

Service Description:

Atlas AIDR

AI Detection and Response

Service Overview

eSentire's Atlas AIDR (the "Service") is a capability within the Atlas Platform that gives Client visibility of AI activity across its environment — which AI Applications, models, and MCP Servers are in use and by which teams, what data those sessions touched, what the AI is costing, and any threatening behaviour observed. The Service inventories AI Applications, models, and MCP Servers in use, screens for sensitive data and prompt injection, and surfaces Findings in the Atlas UI.

The Service is delivered through an Atlas Platform-hosted inline proxy; there is no separate console to operate and no separate vendor relationship. The Service is available as an Add On to Atlas Professional, Enhanced, and Elite Packages, or as a standalone service, and is licensed on a per-user basis, as detailed on the Order Form.

Service Definitions

In addition to the below, capitalized terms contained in this Service Description are as defined herein, or as defined in the "Managed Detection Response ("MDR") Services - General Information" document (the "MDR General Information" document), available at <https://www.esentire.com/legal/documents>. The MDR General Information document contains information applicable to all eSentire MDR services, including this Service.

"AI Application" means any Client-operated coding agent, copilot, application, or service that sends requests to a Model Provider, including without limitation Claude Code, Cursor, GitHub Copilot, Codex, and Gemini CLI.

"Atlas UI" means the user interface of the Atlas Platform where interaction with and configuration of the Atlas Platform and MDR services is performed.

"Client Environment" means the Client systems from which Client Data is collected for the delivery of the Service.

"Findings" means investigation results delivered to Client in the Atlas UI.

"MCP Server" means a Model Context Protocol server, plugin, or hook that provides an AI Application or agent with tools or system access.

"Model Provider" means the third-party AI model provider(s) in scope for the Service, including Anthropic, OpenAI, Google, and Amazon Bedrock, and additional OpenAI-compatible providers available on request.

"Routed Traffic" means traffic generated by an in-scope AI Application that Client has configured to pass through the Service.

Service Capabilities

AI, Agent & MCP Server Inventory

The Service inventories the AI **Applications**, models, and MCP Servers, plugins, and hooks in use across Client Environment, including which teams are using them.

Sensitive Data & Prompt Injection Screening

Routed Traffic is screened using pattern-based detection for prompt injection and sensitive data, before content reaches the Model Provider. eSentire-defined, Client-enabled block rules, business-logic rules and security rules are applied.

AI Spend Attribution

The Service reports AI usage and associated cost attributed by team, giving Client's finance and procurement functions a cost view they typically cannot otherwise produce, and giving eSentire's SOC an early signal of anomalous token consumption indicative of a hijacked or looping agent.

Data Access and Reporting

Findings, the AI and MCP Server inventory, and AI spend data are available through the Atlas UI, correlated with other telemetry from the rest of Client's Atlas estate.

SOC Alerting and Investigation

Where Client has purchased and implemented other eSentire MDR services, Findings are delivered to eSentire's SOC and investigated alongside Client's broader MDR telemetry, consistent with the MDR General Information document.

Coverage and Scope

The Service supports Anthropic, OpenAI, Google, Cursor, GitHub Copilot, Codex, and covers any AI Application instance routed through the proxy. Support for additional applications/providers will be added over time.

The Service is scoped to desktop AI agents and applications. Browser-based AI activity is out of scope for the Service.

Deployment

eSentire delivers the Service through an Atlas Platform-hosted inline proxy. During onboarding, eSentire and Client agree on the AI Applications and Model Providers in scope and configure initial screening, redaction, and block rules, together with any business-logic rules Client wishes to define.

Client then adds a configuration to its in-scope AI Applications, pushed centrally through Client's existing policy tooling, directing Routed Traffic to the Service. The first captured session typically arrives within minutes of the configuration being applied.

Maintenance and Support

The proxy is Atlas Platform-hosted and delivers the Atlas Platform uptime SLA described in the MDR General Information document. eSentire maintains and updates the Service's screening and detection logic and supports Client in tuning block and business-logic rules, onboarding additional AI Applications or Model Providers, and investigating flagged detections.

Responsibilities

Client is responsible for:

- identifying and inventorying the AI Applications, Model Providers, and teams in scope for the Service;
- deploying and maintaining the configuration that routes in-scope AI Applications to the Service, including via Client's own policy tooling;
- defining and maintaining any business-logic or block rules specific to Client's environment;
- granting eSentire the access necessary to configure and support the Service;

- ensuring Client's deployment of the Service complies with Client's particular legal, security, or data privacy requirements; and
- providing the resources, information, documentation, and access to personnel, equipment, and systems reasonably required by eSentire to perform the Service.

Service Level Objectives

Service levels of this Service are set out in the MDR General Information document.

Service Turndown

Upon expiration or termination of the Service, eSentire will disable the Service configuration and cease accepting routing of Client's AI Application traffic through the Service proxy. Instructions will be provided to restore default settings to Client's AI Applications. Session data, metadata, findings, and investigations collected by the Service are destroyed or allowed to expire, in accordance with Client's standard Atlas Platform data retention and destruction policies.

Service Terms

Exclusions

The Service excludes: monitoring of browser-based AI activity, governance of AI Applications, models, or MCP Servers that Client has not configured to route through the Service.

Appendix A: Responsibilities Matrix

Function	Client	eSentire
AI Application / Model Provider inventory & scoping	RA	C
Service configuration & routing deployment	RA	C
Business-logic / block rule definition	RA	C
Screening & detection content (standard library)	I	RA
MCP Server approval workflow	RA	C
MCP Server behavior monitoring (post-approval)	I	RA
PII detection & redaction	I	RA
Alert monitoring & investigation*	I	RA
AI spend reporting	I	RA
Atlas proxy uptime & maintenance	I	RA

R = Responsible; responsible for action and implementation. Responsibility can be shared.

A = Accountable; ultimately answerable for the activity or decision, including "yes" or "no" authority and veto power.

C = Consulted; subject matter experts consulted prior to a final decision or action.

I = Informed; needs to be informed after a decision or action is taken.

* Applies where Client has purchased and implemented other eSentire MDR services.