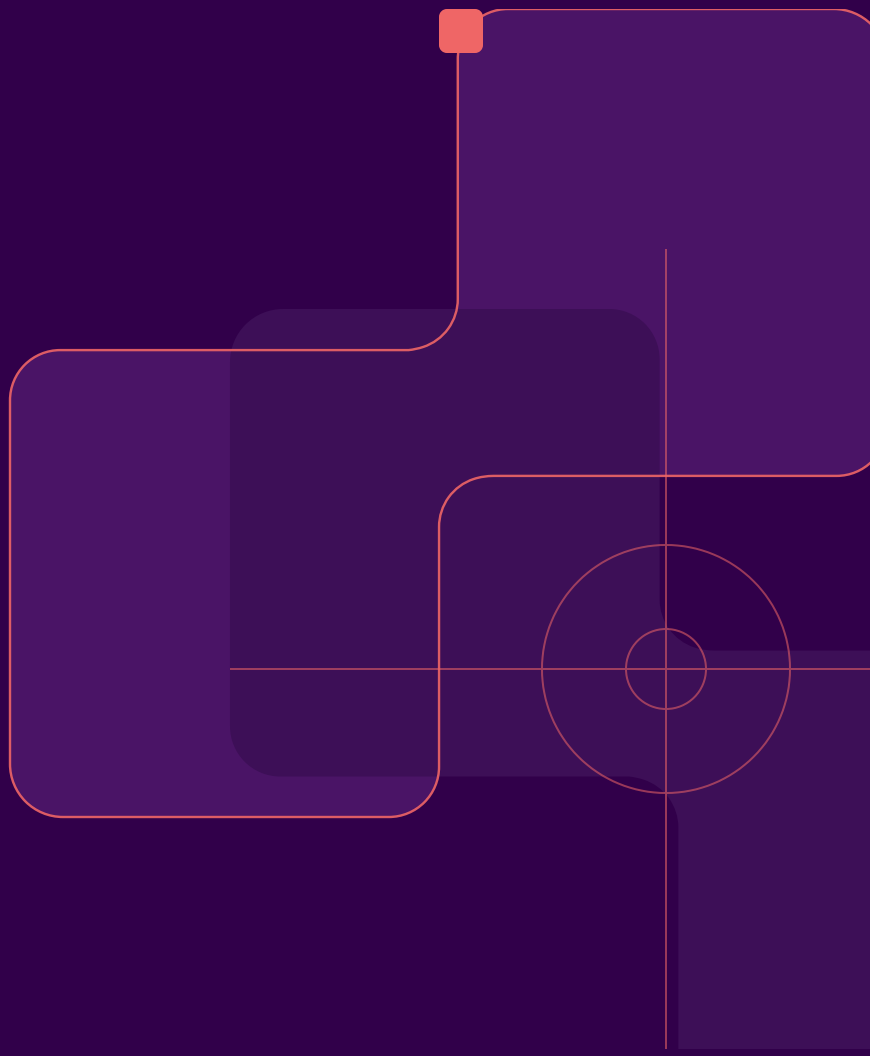
ELUCIENT



WHITEPAPER

AI Governance in the Firms You Hire

What to expect from the agencies and consultants who handle your data

Published
November 2026

Author
John Marcinuk, Founder and CEO, Elucient

Web
elucient.com

EXECUTIVE SUMMARY

Your agencies and consultants are already using AI on your work

The firms you hire now handle your data, your strategy and your brand with AI, so their governance has become part of yours.

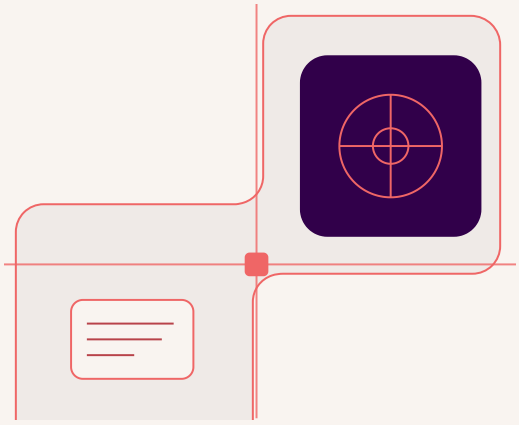
In Thomson Reuters' 2026 survey of more than 1,500 professionals, 40% said their organization uses generative AI, up from 22% a year earlier. About two-thirds of corporate respondents said their outside firms should be using AI, and fewer than 20% require anything about it through guidelines or RFPs.[1] In corporate legal departments, 59% of in-house teams don't know whether their law firms use generative AI on their matters.[2]

Control inside firms is uneven. In a survey of 2,000 workers published in January 2026, 49% said they had adopted AI tools without approval,[3] and IBM found that 92% of organizations with an AI-related breach lacked proper AI access controls.[4] Only 15% of organizations track how all of their relevant third parties use or govern AI.[5]

This paper shows how a firm's AI use reaches your business, what good governance looks like inside a firm and what to write into your contracts. It closes with a 90-day plan and nine questions to ask every firm you hire.

WHAT'S INSIDE

Your firms are already using AI	3
How a firm's AI use reaches you	4
What good governance looks like	7
Where your data should run	8
Write it into the contract	10
A 90-day plan	11
Nine questions to ask	12



DEMAND

2 in 3

corporate clients say their outside firms should be using AI[1]

REQUIREMENTS

Under 20%

of corporate clients require anything about AI through guidelines or RFPs[1]

OVERSIGHT

15%

of organizations track how all relevant third parties use or govern AI[5]

THE ADOPTION SURGE

Your firms are already using AI on your work

Research, analysis, proposal writing and first drafts now run through AI at a growing share of professional services firms. In Thomson Reuters' 2026 survey of professionals in legal, tax, accounting, risk and government work, 40% said their organization uses generative AI, up from 22% a year earlier. Among those organizations' users, 82% use it at least weekly, and more than half use publicly available tools such as ChatGPT and Claude for work.[1]

Guidance inside those firms is inconsistent. Forty percent of firm professionals say they've been told both to use and not to use AI on client work.[1] On the client side, 59% of in-house legal teams don't know whether their law firms use generative AI on their matters.[2] Fewer than 20% of corporate clients require anything about AI through guidelines or RFPs, so each firm decides for itself how AI touches the work it does for you.[1]

Agents are the next step. Fifteen percent of organizations use agentic AI, which can take actions across systems instead of only answering prompts, and another 53% are planning or considering it.[1] Microsoft counted 15 times as many active agents across Microsoft 365 in March 2026 as it had a year earlier.[6] An agency's agent with access to your shared drive or CRM is a larger exposure than a person drafting copy, and it calls for its own rules.

MIXED SIGNALS

40%

of firm professionals have been told both to use and not to use AI on client work[1]

59%

of in-house legal teams don't know whether their law firms use generative AI on their matters[2]

15x

growth in active agents across Microsoft 365, March 2025 to March 2026[6]

Four ways a firm's AI use reaches your business

Exposure through an outside firm usually starts with ordinary work. An account lead pastes your quarterly numbers into a chatbot to build a forecast faster, or an analyst uploads your customer interviews to summarize them overnight. Each step makes sense to the person taking it. Without governance at the firm, your information ends up in tools neither of you has reviewed.



Your data in their staff's tools

In a BlackFog survey reported in January 2026, 49% of workers said they'd adopted AI tools without approval, and 58% of those using unapproved tools relied on free versions. A third had shared enterprise research or datasets with them, and 23% had entered company financial information.[3] When those workers are on your account, the research and financials are yours.



Their vendors' terms

Free and consumer AI accounts carry different terms than enterprise agreements, including how long the vendor keeps data, whether it trains on it and when it discloses it. In February 2026, a federal judge in New York ruled in *United States v. Heppner* that documents created with an AI chatbot weren't privileged, citing the provider's policy allowing training and disclosure to authorities.[7]



Their output under your name

AI drafts come back fluent and confident, including when they're wrong. In Microsoft's 2026 research, 86% of AI users said they treat its output as a starting point, and quality control of AI output topped their list of the human skills that matter most.[6] A firm without a review step hands that starting point to you as finished work.



Their agents in your systems

An agent that can read a shared drive, an inbox or a CRM can be redirected by instructions hidden in a document or web page. OWASP ranks this attack, agent goal hijack, first in its Top 10 for Agentic Applications.[8] IBM put the average loss from a prompt injection attack at $5.89 million.[4] Any agent a firm connects to your accounts brings that risk into your environment.

WHAT CAN GO WRONG

Missing controls show up in breaches and in deliverables

92%

of organizations reporting an AI-related breach lacked proper AI access controls[4]

1 in 5

breached organizations reported an attack on their AI models or applications[4]

40%

of breached organizations use access controls on their AI models and data[4]

IBM, Cost of a Data Breach Report 2026, based on 602 organizations affected by a data breach.

CASE STUDY

What two government clients found in reports they paid for

AUSTRALIA, 2025

Australia's Department of Employment and Workplace Relations paid AU$440,000 for a review by Deloitte Australia. A University of Sydney researcher found a fabricated quote attributed to a federal court judge and references to academic work that doesn't exist. Deloitte agreed to repay the contract's final installment, and the revised report disclosed that it had been written using Azure OpenAI.[9]

CANADA, 2025 TO 2026

Newfoundland and Labrador paid about CA$1.6 million for a health workforce plan from Deloitte Canada. Reporting by The Independent and Radio-Canada found citations to sources that appear not to exist. Deloitte said AI was used "to support a small number of research citations" and issued corrections, and the province is reviewing options for compensation.[10]



Newfoundland and Labrador now requires vendors to state whether and how they plan to use AI on provincial contracts, and it can approve, deny, modify or audit that use.[10]

How the clients found out

In both cases, outsiders found the errors after the reports were published: a researcher in Australia and reporters in Canada. The province's response moved the question to the start of the work, where AI use is disclosed and approved before anything is delivered.

What it means for a private company

A company that receives AI-assisted work with unchecked sources may hear about it from a customer, its board or a regulator, after the material is already in use.

WHY THIS MATTERS NOW

Rules for AI use are arriving from regulators, buyers and the courts



Regulators

The EU AI Act's transparency rules, which cover telling people when they're dealing with AI and marking AI-generated content, apply from August 2, 2026. The Digital Omnibus on AI, in force since July 27, 2026, keeps a duty for companies that use AI to support staff AI literacy and moves most high-risk obligations to December 2027.[11] In May 2026, Colorado replaced its 2024 AI Act with a law focused on transparency when automated systems shape decisions about jobs, housing, health care and insurance.[12]



Buyers and procurement

Some buyers are already writing AI terms into procurement. Newfoundland and Labrador now requires bidders to disclose planned AI use and reserves the right to audit it.[10] Across organizations, only 15% track how all of their relevant third parties use or govern AI,[5] and fewer than 20% of corporate clients require anything about AI through guidelines or RFPs.[1]



Courts and professional standards

ABA Formal Opinion 512 says lawyers need a client's informed consent before entering information about the client's matter into AI tools that can learn from it.[13] The *Heppner* ruling showed that a provider's terms can decide whether information shared with AI stays confidential.[7] Both give your legal team a clear reference point for what to require of any firm that handles your confidential information.

EU AI ACT DATES AFTER THE DIGITAL OMNIBUS[11]

Feb 2, 2025
Prohibited practices and AI literacy duty apply

Aug 2, 2026
Transparency obligations apply

Dec 2, 2026
Content-marking deadline for systems already on the market

Dec 2, 2027
Stand-alone high-risk systems, such as hiring tools

Aug 2, 2028
High-risk AI built into regulated products

WHAT GOOD GOVERNANCE LOOKS LIKE

Five layers of AI governance inside a firm

A well-run firm can describe its AI governance in five layers. The structure draws on the NIST AI Risk Management Framework and ISO/IEC 42001,[14,15] adapted to firms of 20 to 200 people. Each layer answers a question you can put to any agency or consultant, and each one leaves evidence you can ask to see.

Ownership WHO IS ACCOUNTABLE	A named partner or the COO owns AI policy, supported by a small group from delivery, IT and legal that meets on a set schedule and rules on new tools and edge cases.	ASK TO SEE	The person accountable for AI on your account
Data HOW YOUR INFORMATION IS CLASSIFIED	Every piece of client information sits in a tier, from public material to restricted data, and each tier has a plain definition staff can apply on a deadline.	ASK TO SEE	The firm's data tiers and where your information falls
Tools WHERE YOUR WORK CAN RUN	Each tier maps to approved deployments: enterprise tools under contracts that bar training and limit retention, a private cloud account or models on the firm's own hardware.	ASK TO SEE	The approved tool list and the vendor terms behind it
Workflow HOW AI-ASSISTED WORK IS REVIEWED	Review scales with risk. Every client deliverable has a named reviewer who checks figures, quotes and citations against their sources before anything is sent.	ASK TO SEE	The review step for deliverables on your account
Proof WHAT THE FIRM CAN SHOW YOU	The firm keeps a written AI statement for clients, a log of which tools touched which engagement and answers ready for security reviews.	ASK TO SEE	The AI statement and the usage log for your work

WHERE YOUR DATA SHOULD RUN

Decide where each kind of information may be processed

You know better than any firm which of your information is sensitive, so classification should start with you. Four tiers cover what a company typically hands to an agency or consultant, and each one maps to the places a firm may process it. Once the tiers are in your contract, they set the floor for every firm you hire, whatever its own defaults allow.

Open Published content, public research, live campaigns and press materials	CAN RUN ON Any approved tool
Internal Brand guidelines, approved messaging, campaign briefs and other working material that isn't sensitive	CAN RUN ON Enterprise SaaS
Confidential Customer data, research findings, sales performance, and unreleased product and pricing plans	CAN RUN ON Private cloud or local or an enterprise tool you've approved in writing
Restricted Unreleased financials, deal and M&A material, regulated customer data and anything under NDA	CAN RUN ON Local only

Enterprise SaaS

A vendor's service under a business contract that bars training on the firm's data and limits how long the vendor keeps it.

Private cloud

Models running inside the firm's own cloud account, such as Azure, AWS Bedrock or Google Vertex AI, so your data stays in an environment it controls.

Local

Open-weight models running on servers or laptops the firm owns, with nothing sent to an outside service.

LOCAL MODELS

When to ask a firm to keep your work on its own hardware

Sensitive work used to force a choice between a frontier model in someone else's cloud and no AI at all. Epoch AI found in May 2026 that the most capable open-weight models, which anyone can download and run, trail the leading closed models by about four months on average.[16] A firm can run models at that level on its own servers, and smaller ones run on a standard work laptop. A firm with that capability can still use AI on your Restricted information.

Small, task-specific models handle redaction, file classification, transcript summaries and contract extraction on ordinary hardware. Gartner predicts that by 2027, organizations will use them at least three times more than general-purpose LLMs.[17]

CAPABILITY GAP

4 months

average lag between the best open-weight models and the leading closed models[16]

A firm can also pair a local model with a cloud model

ON THE FIRM'S HARDWARE		IN THE CLOUD	ON THE FIRM'S HARDWARE
Mask	**Check**	**Reason**	**Restore and review**
A local small model swaps your names, figures and identifiers for placeholders.	A reviewer or automated scan confirms nothing identifying remains.	A frontier model works on the masked text and returns its draft.	The local model puts your details back, and a named reviewer checks the result.

Controls to ask about

Local models need their own controls. Ask a firm that runs them for:

- Its approved list of models and versions
- Where it downloads models from
- Access controls and usage logs
- Its schedule for updates and retirement

MODEL SUPPLY CHAIN

352,000

unsafe or suspicious issues flagged across 51,700 models when Protect AI and Hugging Face scanned more than 4 million[18]

Write it into the contract

A contract turns a firm's governance into something you can enforce. Fewer than 20% of corporate clients set AI requirements through guidelines or RFPs,[1] so today a firm's own policy usually decides how AI touches your work. Newfoundland and Labrador now asks every bidder to state how it will use AI and keeps the right to audit that use.[10] Marketing and operations buyers can take the same approach through documents they already use.



In the RFP

- Ask for the firm's AI policy and approved tools
- Ask which parts of the proposed work will involve AI
- Ask how AI-assisted work is reviewed
- Ask who is accountable for AI on your account



In the SOW

- Name the approved tools for the engagement
- Set the data tier for each kind of material you share
- Require disclosure of AI use in each deliverable
- Require source checks for every figure, quote and citation



In the master agreement

- Bar training on your data by the firm and its vendors
- Set breach notice terms for AI tools
- Require your approval before any agent connects to your systems
- Keep a right to audit AI use on your work

3 in 4 professionals say firms should be the ones to start the AI conversation with clients.[1] Expect the firms you hire to raise it first, with their answers ready.

A REFERENCE PLAN

A 90-day plan for your outside firms

Days 0 to 30

Map who touches what

- List every agency, consultancy and contractor with access to your data
- Note which data tiers each one handles
- Pull any AI terms in current contracts
- Find which firms connect tools or agents to your systems

Days 31 to 60

Ask and require

- Send the nine questions to each firm
- Add AI terms to new RFPs and SOWs
- Set the data tiers for each engagement
- Approve or decline each request for agent access

Days 61 to 90

Review and renew

- Review the answers and follow up on gaps
- Add AI terms to master agreements at renewal
- Collect each firm's written AI statement
- Schedule an annual review of every firm's answers

Decide how much an outside firm's agent may do in your systems

Gartner sorts AI agents into four levels of autonomy, each needing its own controls, and predicts that by 2027, 40% of enterprises will demote or decommission autonomous agents because governance gaps surface only after incidents in production.[19]

Level 1

Observe

Reads defined sources, and only the requesting user sees its output.

CONTROLS

Limited data access, authentication, usage logs

Level 2

Advise

Drafts and recommends while people take every action.

CONTROLS

Accuracy testing, quality checks, training on how far to rely on it

Level 3

Act with approval

Writes, sends or changes things after a person approves each action.

CONTROLS

Approval workflow with audit trail, security testing, incident response

Level 4

Act autonomously

Works inside guardrails while people review exceptions and logs.

CONTROLS

Continuous monitoring, fast rollback, automatic shutoff at set thresholds

> Limit an outside firm's agents to Observe or Advise in your systems by default, and approve anything higher for a specific, named use.

READINESS CHECK

Nine questions to ask every firm you hire

Ask these before your next SOW or at your next renewal. A well-governed firm can answer each one in a sentence or two and show you the evidence behind it.

Which AI tools touch our data, and under what terms?

A GOOD ANSWER

A named list of enterprise tools, with contract terms that bar training and limit retention.

Where is our data processed, and does it ever leave your environment?

A GOOD ANSWER

Specific locations for each data tier, including any work that runs on the firm's own hardware.

Is any of our information used to train a model?

A GOOD ANSWER

A clear no, backed by vendor terms the firm can show you.

Who reviews AI-assisted work before it reaches us?

A GOOD ANSWER

A named reviewer for each deliverable and a defined review step.

How do you verify facts, figures and citations?

A GOOD ANSWER

Every source checked against the original before delivery, by someone accountable for it.

Can you tell us which parts of a deliverable involved AI?

A GOOD ANSWER

Yes, noted on each deliverable or available on request.

How do you handle our contract's limits on third-party AI?

A GOOD ANSWER

A specific plan, such as moving that work to local models.

What happens if an AI tool you use has a breach?

A GOOD ANSWER

A notice timeline, a named contact and a record of which of your data the tool held.

Who at your firm is accountable for AI use on our account?

A GOOD ANSWER

One person, by name.

HOW ELUCIENT APPLIES THIS

The same five layers, on our own work

Elucient applies these five layers to its own work and to client engagements, including keeping sensitive client material out of cloud AI.

We also help companies put these questions, data tiers and contract terms in place with the agencies and consultants they already work with, so the firms they hire can use AI on their behalf with clear rules.

Elucient is a strategy and advisory consultancy. We help businesses bring new products to market, improve marketing and business performance, and modernize the people, processes and platforms that run them.

ELUCIENT

Let's talk.

John Marcinuk
Founder and CEO
john@elucient.com
elucient.com

SOURCES

Notes and sources

1 Thomson Reuters Institute, "2026 AI in Professional Services Report," February 2026. Online survey of 1,514 professionals in 27 countries, fielded October and November 2025.
thomsonreuters.com/en/reports/2026-ai-in-professional-services-report

2 Association of Corporate Counsel, release on the ACC and Everlaw generative AI report, October 14, 2025. Survey of 657 in-house professionals in 30 countries.
acc.com/about/newsroom/news/acc-genai-report-corporate-law-departments-ai-use-everlaw

3 Taryn Plumb, "Roughly half of employees are using unsanctioned AI tools, and enterprise leaders are major culprits," CIO, January 29, 2026. Reports a BlackFog survey of 2,000 workers at companies with more than 500 employees.
cio.com/article/4124760

4 IBM, "AI-powered adversaries and the enterprise risk challenge," IBM Think, July 29, 2026, on the Cost of a Data Breach Report 2026 conducted with the Ponemon Institute.
ibm.com/think/x-force/2026-cost-of-a-data-breach-ai-adversaries-enterprise-risk

5 Compliance Week, "The State of AI Governance in Third-Party Risk," November 4, 2026, citing interim findings from the ECI State of Third-Party Risk and Integrity Survey with Compliance Week and Diligent.
complianceweek.com/third-party-risk/the-state-of-ai-governance-in-third-party-risk

6 Microsoft WorkLab, "2026 Work Trend Index Annual Report: Agents, human agency, and the opportunity for every organization," May 5, 2026. Survey of 20,000 knowledge workers in 10 markets, plus Microsoft 365 telemetry.
microsoft.com/en-us/worklab/work-trend-index

7 Justin Curl and Mihir Kshirsagar, "AI and Privilege After United States v. Heppner," Lawfare, March 30, 2026, on the February 17, 2026 opinion of Judge Jed Rakoff, U.S. District Court for the Southern District of New York.
lawfaremedia.org/article/ai-and-privilege-after-united-states-v.-heppner

8 OWASP GenAI Security Project, "OWASP Top 10 for Agentic Applications," December 9, 2025.
genai.owasp.org

9 Associated Press, "Deloitte to partially refund Australian government for report with apparent AI-generated errors," October 7, 2025.
Via Spectrum News NY1: ny1.com/nyc/all-boroughs/ap-top-news/2025/10/07

10 Sarah Smellie, The Canadian Press, "Newfoundland and Labrador overhauls procurement after report errors thought to be AI," March 11, 2026.
Via CP24: cp24.com/news/canada/2026/03/11

11 Regulation (EU) 2026/1744 of the European Parliament and of the Council (Digital Omnibus on AI), amending Regulation (EU) 2024/1689, in force July 27, 2026.
eur-lex.europa.eu/eli/reg/2026/1744/oj/eng

12 Office of Governor Jared Polis, "Governor Polis Signs Bills into Law Making Colorado An Even Better Place to Do Business," May 14, 2026, on SB26-189, Automated Decision-Making Technology.
governorsoffice.colorado.gov · leg.colorado.gov/bills/sb26-189

13 American Bar Association Standing Committee on Ethics and Professional Responsibility, Formal Opinion 512, "Generative Artificial Intelligence Tools," July 29, 2024.
americanbar.org

14 National Institute of Standards and Technology, "Artificial Intelligence Risk Management Framework (AI RMF 1.0)," NIST AI 100-1.
nist.gov/itl/ai-risk-management-framework

15 International Organization for Standardization, ISO/IEC 42001:2023, "Information technology, Artificial intelligence, Management system."
iso.org/standard/81230.html

16 Jack Edwards and Luke Emberson, "Open models lag state-of-the-art closed models by 4 months," Epoch AI Data Insight, May 29, 2026.
epoch.ai/data-insights/open-closed-eci-gap

17 Gartner, "Gartner Predicts by 2027, Organizations Will Use Small, Task-Specific AI Models Three Times More Than General-Purpose Large Language Models," press release, April 9, 2025.
gartner.com/en/newsroom

18 Ana Maria Constantin, "The AI industry's model and agent skill repositories are full of malware," The Next Web, May 8, 2026.
thenextweb.com/news/hugging-face-clawhub-malware-ai-supply-chain

19 Gartner, "Gartner Says Applying Uniform Governance Across AI Agents Will Lead to Enterprise AI Agent Failure," press release, May 26, 2026.
gartner.com/en/newsroom