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

ARGENT FOR UTILITIES

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 utility. FERC audit staff report growing use of AI tools inside grid networks,[1] and IBM puts the average energy breach, utilities included, at $5.24 million.[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 UTILITIES ASK FOR FIRST

1	**Budget controls** An outage 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 through your ServiceNow or Jira tickets. Project cost tracking shows what each project or business unit spends on AI.
2	**Data protection** An engineer pastes a substation device list into an AI.	Argent reads each request first and blanks out IP addresses, phone numbers, Social Security, card and bank account numbers and your own formats before anything leaves the utility. The record notes the kind of data, never the data.
3	**Auditing for auditors** A CIP auditor asks which AI tool sent what, and when.	Every AI decision goes into a signed, tamper-evident record that cannot be quietly changed. A CIP auditor or internal audit 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 grid data and general use. A model can be swapped for an approved one without staff changing anything.
5	**Air-gapped** Grid data must never leave the utility's network.	Argent runs entirely inside your own utility, 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

NERC CIP-004-7. When a person leaves, a utility must remove their remote and physical access to its most critical grid systems within **24** hours.[3]

FERC audit staff, October 2026. AI tools in the grid's cyber networks may create access risk when access is not "properly authorized, verified and monitored."[1]

HOW A UTILITY 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: grid and customer 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 UTILITY'S OWN USE CASES.

NYEDIS.AI/DEMO

Sources. 1 FERC staff report, 2026 Lessons Learned from Commission-Led CIP Reliability Audits, Oct 13, 2026. 2 IBM / Ponemon, Cost of a Data Breach Report 2026, Jul 2026 (energy includes utilities). 3 NERC CIP-004-7, Table R5 Part 5.1, effective Jan 1, 2024.