.

Real Estate and Finance. Complex transaction document sets including loan packages, title materials, and regulatory filings become navigable through natural language queries. “What conditions remain unsatisfied in this deal?” produces a structured response built from the transaction documents in iManage, not from a language model's general knowledge.

© 2026 Symbologic, Inc. All rights reserved.

7. What Firms Can Expect

The value of Symbologic Librarian flows directly from the architecture. Because no data moves and no index is built outside iManage, the deployment friction that consumes months with conventional platforms is eliminated on day one.

Area	Without Librarian	With Librarian
Case Research	Hours of manual iManage navigation per query	Minutes via natural language with source citations
Precedent Access	Limited to workspaces the associate knows to search	Full iManage library, scoped to user permissions
Contract Review	Associate reviews agreements without firm precedent	AI surfaces comparable prior positions from iManage
Deployment	Months of data ingestion and security review	One business day from integration to first query
Access Control	Must be configured separately for AI layer	Extended automatically from iManage ACL
Privilege Posture	Documents disclosed to vendor infrastructure	No documents leave the firm's environment

Case research changes fundamentally. Today, an associate researching how the firm has handled a specific clause type across prior matters runs keyword searches in iManage, opens documents one by one, and synthesizes findings manually. The process takes hours and depends heavily on knowing which workspaces matter in the context of what they are researching. With Librarian, the same question is asked in natural language and answered in minutes, with citations to the specific iManage documents that support each point. The associate's time shifts from retrieval to analysis.

Contract review becomes a collaborative process between the lawyer and the firm's own precedent. Rather than reviewing a new agreement in isolation, lawyers can ask Librarian what positions the firm has previously accepted on specific clause types, which prior agreements contained similar language, and how those negotiations were resolved. The institutional knowledge that currently lives in senior partners' memories becomes accessible to every associate working within their permitted matter scope.

© 2026 Symbologic, Inc. All rights reserved.

Cost Predictability, Not Cost Creep. Technology leaders are under growing pressure to show measurable return on AI investment: in Deloitte's 2026 Global Technology Leadership Study[10], 79% of technology leaders name delivering measurable enterprise value as their top priority, yet 42% report low or no ROI on their AI investments so far. Unpredictable per-query token spend is a major driver of that gap. Librarian is built to keep AI cost predictable as usage scales: repeated questions are served from cached responses instead of being re-run against a model, answers are shared across users asking similar questions instead of each user paying for the same synthesis twice, and deterministic questions are routed to direct SQL execution with no model inference cost at all. A firm scaling Librarian from one practice group to the whole firm is not multiplying its token bill by the same factor.

The deployment experience itself sets a different expectation. A firm that connects Librarian to iManage on a Monday and has lawyers running queries by Tuesday morning has a fundamentally different relationship with its AI vendor than one that requires an exhaustive, multi-month security review before a single query was run. The ability to start, evaluate, and iterate quickly changes how firms think about AI adoption entirely.

Governance is built in from the start, not bolted on later. Because Librarian embraces and extends the firm's existing iManage access controls rather than introducing a parallel system, every query is already bound by the permissions the firm established for human access. There is no separate permission system to maintain, no gap that could let AI surface documents across an ethical wall, and no audit gap between what lawyers can access and what Librarian can return.

8. The Case for Acting Now

8.1 Inaction Has Its Own Risk Profile

Firms waiting for the legal AI landscape to clarify before committing face a risk that is easy to underestimate: their lawyers are not waiting.

Consumer AI tools are already in use across every large firm, on personal devices, through individual accounts, outside any IT visibility or governance framework.[11] The firm's caution is not preventing AI adoption. It is ensuring that AI adoption happens in the worst possible way, without privilege protections, without access controls, and without any audit trail.

The question for firm leadership is not whether to have an AI strategy. It is whether to have a governed AI strategy before the ungoverned alternative becomes entrenched.

© 2026 Symbologic, Inc. All rights reserved.

8.2 The Firms That Will Win This Decade

The firms building durable competitive advantage through AI are not necessarily the fastest movers. They are the firms moving correctly, in a way that clients, bar associations, and partners can uniformly endorse.

A firm that can tell its clients "Your documents never leave our infrastructure, and for our most sensitive matters, neither does the AI itself: the model runs entirely inside our own boundary. Every query is bound by the same access controls that govern human access to your files" is making a claim no conventional AI platform can match.

Symbologic Librarian is built to make that claim defensible, technically, legally, and operationally.

Conclusion

The attorney-client privilege problem is not an obstacle to AI in legal services. It is a design specification.

The right AI platform for law firms must, by construction, leave privileged documents in place and extend existing access control frameworks. It must operate in a way that is built to withstand the analysis Heppner established. Contractual workarounds do not meet that specification. An architecture that never creates the disclosure in the first place is a far stronger starting point.

Symbologic Librarian meets it. Queries go to the data. Data stays in iManage or other document repositories. Access control follows rules the firm already established. Deployment completes in a day. Cost follows the same logic: response caching, shared answers, and deterministic routing for analytical queries keep AI spend predictable as usage scales.

The result is an AI practice that is ready on day one, governed on day one, and defensible to every audience that matters.

© 2026 Symbologic, Inc. All rights reserved.

References

1. American Bar Association, “The Legal Industry Report” (2025) and “What the 8am Legal Industry Report Tells Us About AI Use,” ABA Law Practice Magazine (March/April 2026). Legal industry report data: 69% of legal professionals report personal use of generative AI tools, while only 46% of firms have adopted general-purpose AI tools firm-wide, 34% have adopted legal-specific AI platforms, and 43% report no formal AI policy.

2. Thomson Reuters Institute, “The Future of Professionals” report (2025). Amongst the key findings of the report were that 34% of professionals use AI their organization hasn't approved, creating invisible risks organizations cannot monitor or control.

3. ABA Journal, “Rising AI expenses could hinder BigLaw performance, new survey shows,” August 10, 2026 news report.

4. Thomson Reuters Institute, “The Future of Professionals” report (2025) executive summary. The report warned that “Law firm leaders have moved past AI adoption; now they must close the execution gap across strategy, client delivery, governance, and talent before it becomes a competitive risk.”

5. United States v. Heppner, No. 25-cr-00503-JSR (S.D.N.Y. Feb. 17, 2026). Judge Jed S. Rakoff ruled that documents generated using a public generative AI platform were not protected by attorney-client privilege or the work product doctrine. The opinion addressed “a question of first impression nationwide.”

6. American Bar Association, Model Rules of Professional Conduct, Rule 1.6: Confidentiality of Information (2023). Requires lawyers to make reasonable efforts to prevent unauthorized disclosure of client information, including when using third-party cloud-based services.

7. American Bar Association, Formal Opinion 512: Generative Artificial Intelligence Tools (July 2023). Addresses lawyers' ethical obligations when using generative AI, including confidentiality duties under Rule 1.6.

8. State Bar of Texas AI Toolkit includes a link to Opinion 705 of the Professional Ethics Committee for the State Bar of Texas, which addresses ethical issues under the Texas Disciplinary Rules of Professional Conduct raised by a lawyer's use of generative AI, including confidentiality (Rule 1.05) and competence (Rule 1.01).

9. iManage Reports Strong Global Growth as Organizations Anchor AI Investments in a Trusted Knowledge Foundation,” Yahoo Finance/GlobeNewswire (Mar. 10, 2026). Reports that 79% of the Am Law 100, 83% of the Global 100, and 40% of the Fortune 100 anchor their AI strategy on iManage's platform.

10. Deloitte, 2026 Global Technology Leadership Study (April 2026). Survey of 660+ technology leaders: 79% prioritize measurable enterprise value from AI investment, while 42% report low or no ROI to date.

11. North Carolina Bar Association Blog, “Beyond the Ban: Why Your Law Firm Needs a Realistic AI Policy in 2026. Noted that while 79% of legal professionals utilized AI tools, 44% of law firms had not yet implemented formal governance policies.

© 2026 Symbologic, Inc. All rights reserved.

Additional resources

1. State Bar of California, Practical Guidance for the Use of Generative Artificial Intelligence in the Practice of Law (2024). Outlines duties of confidentiality and competence applicable to AI tool selection and use.

2. Gibson, Dunn & Crutcher LLP, “AI Privilege Waivers: SDNY Rules Against Privilege Protection for Consumer AI Outputs” (February 2026). Analysis of the Heppner ruling and implications for enterprise AI deployment in legal practice.

3. McDermott Will & Emery, “Using AI Without Waiving Privilege: Lessons from United States v. Heppner” (March 2026).

SYMBOLOGIC

Symbologic's secure, enterprise AI solutions and professional services deliver semantically and analytically accurate conversational intelligence within existing client infrastructure. Combining Data-in-Place architecture with zero-trust security, Symbologic AI solutions and professional services transform siloed enterprise data into continuously expanding, reusable corporate memory, empowering your employees while reducing the cost of AI.

© 2026 Symbologic, Inc. All rights reserved.