



VENIO SYSTEMS BRIEFING FOR LAW FIRM TECHNOLOGY LEADERS

AI Adopted. Now Defend It.

What the 2026 data means for the people who own a law firm's AI stack, and a working playbook for the data where the stakes are highest: discovery.

September 2026. Built on independent and third-party research, with every figure sourced and dated.

94%

of law firms are using or exploring
generative AI tools

18%

of professionals know their
organization tracks AI return

1,397

US decisions to date address
AI-fabricated material in filings

Your firm has already adopted AI. The question now is whether you can show what it touched, what it cost and why its output can be trusted.

Partners will ask that question about the budget. Clients will ask it about their data. And in 2026, courts increasingly asked it about discovery material, in protective orders that read like technical specifications.

This briefing sets out what the latest research shows, what each finding means for the technology function, and the tools to answer those questions before someone else asks them.

INSIDE

The short version	2
The adoption question is closed	3
Comfort is rising faster than control	4
Discovery data is where it gets tested first	5
Twelve questions for any AI vendor	6
Tier the risk, measure what you can defend	7
Count your handoffs, plan the next 90 days	8
Sources and method	9

The short version

Five findings from the 2026 research, and what they add up to for the person accountable for the firm's technology.

1 Adoption is universal. Deployment is not.

94% of firms are using or exploring generative AI, yet only two AI products have reached full deployment to lawyers at half or more of firms.¹ Most AI in law firms is still sitting in pilots, each with its own data flow.

2 Budgets are up. Proof is missing.

Law firm technology spending grew 9.7% in 2025, likely the fastest real growth the industry has seen,² while only 18% of professionals know their organization tracks AI return.³

3 Comfort is rising faster than control.

Respondents naming accuracy as their top AI concern fell from 82% to 63% in a year.¹ Meanwhile, the court record keeps growing: 1,397 US decisions to date address AI-fabricated material in filings.⁴

4 Courts have started specifying AI controls for discovery data.

One 2026 protective order allowed AI on confidential material only if the provider is contractually barred from storing or training on inputs, limits onward disclosure and allows deletion. Another barred open AI tools from all discovery material.^{7,8}

5 The technology function holds the evidence.

Every question partners, clients and courts are now asking resolves to data only IT can produce: which tools, which data, which logs, which controls.

Three moves that follow

1. **Inventory** every AI surface that touches client or discovery data, including pilots and tools you did not choose.
2. **Hold every vendor to the same standard.** The twelve questions on page 6 are built from what courts and security leaders are now asking.
3. **Measure defensibility, not just hours saved.** Page 7 sets out six metrics a management committee can act on.

About this briefing. Statistics come from industry associations, market research, peer-reviewed study and court records, listed with dates and sample notes on page 9. No survey sponsored by an eDiscovery software vendor was used. Surveys cover different populations, so figures from different sources are never combined into one number. Throughout, the plain text states what a source found and the shaded "What it means for you" panels are our analysis of what it means for you.

The adoption question is closed. The accountability question is open.

For three years the question facing law firm technology leaders was whether to adopt AI. The 2026 data says that decision has been made, often by the lawyers. What remains open is harder: which tools are really deployed, what they cost, and whether anyone can show what they return.

94%

ILTA 2026, 508 firms

of law firms are using or exploring generative AI tools, up 14 points on last year. Among firms with 700 or more lawyers the figure is 100%, and among firms under 50 lawyers it is 85%.¹

WHAT IT MEANS FOR YOU

AI is already in your environment, sanctioned or not. The job has moved from selecting tools to governing the ones in use, including the ones IT did not choose.

2

ILTA 2026

AI products have reached full deployment to lawyers at half or more of surveyed firms. The long tail of tools sits in testing or pilot.¹

WHAT IT MEANS FOR YOU

Pilots are cheap to start and costly to leave running. Each carries its own contract, retention settings and data path, so a long pilot list is a long list of places client data may sit without an end-to-end review.

9.7%

Thomson Reuters and
Georgetown Law, US market

growth in average law firm technology spending in 2025, with knowledge management spending up 10.5%. The report describes it as likely the fastest real growth the industry has experienced.²

WHAT IT MEANS FOR YOU

Technology is now one of the fastest-growing lines in the firm's cost base, and partners notice fast-growing lines. Expect the next budget cycle to ask what this spend returned, in terms the management committee recognizes.

18%

Thomson Reuters Institute,
1,500+ professionals

of respondents said they knew their organization was tracking the return on its AI tools. Another 40% did not know whether return is measured at all.³

WHAT IT MEANS FOR YOU

Most firms cannot yet answer the question the spending growth invites. The technology leader who arrives with a measurement framework sets the terms of that conversation instead of reacting to it. Page 7 offers one.

90%

Thomson Reuters and
Georgetown Law

of legal dollars still move through hourly billing, while general counsel spending expectations have dropped to levels not seen since the pandemic.²

WHAT IT MEANS FOR YOU

Efficiency that is not measured gets absorbed by the hourly model and never registers as value. Clients under budget pressure will ask where the AI savings went, and the answer starts with data the technology function controls.

Comfort is rising faster than control

Familiarity is doing what familiarity does: as AI use spreads, reported worry about it is falling. The record outside the firm is moving the other way, and so is the security picture inside it.

82% to 63%

ILTA 2026

Share of respondents naming accuracy as the biggest area of concern about generative AI, year over year. Accuracy still tops the list.¹

WHAT IT MEANS FOR YOU

Falling concern is a behavioral signal, not a quality signal. Nothing in the data shows outputs became 19 points more reliable. The risk is that familiarity turns into checking less. Verification has to be built into the workflow, because it will not survive as a habit.

1,397

AI Hallucination Cases
Database, Sept 2026

US court and tribunal decisions addressing AI-hallucinated material in filings. Worldwide, 813 of the logged decisions involve a lawyer rather than a self-represented party.⁴

WHAT IT MEANS FOR YOU

This risk now lives in the court record, with names attached. When it happens inside a firm, the management committee's first questions are technical ones: which tool, which workflow, which control failed. Have the logs to answer them.

17% to 33%

Stanford and Yale, peer-reviewed

of queries produced hallucinated answers across leading AI legal research tools in a preregistered evaluation, lower than a general-purpose chatbot but far from zero.⁶

WHAT IT MEANS FOR YOU

Purpose-built is not the same as error-free. The study tested 2024 versions and tools have improved since, but the procurement lesson holds: ask vendors for validation you can reproduce on your own data, not for the word "grounded".

#2

ILTA 2026

Generative AI entered ILTA's tracked list of law firm security challenges for the first time this year, and landed second, behind only user behavior and ahead of malware and compliance.¹

WHAT IT MEANS FOR YOU

Your peers now rate AI as a top-tier security problem. That is the mandate to treat every AI data flow like any other system holding client data: inventory, access control, retention and audit.

13%

ILTA 2026

of respondents say they already use agentic AI, on the business side or in practice.¹

WHAT IT MEANS FOR YOU

Agents take actions rather than only producing text. Supervising a chain of automated steps needs a record of what the system did and why. If your audit trail cannot show that for today's tools, agents will expose the gap.

On policy. 78% of ILTA survey respondents report a formal generative AI policy.¹ That is a real achievement and also the floor. A policy states what people should do. Only systems and logs show what they did, and those are what a client audit or a court will ask for.

Discovery data is where AI governance gets tested first

Research and drafting touch the firm's own work. Discovery touches everyone else's data: client custodians, opposing productions, third parties who never agreed to anything. In 2026, rules on which AI may touch it became an increasingly routine part of protective orders.

13

AI Discovery Cases Database,
to 10 Sept 2026

US rulings logged in 2026 on AI in discovery: whether AI prompts and outputs are protected or must be produced, and how protective orders should govern AI tools.⁵

WHAT IT MEANS FOR YOU

Several turn on technical facts your team controls: where a tool runs, what it retains and whether data can be deleted.

Three rulings a technology leader should know

Morgan v. V2X, Inc.

D. Colo., March 2026

Amended the protective order so confidential information may enter an AI tool only if the provider is contractually barred from storing or training on inputs and from disclosing them to third parties except where essential to the service, and must allow deletion on request. Parties must keep written documentation of those terms. The court acknowledged this would, for now, bar most if not all mainstream low-cost AI for that material.^{8,10}

Jeffries v. Harcros Chemicals, Inc.

D. Kan., March 2026

Extended the protective order's restriction on open, publicly accessible AI tools to all discovery materials, confidential or not, citing the practical impossibility of retrieving data once a public model has taken it in. The order also requires notice before a party uses a particular AI tool, and use within a secure environment.^{7,10}

United States v. Heppner

S.D.N.Y., February 2026

Held that documents a criminal defendant generated with a consumer AI platform, on his own initiative and without direction from counsel, were protected by neither attorney-client privilege nor the work product doctrine. The court pointed to the platform's privacy policy, which permitted training on inputs and disclosure to third parties.⁹

What this means for the technology function

PROTECTIVE ORDERS ARE BECOMING TECHNICAL SPECIFICATIONS

No storing or training on inputs, limits on onward disclosure, deletion on demand, and written proof of all three. These are requirements your stack must meet on a timetable set by someone else's motion. Know today which of your tools pass.

THE APPROVED-TOOL LIST IS NOW LITIGATION INFRASTRUCTURE

One order requires notice before a particular AI tool is used; another compelled a party to name the AI tools it had used. The litigation team needs an accurate list of approved tools, and where each runs, on day one of discovery. That is an IT inventory, not a legal memo.

THE UNAPPROVED PATH IS THE WEAKEST LINK

Material created in consumer AI tools, outside a controlled workflow, may not be protected. Policy alone will not stop it. Make the approved, contained path the fastest one for the work people are trying to do.

Twelve questions for any AI vendor that touches client data

Built from the safeguards courts are now writing into protective orders and the risks security leaders ranked highest this year. Send the same list to every vendor, including the incumbents and the pilots, and keep the answers on file. A vague answer is itself an answer.

Ask	Why it matters
DATA USE	
1 Are our inputs, outputs or documents stored for, or used to train or improve, any model, yours or a third party's? Is that commitment in a contract we can produce?	The first safeguard in Morgan, which also required written documentation of it. A policy page can change; a contract term is enforceable.
2 Which sub-processors and model providers can see our data, and in which jurisdictions?	Onward disclosure is the second Morgan safeguard, and each sub-processor is a new place data lives.
3 Can the service run inside our cloud tenant, a specified region or on our own servers?	Some matters and clients will not allow data to leave a defined environment.
CONTAINMENT	
4 Can we delete a matter's data on demand, including derived data such as indexes, embeddings and logs, and receive written certification?	Deletion is the third Morgan safeguard. Derived data is where it usually fails.
5 How long are prompts and outputs retained by default, and can we set retention per matter?	Default retention quietly outlives the protective order.
6 Is data segregated by matter, so one court's order can be enforced without touching other matters?	Different matters carry different orders. Controls need matter-level granularity.
TRACEABILITY	
7 Can every AI output be traced to the specific source documents behind it?	Verification is only practical when the reviewer can open the source in one step.
8 Is every AI action logged (who, what, when, which model version) and exportable in a form a court or opposing counsel could read?	When something goes wrong, the log is the firm's defense. Agents make this essential.
9 How is accuracy validated on our data, and can we reproduce the validation ourselves?	Published benchmarks show purpose-built tools still err. Your data is the only test that counts.
CHANGE AND EXIT	
10 Will we be notified before the underlying model or its settings change?	A silent model change can invalidate a validation you relied on mid-matter.
11 Can AI features be switched off for a single matter if a court order requires it?	An order may restrict AI on one case while the rest of the firm keeps working.
12 What does a complete export of our data and work product cost, and how long does it take?	Exit terms decide whether you can act on the answers to the other eleven.

Tier the risk. Measure what you can defend.

Not every AI use on litigation data needs the same control. A simple three-tier model lets the technology function apply the heaviest controls where exposure is highest, and approve lower-risk work quickly.

Tier 1: Contained

Internal notes, summaries of public filings, drafting that uses no client data.

Controls: approved tools only; periodic spot checks.

Owner: practice group, under firm policy.

Tier 2: Validated

Early case assessment, clustering and prioritization, first-pass review, summaries of custodian documents.

Controls: closed tools that meet the Morgan safeguards; logged sampling and validation; continuous audit trail.

Owner: litigation support with IT.

Tier 3: Attested

Anything a court or opposing party relies on: productions, privilege logs, certified responsiveness calls, citations in filings.

Controls: documented human sign-off; source trace for every assertion; retained logs.

Owner: supervising attorney, with IT evidence.

A starting framework, not legal advice. Set tier boundaries with litigation leadership and general counsel, and revisit them when orders in your matters change.

Six metrics a management committee can act on

Hours saved is the metric everyone reports and nobody can defend. These six connect AI spend to the questions partners, clients and courts are actually asking.

Metric	What it measures	Question it answers
Time to defensible output	Elapsed time from request to an AI-assisted output that has been reviewed and traced to its sources.	Are we faster at work we can stand behind?
Source-trace rate	Share of AI outputs linked to specific source documents.	Can a reviewer verify this in one step?
Verification rework rate	Share of AI outputs materially changed at review.	Is the tool improving, and is anyone checking?
Data egress points per matter	Number of systems outside firm control that hold a copy of the matter's data.	How many places must a deletion order reach?
Sanctioned-tool coverage	Share of AI activity on client data that runs through approved tools, estimated from logs and network data.	Is policy matching behavior?
AI terms agreed at outset	Share of active matters where AI use is addressed in the protective order or ESI protocol before discovery begins.	Are we setting terms, or reacting to motions?

Count your handoffs. Plan the next 90 days.

Pick one recent matter and fill in the grid. Every copy of the data is a place a deletion order must reach. Every change of system is a gap in the audit trail that someone may one day ask you to explain.

Stage	System used	Where the data sits	Copy created?	AI used?	Audit trail unbroken?
Legal hold					
Collection					
Processing					
Early case assessment					
Review					
Production					

Inventory

Days 1 to 30

- List every AI tool that touches client or discovery data, pilots included.
- Map where each one runs and what it retains.
- Complete the handoff grid for two live matters.

Standardize

Days 31 to 60

- Send the twelve questions to every vendor.
- Agree tier boundaries with litigation leadership.
- Draft model AI language for protective orders and ESI protocols.

Measure

Days 61 to 90

- Baseline the six metrics.
- Retire or contain tools that fail the questions.
- Report to the management committee in their terms.

HOW VENIO APPROACHES THIS

We wrote the twelve questions to be vendor-neutral, and we expect to be asked them. Venio runs legal hold, processing, early case assessment, review and production on one platform, so a matter's data stays in one system with one audit trail instead of being copied between tools. Deployment is your choice of cloud, on-premises or hybrid, depending on where the data is allowed to live.¹¹

~40%

less review volume before first eyes-on at an AmLaw 50 firm

35%

reduction in operator hours and infrastructure for the same firm

72%

reduction of a 120+ TB dataset for a global bank

Put your discovery stack through the twelve questions with us. A 30-minute working session on your own matter profile, not a slide deck.

[Book a working session →](#)

Sources and method

- 1 **International Legal Technology Association (ILTA), 2026 Technology Survey**, released September 2026. 508 firms representing more than 139,000 attorneys across about a dozen countries; not a US-only sample. Figures as published in D. Austin, eDiscovery Today (14 Sept 2026) and R. Ambrogio, LawSites (14 Sept 2026). iltanet.org/resources/publications/surveys/ts26
- 2 **Thomson Reuters Institute and Georgetown Law Center on Ethics and the Legal Profession, 2026 Report on the State of the US Legal Market**, 7 January 2026. US law firm financial and buyer-sentiment data. [thomsonreuters.com press release, 7 Jan 2026](https://thomsonreuters.com/press-release/7-Jan-2026)
- 3 **Thomson Reuters Institute, 2026 AI in Professional Services Report**. 1,500+ professionals across legal, tax, accounting, risk, fraud and government in 27 countries; global sample. thomsonreuters.com/en/institute/reports/2026-ai-in-professional-services-report
- 4 **D. Charlotin, AI Hallucination Cases Database**, last updated 19 September 2026; 2,044 decisions worldwide at that date. Counts change daily. Licensed CC BY 4.0. damiencharlotin.com/hallucinations
- 5 **D. Charlotin, AI Discovery Cases Database**, last updated 10 September 2026. Licensed CC BY 4.0. damiencharlotin.com/ai-discovery
- 6 **V. Magesh, F. Surani, M. Dahl, M. Suzgun, C. D. Manning and D. E. Ho, Hallucination-Free? Assessing the Reliability of Leading AI Legal Research Tools**, 22 Journal of Empirical Legal Studies 216 (2025). Tested 2024 product versions. doi.org/10.1111/jels.12413
- 7 **Jeffries v. Harcros Chemicals, Inc.**, No. 25-2352-KHV-ADM, 2026 WL 820218 (D. Kan. Mar. 25, 2026), memorandum and order granting motion to amend protective order.
- 8 **Morgan v. V2X, Inc.**, No. 1:25-cv-01991-SKC-MDB, ECF No. 65, 2026 WL 864223 (D. Colo. Mar. 30, 2026), order granting in part motion to amend stipulated protective order.
- 9 **United States v. Heppner**, No. 25 Cr. 503 (JSR), ECF No. 27, 2026 WL 436479 (S.D.N.Y. Feb. 17, 2026), memorandum.
- 10 **Sidley Austin LLP, Generative AI in Discovery: Protective Orders as an Emerging Point of Dispute**, 6 April 2026. sidley.com/en/insights/newsupdates/2026/04/generative-ai-in-discovery-protective-orders-as-an-emerging-point-of-dispute
- 11 **Venio Systems, platform pages and published case studies** (AmLaw 50 firm; global bank). These are Venio's own reported results. veniosystems.com

Method notes

Independence. No survey sponsored or co-produced by an eDiscovery software or services vendor was used. Surveys from legal information publishers are identified by name so readers can weigh them.

Populations. Each figure describes its own survey population, noted beside it. Figures from different surveys are not combined or directly compared.

Court decisions. Summaries describe what each order required or held and are not a statement of law in any jurisdiction. Read the orders before relying on them.

Frameworks. The twelve questions, the risk tiers, the six metrics and the 90-day plan are Venio's recommendations derived from the sources above. They are not survey findings.

Not legal advice. This briefing is for information. Consult your general counsel and litigation leadership on matter-specific obligations.