BUILT BY A TEAM THAT HAS SECURED MAJOR LEGAL INSTITUTIONS WORLDWIDE FOR 25 YEARS.

ARGENT FOR LAW FIRMS

AI AGENTS NOW ACT LIKE STAFF. PUT A CHECKPOINT IN FRONT OF EVERY ONE.

Argent Strategic Intelligence puts a checkpoint between your AI agents and the outside AI services they use. It checks every request before it leaves, and it runs entirely inside your own firm. The ABA says a self-learning AI tool needs the client's informed consent first,[1] and a professional-services breach averages $5.08 million.[2]

One pane of glass for your entire agent ecosystem: every agent, model, tool, owner and dollar in one view. **No rip and replace.** Argent works with what you already run: your identity provider (Okta, Microsoft Entra, LDAP), your own secrets vault (HashiCorp Vault, AWS Secrets Manager, CyberArk Conjur) and the AI tools your staff already use.



SEE.
Every agent, its owner, where it sends data.



PROVE.
A signed record of every decision.



STOP.
Unapproved models, tools and leaks are refused.

THE FIVE THINGS LAW FIRMS ASK FOR FIRST

1	**Budget controls** A document-review agent runs up a large bill overnight.	Set spend caps by team, person and agent. An agent that hits its cap is stopped; its owner asks for more, and nothing is granted until a manager approves, in Argent or through your ServiceNow or Jira tickets. Project cost tracking shows what each project or practice group spends on AI.
2	**Data protection** An associate pastes a client's settlement file into an AI.	Argent reads each request first and blanks out Social Security, card and bank account numbers before anything leaves the firm, plus your own matter numbers once you add a pattern. The record notes what kind of data, never the data itself.
3	**Auditing for auditors** An auditor asks which agent sent what, and when.	Every AI decision goes into a signed, tamper-evident record that cannot be quietly changed. A client, an auditor or your general counsel can check a package of it on their own computer and see whether anything was altered.
4	**Model control** Lawyers start using an AI model nobody approved.	Only models you approve are allowed; everything else is refused. Group them into model classes, such as approved for client data and general use. A model can be swapped for an approved one without staff changing anything.
5	**Air-gapped** Client files must never leave the firm.	Argent runs entirely inside your own firm, even on a fully air-gapped network. The install bundle is checked at every step as it crosses into your closed environment.

WHAT COURTS AND STATES ALREADY EXPECT

Federal Rule of Evidence 502(b). In a federal case, an accidental disclosure does not waive privilege if the holder took reasonable steps to prevent it and promptly fixed it.[3]

New York GBL 899-aa. Affected residents must be told within **30** days of discovery; a knowing or reckless failure can cost up to **$250,000**.[4]

HOW A LAW FIRM STARTS

1. Watch mode. A real-time audit and view of what your AI is doing today. Nothing is blocked.

2. Enforce mode. Switch on blocking one area at a time: client data, budgets, models, tools.

3. Inside your walls. Runs on your own hardware or in your own cloud VPC. Nothing leaves the perimeter without your explicit approval.

BOOK A DEMO. WE WILL SHOW YOU ARGENT WORKING ON YOUR FIRM'S OWN USE CASES.

NYEDIS.AI/DEMO

Sources. 1 ABA Formal Opinion 512, Generative Artificial Intelligence Tools, July 29, 2024. 2 IBM / Ponemon, Cost of a Data Breach Report 2026, July 2026 (Services: legal, accounting and consulting firms). 3 Federal Rule of Evidence 502(b), read 2026-11-19. 4 N.Y. Gen. Bus. Law 899-aa(2), (6)(a), NY State Senate, read 2026-11-19.