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

ARGENT FOR HEALTHCARE

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. It checks every request before it leaves and runs inside your own walls. Healthcare breaches cost the most of any industry on average,[1] and Medicare tells Medicare Advantage plans AI alone cannot be the basis to deny an admission.[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 HEALTH SYSTEMS ASK FOR FIRST

1	**Budget controls** A billing 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 department spends on AI.
2	**Data protection** An agent sends a discharge note to an outside AI.	Argent reads each request first and blanks out Social Security numbers, plus labelled medical record numbers, member IDs and diagnosis codes once you switch on the healthcare set. The record notes what kind of data, never the data itself.
3	**Auditing for auditors** The privacy office asks which agent sent what, and when.	Every AI decision goes into a signed, tamper-evident record that cannot be quietly changed. An auditor or your privacy office 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.	Only models you approve are allowed; everything else is refused. Group them into model classes, such as approved for patient data and general use. A model can be swapped for an approved one without staff changing anything.
5	**Air-gapped** Patient data must never leave the building.	Argent runs entirely inside your own health system, 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

HIPAA Breach Notification Rule. A disclosure the rule does not permit is presumed a breach, and each affected patient must be told within **60** days of discovery.[3]

HHS civil penalties, 45 CFR 102.3. HIPAA penalties reach up to **$2,190,294** a year for each type of violation, unchanged for 2026.[4]

HOW A HEALTH SYSTEM 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: patient 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 HEALTH SYSTEM'S OWN USE CASES.

NYEDIS.AI/DEMO

Sources. 1 IBM / Ponemon, Cost of a Data Breach Report 2026, July 2026. 2 CMS, HPMS memo, FAQs on Coverage Criteria and Utilization Management (CMS-4201-F), Feb 6, 2024. 3 45 CFR 164.402 and 164.404, eCFR, read 2026-11-14. 4 45 CFR 102.3 (2025 amounts), eCFR, read 2026-11-19; Federal Register 2026-21524, Oct 21, 2026.