

eSENTIRE

ATLAS 

Frequently Asked Questions

What is eSentire Atlas AI and how does it enhance security operations?

Atlas AI is an advanced multi-agent Generative AI system embedded within eSentire's Atlas Security Operations Platform. It's designed to scale human expertise rather than replace it, acting as an integral component to automate and enhance security operations.

Atlas AI accelerates threat detection and response through autonomous reasoning and action, aiming to think, act, and improve outcomes like an expert. It's purpose-built on real-world workflows validated by investigations across a global customer base, providing transparency, context, and validation in minutes.

How does eSentire's Atlas Security Operations Platform integrate with AI/LLM systems?

The eSentire Atlas SecOps Platform incorporates Generative AI models into its architecture to enhance various aspects of security operations for both eSentire delivery teams and customers. This integration facilitates functions such as Agentic AI for initial alert investigations, AI Auditor for providing feedback on analyst findings, and Point Tooling for enhancing investigations (e.g., Command Line Explainer, Description Writer, Log Search Query Builder).

It also enables natural language interaction with Atlas for users (e.g., AI Investigator, Help Assistant) and allows internal eSentire users to leverage Gen AI for building low-code Atlas Actions.

How secure is Atlas AI? How do you govern and control AI in your security operations?

eSentire's Atlas AI utilizes enterprise-ready security architecture built specifically for sensitive environments. Using commercial grade models and applying them across our data mesh, we have a containerized system with 24/7 monitoring within SOC operations and a comprehensive AI Governance/Oversight program. Our secure AI implementation maintains complete records of every investigation step, reasoning process, and tool call for full transparency.

Can Atlas AI perform actions unilaterally without human input?

No, the Atlas AI system cannot perform response or remediation actions unilaterally. All output from Atlas AI is reviewed and approved by a human. The system is designed to facilitate investigations and analysis, but it cannot take any direct action, such as isolating a host, executing financial transactions or posting information publicly, without human input and approval. Its current capabilities are limited to the collection of investigative data.

What role do human experts play in the Atlas AI system?

Despite the advanced AI capabilities, human experts are integral to the Atlas AI multi-agent system and are “in the loop every step of the way.” For threat investigations, Atlas AI is implemented to enable SOC and customers to have a preliminary investigation with a great level of details, designed to be reviewed and approved by a human, as Atlas AI currently cannot take direct actions beyond collecting investigative data.

eSentire SOC Analysts and customers can review, validate, enrich, and tailor findings generated by the AI, ensuring they are accurate and align with the customer’s business context and risk tolerance. This human oversight ensures accuracy, quality, and relevance of the AI’s output.

How does eSentire ensure the transparency, accuracy, and fairness of data processing by the Atlas AI system?

The outputs from Atlas AI investigations are available to users in the Findings view in Atlas. This provides details of the analysis and investigation steps Atlas AI used in its investigation. Further, eSentire ensures the quality of the output from Atlas’s Gen AI layer through its rigorous software development, architecture, and quality assurance processes.

Every AI-driven decision undergoes review, refinement, and action by a SOC expert. Furthermore, eSentire employs ongoing testing and monitoring to secure interactions with Atlas’s AI, including controls to prevent “data poisoning” even though eSentire does not directly train the AI models.

What data sets are the AI models used by Atlas AI originally trained on, and how does Atlas AI handle customer data?

eSentire utilizes leading third-party generalized AI models and does not train these models directly. Consequently, customer data is explicitly not used to train these AI models. Atlas AI interacts with current data within the Atlas Security Operations Platform, queries data from connected technologies, and can also query external sources. It may use historical data when necessary for analysis to provide context. Atlas AI also has the capability to decrypt common encryption used to obfuscate code for further analysis.

Does Atlas AI integrate with existing security infrastructure and tooling?

Yes, Atlas AI is fully embedded into eSentire’s Atlas SecOps Platform and is included as part of the eSentire MDR service. The Atlas SecOps Platform supports over 300 best-of-breed technology integrations, with a continuously expanding list of new integration partners. This extensive platform integration capability allows Atlas AI to utilize data from EDR/EPP, Network, Log, Identity, Email security, SaaS platforms, VPN providers and Web gateway technologies connected to the platform for investigations.

What is eSentire’s Predictive Threat Defense Network?

eSentire’s Predictive Threat Defense Network describes a dynamic and adaptive security ecosystem where an attack on one customer strengthens the defenses for all. Unlike competitors who offer standalone tools, eSentire provides a living, learning defense network that evolves with emerging threats. This network operationalizes intelligence, predicts workflows, prevents threat disruption, and continuously improves, leveraging the collective experience from its global base of over 2000 customers across various industries and countries.

eSentire Atlas Security Operations Platform, Powered by Atlas AI

AI You Can Trust. Outcomes You Can Prove.

eSentire's Agentic AI isn't just AI — it's 25 years of SecOps expertise in action. While others chase AI hype, we deliver certainty — faster out of the investigation starting blocks, expert-validated, outcome-driven responses at scale. Our Agentic AI doesn't just detect threats; it thinks, acts, and improves your protection 24/7. We show up and prove results:

- 35% faster threat intelligence vs commercial feeds
- 99% noise reduction across customer environments
- 95% SOC expert alignment with Atlas AI investigations
- 99.3% of threats isolated at the first host
- 200 new threat protections added per day to harden customer defenses
- 43X investigation acceleration with 5 hours of investigation work achieved in less than 7 minutes
- 96% SOC analyst retention, with an average tenure of 6 years

LEARN MORE →

Ready to build an MDR quote?

Our Cybersecurity Specialists are ready to help you understand our MDR pricing packages and determine the ideal solution for your organization.

**Request a free MDR quote or try the free
eSentire MDR Package Selector.**

BUILD A QUOTE →

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

eSENTIRE

eSentire, Inc., the Authority in Managed Detection and Response (MDR), protects the critical data and applications of 2000+ organizations in 80+ countries, across 35 industries from known and unknown cyber threats by providing Exposure Management, Managed Detection and Response and Incident Response services designed to build an organization's cyber resilience & prevent business disruption. Founded in 2001, eSentire protects the world's most targeted organizations with 65% of its global base recognized as critical infrastructure, vital to economic health and stability. By combining open XDR platform technology, 24/7 threat hunting, and proven security operations leadership, eSentire's award-winning MDR services and team of experts help organizations anticipate, withstand and recover from cyberattacks. For more information, visit www.esentire.com and follow [@eSentire](https://twitter.com/eSentire).