BUILT BY A TEAM THAT HAS SECURED MAJOR CONSULTING AND ADVISORY INSTITUTIONS WORLDWIDE FOR 25 YEARS.

ARGENT FOR CONSULTING 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 outside AI services and checks every request before it leaves, inside your firm. The professional, business and consumer services sector, consultancies included, was the third most attacked industry in 2025,[1] and one in four malicious breaches is AI-enabled.[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 CONSULTING FIRMS ASK FOR FIRST

1	**Budget controls** A proposal team's agent runs up a large bill before a pitch.	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. Each session's AI spend is booked to the project it declared.
2	**Data protection** A consultant pastes a client's unannounced deal into an AI.	Argent reads each request first and blanks out your client and project code names (patterns you write), plus Social Security, card and bank account numbers, before anything leaves the firm. The record notes the kind of data, never the data.
3	**Auditing for auditors** A client asks which agent sent what about its deal, and when.	Every AI decision goes into a signed, tamper-evident record that cannot be quietly changed. A client, an auditor or internal audit can check a package of it on their own computer and see whether anything was altered.
4	**Model control** A government-practice team uses a 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 consultants changing anything.
5	**Air-gapped** Client material must never leave the building.	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 REGULATORS ALREADY EXPECT

DFARS 252.204-7012. A contractor holding U.S. defense information must report a cyber incident affecting it within **72** hours of discovery; the clause flows to subcontractors.[3]

EU GDPR, Article 33. A firm handling a client's personal data tells the client "without undue delay"; the client then has **72** hours, where feasible, to tell the regulator.[4]

HOW A 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, then 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 IBM X-Force, Threat Intelligence Index 2026, industry trends (2025 incidents). 2 IBM, Cost of a Data Breach Report 2026, July 2026. 3 DFARS 252.204-7012(a) and (m), Cornell LII copy, read 2026-11-19. 4 Regulation (EU) 2016/679, Article 33, EU Publications Office, read 2026-11-19.