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

ARGENT FOR CONSUMER BRANDS

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 company. In IBM's 2026 study, 43% of security incidents involved AI tools staff used without approval.[1]

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 BRANDS ASK FOR FIRST

1	**Budget controls** A launch team's 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 in ServiceNow or Jira. Project cost tracking shows what each project spends on AI.
2	**Data protection** A chemist pastes a new formula sheet into an AI.	Argent reads each request first and blanks out customer emails, phone and card numbers before anything leaves. Add your own formula and supplier code formats, or register a formula sheet so a pasted copy is caught. The record notes what kind of data, never the data.
3	**Auditing for auditors** Counsel asks when an injury complaint reached the agent.	Every AI decision goes into a signed, tamper-evident record, stamped with the time Argent saw it. Counsel or an auditor can check a package of it on their own computer and see whether anything was altered.
4	**Model control** A care agent sends complaints to an unapproved model.	Only models you approve are allowed; everything else is refused. Group them into model classes, such as approved for customer data and general use. A model can be swapped for an approved one without staff changing anything.
5	**Air-gapped** A new formula must never leave the building.	Argent runs entirely in-house, 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

CPSC, 16 CFR 1115.14. Once information supports a defect that could create a substantial risk of injury, firms should report "immediately, that is, within **24** hours". Complaints can count.[2]

FTC reviews rule, 16 CFR 465. Fake consumer reviews are illegal, and AI-generated ones are covered: up to **$53,088** per violation.[3]

HOW A BRAND 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: data such as formulas, 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 COMPANY'S OWN USE CASES.

NYEDIS.AI/DEMO

Sources. 1 IBM / Ponemon, Cost of a Data Breach Report 2026, July 2026 (shadow AI; Consumer sector, Figure 4). 2 16 CFR 1115.14(c), (e), CPSC, eCFR, read 2026-11-13. 3 16 CFR 465.2; FTC final rule, 89 FR 68034, Aug 2024; FTC release, Dec 22 2025, read 2026-11-13.