ATLAS PLATFORM

Atlas AIDR

AI usage observability, threat detection and control. Every AI agent, model, and MCP server your developers run. Seen, attributed, and accounted for.

Rapid AI Adoption Requires Advanced AI Security

Knowing which AI tools your people signed up for is the easy part. AI coding and agentic AI act with real credentials and broad permissions pulling in third-party MCP servers without being reviewed, touching code and data, and taking autonomous action. Most organizations can't say what their AI agents are doing, and at what cost. Atlas AIDR turns every instrumented AI session into an attributable record, screens every prompt for secrets, PII and injection, and enforces redaction and block rules inline for routed traffic. When a detection fires on AI telemetry, our SOC investigates and responds to mitigate risk associated with unruly AI activity.

76%
of employees now use AI at work, up from 30% two years ago.
McKinsey, The State of AI, 2025

43%
of breached organizations had shadow AI involved, more than double last year.
IBM Cost of a Data Breach Report, 2026

33%
of 1,000 scanned MCP servers carried critical vulnerabilities.
Enkrypt AI MCP Server Scan, October 2025

What Atlas AIDR Delivers

AI Estate Visibility

- Users, teams, and workstations tied to every AI session
- Agents, models, and providers inventoried across the estate
- Full session timelines with every prompt, response, and tool call
- A live inventory you can hand to an auditor, not an annual survey

Agent Supply Chain Governance

- The MCP servers and tools your agents load
- Which teams and agents depend on each one
- Supply-chain drift fed into detection, not just catalogued

Inline Control and Spend

- PII detected in every prompt, and redacted before the model sees it on routed traffic
- Prompt-injection screening and block rules enforced at the proxy
- Your own business-logic rules applied in-flight
- Token economics per request, with spend anomolies surfaced in costs

How It Works

1. **Connect**
Install once per workstation, by policy or a single command. Capture runs the tools' own hooks; route traffic through inline proxy where you want redaction and cost per request.

2. **Capture**
Every prompt, response, tool call, and token becomes a record, tied to a user, team, and workstation.

3. **Inspect**
PII is redacted and prompt injection screened in the request path, before content reaches the model.

4. **Enforce**
Block rules, budgets, and your own business-logic rules applied in-flight, on an SSO-governed control plane.

5. **Respond**
High-severity findings are surfaced to the same Atlas console your SOC works from. TRU hunts proactively, before the attacker moves.

See Atlas AIDR in Action

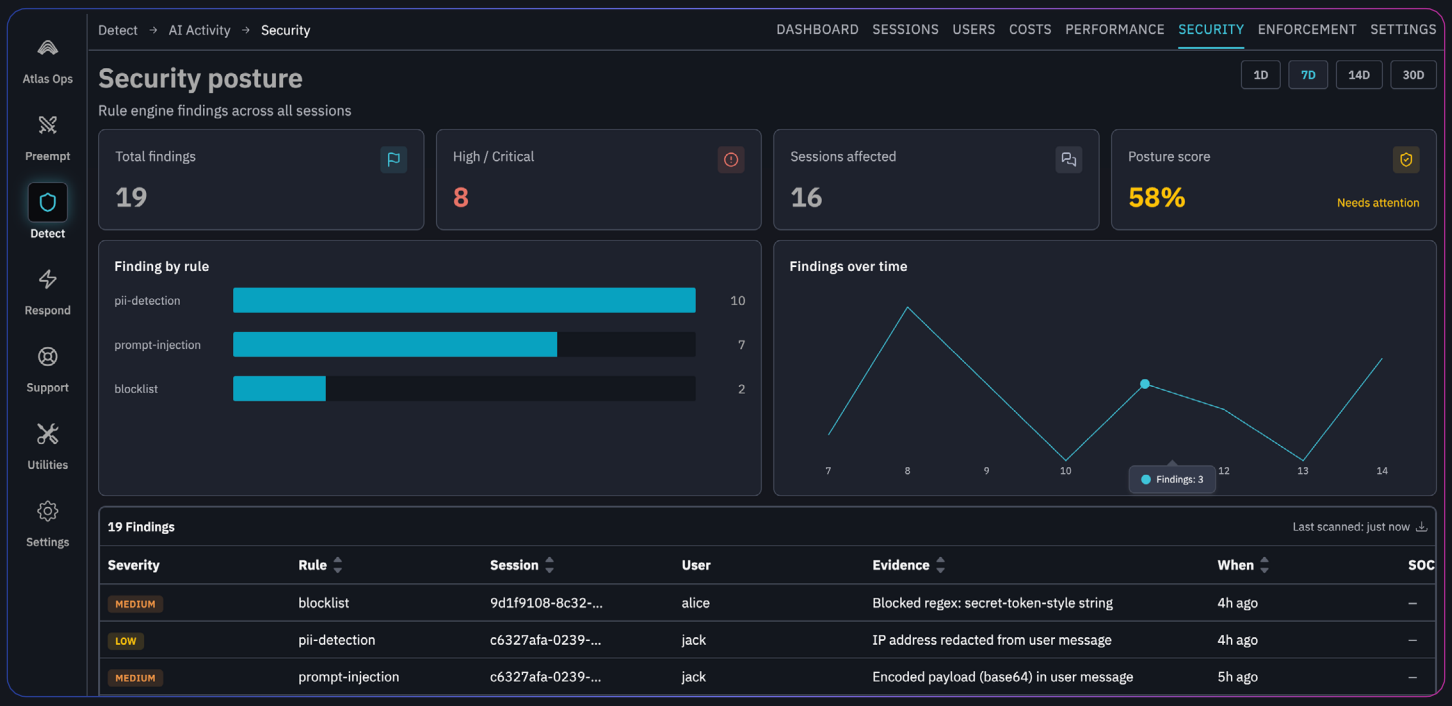

Typical AI Security Tooling vs Atlas AIDR

	Typical AI Security Tooling	Atlas AIDR
Capability		
Inventory of AI agents, models, and providers	✓	✓
Prompt, response, and tool-call capture	✓	✓
PII redaction and prompt-injection screening in the request path	✓	✓
MCP server and tool inventory attributed to teams	–	✓
Monitoring for server behavior change after approval	–	✓
Token economics per request and anomalous-spend alerting	–	✓
Full-session capture in Atlas, with retention set to your policy	–	✓
Correlation with endpoint and network telemetry in the same Atlas console as the rest of your estate	–	✓
24/7 human-verified investigation included, no extra modules	–	✓

Multiple
Model providers covered inline

Minutes
To your first captured session

100%
Prompt to tool call — every step of an instrumented session, attributed to a person and a machine

24/7
SOC and TRU investigation included

Outcomes You Can Expect

- An attributable record of AI activity your auditor will accept
- PII caught and redacted before it reaches a model, not after
- Prompt-injection attempts screened and blocked in the request path
- The MCP servers your agents depend on, inventoried and watched for change
- Anomalous AI spend surfaced as a compromise signal, the day it starts
- AI sessions investigated by the same SOC that watches the rest of your estate

> *"We look at eSentire to be the experts. We trust them implicitly. They're with us through the thick and thin till the end."*

CISO, Financial Services GARTNER PEER INSIGHTS

2,000+ organizations in **80+** countries | **25** years of SOC expertise | **2M+** endpoints protected | MDR Leader, IDC Marketscape for MDR

Assume your agents are already doing things you cannot account for.

See what your AI did yesterday, and what to do about it.

LET'S TALK SECURITY →

IF YOU'RE EXPERIENCING A SECURITY INCIDENT OR BREACH CONTACT US 1-866-579-2200

eSentire

eSentire is a leader in Controlled Autonomy SecOps, protecting 2,000+ organizations across 35+ industries around the world. Founded in 2001, the company's Controlled Autonomy SecOps operating model pairs agentic AI operatives with engineered human-judgment controls, delivering expert-depth security outcomes at machine speed without ceding accountability to opaque automation. Powered by the unified agentic AI Atlas Platform, eSentire's Atlas AI + 24/7 expert human SOC coverage delivers offensive capabilities that preempt exposures before attackers do, detect, and respond to stop threats in real time. For more information, visit **www.esentire.com** and follow **@eSentire**.