BUILT BY A TEAM THAT HAS SECURED MAJOR MILITARY AND DEFENSE INSTITUTIONS WORLDWIDE FOR 25 YEARS.

ARGENT FOR MILITARY & DEFENSE

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 models they use. It checks every request before it leaves and installs inside your enclave. The civil fine for one arms export violation can reach about $1.3 million, or twice the deal's value if that is higher,[1] and Congress barred DeepSeek from Defense Department systems and contract work.[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 DEFENSE TEAMS ASK FOR FIRST

1	**Budget controls** A program's coding agent loops on one task all weekend.	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 ServiceNow or Jira. Project cost tracking shows what each program spends on AI.
2	**Data protection** An engineer's agent sends a design note with an export marking.	Argent reads each request first and blanks out Social Security, card and account numbers before anything leaves (CUI banners, ITAR and EAR legends, CAGE codes and NSNs too, if switched on). The record notes the kind of data, never the data.
3	**Auditing for auditors** An assessor asks what your agents did with controlled data.	Every AI decision goes into a signed record, and any edit is detected. An assessor or your own compliance team can check a package of it offline on their own computer and see whether anything was altered.
4	**Model control** An agent calls a foreign model nobody approved.	Only models you approve are allowed; everything else is refused. Group them into named model classes, such as one for program work and one for general use. A model can be swapped for an approved one without staff changing anything.
5	**Air-gapped** The enclave has no internet, so today it has no AI.	Argent installs inside your enclave and is built to run with no internet. The install bundle is checked at every step as it crosses into your closed environment, and a model you run yourself can serve the enclave.

WHAT REGULATORS ALREADY EXPECT

DFARS 252.204-7012. A contractor must report a cyber incident affecting covered defense information to DoD within **72** hours of finding it.[3]

DFARS 252.204-7021 (CMMC). Keep a current CMMC status at the contract's level for its whole life, and handle FCI or CUI only on systems that have it.[4]

HOW A DEFENSE TEAM 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: controlled 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 PROGRAM'S OWN USE CASES.

NYEDIS.AI/DEMO

Sources. 1 22 CFR 127.10(a)(1)(i), eCFR, read 2026-11-16. 2 Pub. L. 119-60, NDAA FY2026, sec. 1532, Dec 18, 2025. 3 DFARS 252.204-7012(a), (c)(1), eCFR, read 2026-11-16. 4 DFARS 252.204-7021(d), acquisition.gov, read 2026-11-14.