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

ARGENT FOR GOVERNMENT

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 each request before it leaves and runs inside your agency. Generative AI use cases grew about nine-fold in a year at 11 agencies GAO reviewed,[1] and over 40 states introduced about 250 measures on government AI use.[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 AGENCIES ASK FOR FIRST

1	**Budget controls** A program office'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 through your ServiceNow or Jira tickets. Project cost tracking shows what each program spends on AI.
2	**Data protection** A caseworker pastes a family's case notes into an AI.	Argent reads each request first and blanks out Social Security, card and bank account numbers before anything leaves the agency (case numbers, CUI and FTI markings too, if switched on). The record notes the kind of data, never the data.
3	**Auditing for auditors** An inspector general asks which agent sent what.	Every AI decision goes into a signed, tamper-evident record that cannot be quietly changed. An auditor or inspector general can check a package of it on their own computer and see whether anything was altered.
4	**Model control** Staff start using an AI model nobody approved.	Set the list of models you approve; everything else is refused. Group them into model classes, such as approved for case data and general use. A model can be swapped for an approved one without staff changing anything.
5	**Air-gapped** Case data must never leave the agency's network.	Argent runs entirely inside your agency, even on a fully air-gapped network. The install bundle is checked at every step as it crosses into your closed environment.

WHAT OVERSIGHT ALREADY EXPECTS

FISMA, 44 U.S.C. 3554. An agency must notify Congress of a major incident "not later than **7** days" after it has a reasonable basis to conclude one occurred.[3]

OMB memo M-25-21. Most agencies (not Defense or the Intelligence Community) must publish a yearly AI use-case list and shut down high-impact AI lacking safeguards.[4]

HOW AN AGENCY 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: case 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 AGENCY'S OWN USE CASES.

NYEDIS.AI/DEMO

Sources. 1 GAO-25-107653, Generative AI Use and Management at Federal Agencies, Jul 2025. 2 NCSL, New Trends Emerge as States Refine AI Legislation, Jan 2026. 3 44 U.S.C. 3554(b)(7)(C)(iii)(III), via Cornell LII, read 2026-11-13. 4 OMB Memorandum M-25-21, Apr 3, 2025.