AI SOC
AI-driven SecOps

Our AI Agent autonomously triages, investigates, and responds to every alert in your system with the highest quality actions so you finally have time to breathe, sleep, and maybe even take that vacation.

Transform your SecOps into AI SOC with our AI Agent. Improve both your strategic and tactical metrics and proactively hunt for threats.

Explore Our Solutions ↓

Book a Demo

Experience the power of our AI Agents today

Rapidly changing threat landscape

01

AI-Powered cybercrime

Groups that were only qualified to scam your Grandma are now capable of Advanced Persistent Threat (APT) level attacks at *an infinite scale.*

02

Alert fatigue & burnout

You and your colleagues are over-worked and always on the back foot. AI SOC agents do the heavy lifting so you can focus on what matters most proactively.

03

Slow detection & containment

258 days to identify and contain a breach (according to IBM's Cost of a Data Breach Report 2024) is absurd. That ends with Simbian.

04

True positive, false positive - It's all work

258 days to identify and contain a breach (according to IBM's Cost of a Data Breach Report 2024) is absurd. That ends with Simbian.

Simbian AI SOC Improves MSSP Quality of Service While **Reducing Costs By 5X**

3 Collaborators • Created 2 weeks ago

Running	Completed	Failed
2	1,320	0

Let our agents do the heavy lifting

In the rapidly shifting threat landscape powered by attackers with AI, and a security program that was already fighting a barrage of alerts, Simbian Hands lets you take control. With our agents covering all of your bases, you'll finally have the time and energy to prioritize proactive defense.

Learn More ↗

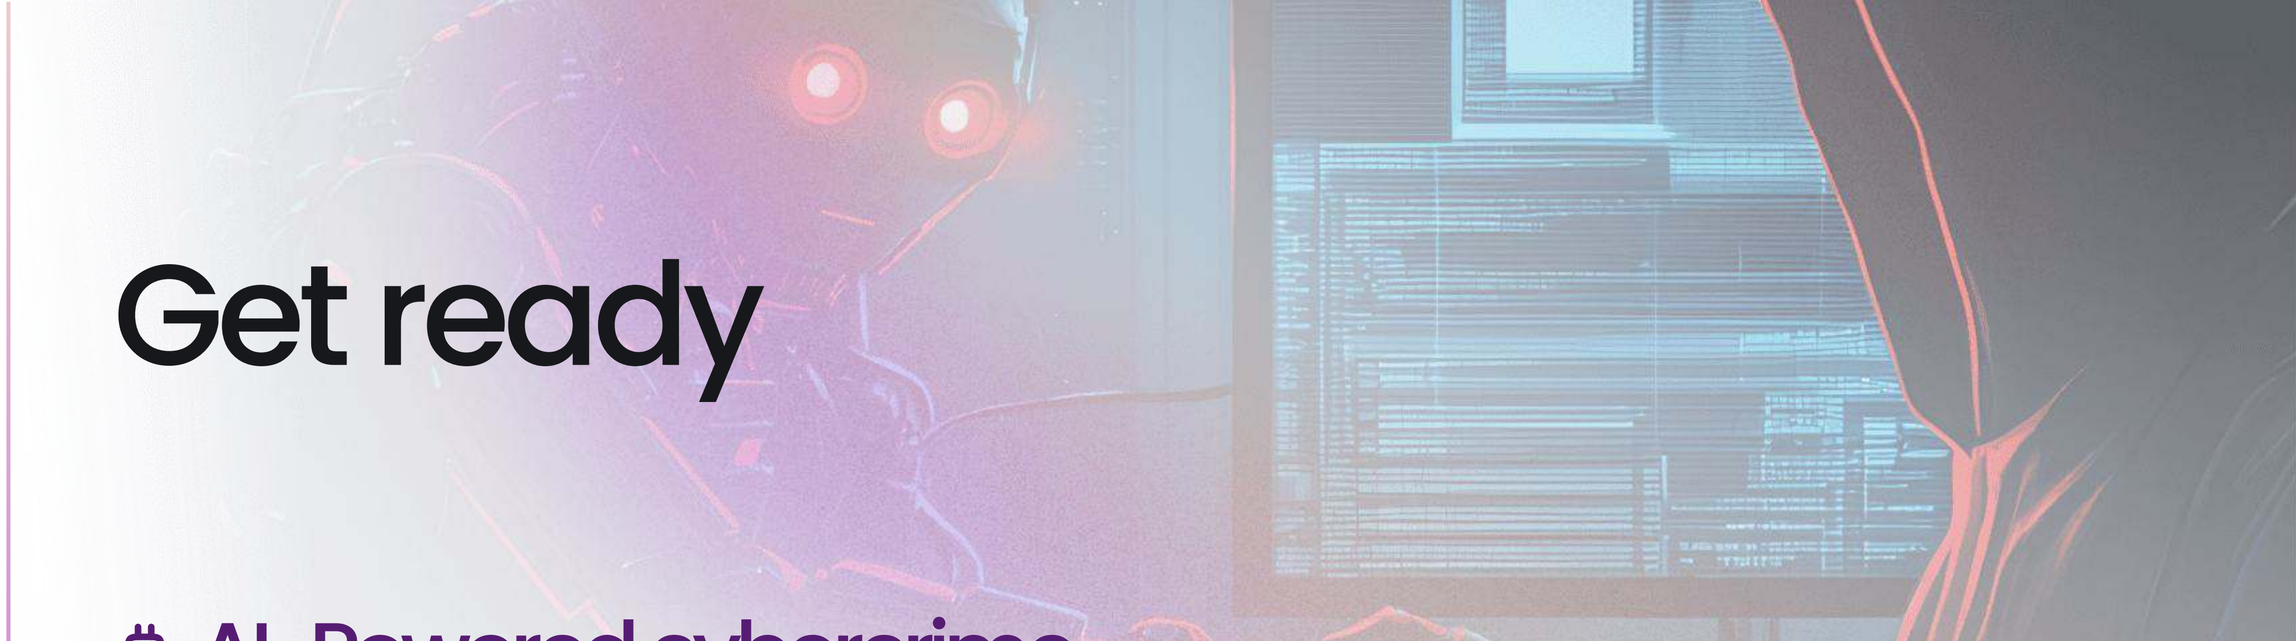

Get ready

AI-Powered cybercrime

Sophisticated social engineering

AI can now generate highly convincing phishing emails with deepfake to trick even the most cautious employees. Attackers can harvest all of the publicly available information about your organization and can further customize the vectors to exploit the unique circumstances of every employee.

Automated vulnerability exploitation

We have always known that everyone has a long list of unpatched vulnerabilities in systems due to operational challenges. Unfortunately, AI-powered bots know that too, and can continuously scan networks for security weaknesses, learning from each successful intrusion and adapting their tactics. They can create incidents more easily and faster than humans can respond.

Automated malware adaptation

Reverse-engineering code from assembly and modifying malware to evade traditional detection systems once required well-funded adversaries; now they are within reach of many LLMs. AI-powered polymorphic malwares analyze their environment and adapt their behavior to bypass security controls, making them particularly difficult to identify and neutralize.

Attackers Use AI
Defeat them with AI

Regain control

Alert fatigue & burnout

Relentless barrage of alerts

AI can now generate highly convincing phishing emails with deepfake to trick even the most cautious employees. Attackers can harvest all of the publicly available information about your organization and can further customize the vectors to exploit the unique circumstances of every employee.

False positive alerts play Russian roulette

Within this torrential downpour of alerts, 90% are "false positive" - notifications unrelated to immediate threats. Often caused by misconfigured rules, imprecise detection signatures, or the inherent limitations of security tools, these should be a starting point for the detection team. However, analysts lose valuable time triaging them, leading to productivity loss and, more critically, real threats potentially getting buried in the backlog. These "crying wolf" alerts also contribute to analyst burnout and decreased vigilance.

Missing investigative context

The last nail in the coffin of analyst morale tends to be the missing context about the alerts. Security information is sensitive and spread across fragmented tools. For example, an analyst may receive a "suspicious travel alert" (one of the most frequent sources of false positives) without details on the user, their reporting structure, their travel plans, or the potential business impact. This forces analysts to expend considerable time and effort manually gathering context from disparate systems. This inefficient, context-starved process significantly slows down response times and increases the cognitive burden on analysts.

Faster, better, more Alert Triage

Unify & analyze

Slow detection & containment

Greater attack surfaces than ever before

Organizations have increasingly complex IT environments (cloud, hybrid, IoT, mobile). There has been greater digital transformation, along with an often globally distributed workforce, many of them working from home -- all contributing to an ever-increasing attack surface. This expanded attack surface makes it harder to monitor everything and see malicious activity across all these diverse environments. These trends are further exacerbated by a growing trend of localization and AI-sovereignty - both of them leading to a distributed digital footprint.

Increasing sophistication of attacks

Attackers are using more advanced techniques (some powered by AI), sophisticated phishing attacks, and stealthier techniques that utilize legitimate system tools and processes to carry out malicious activities without introducing suspicious code. These techniques, on the one hand, are able to trick more employees than ever before, and are also harder to detect since they blend malicious operations with normal system behaviors,

Data silos

Gone are the days where there was a centralized SIEM. Due to the distributed nature of the cloud, many large organizations today have their security logs distributed across a variety of SIEMs and data lakes. Even smaller organizations may have considerable amounts of security logging in their cloud instances.Centralizing these logs is expensive and time-consuming. The traditional way of detection and response doesn't quite work anymore.

Data Analyzed in One Place

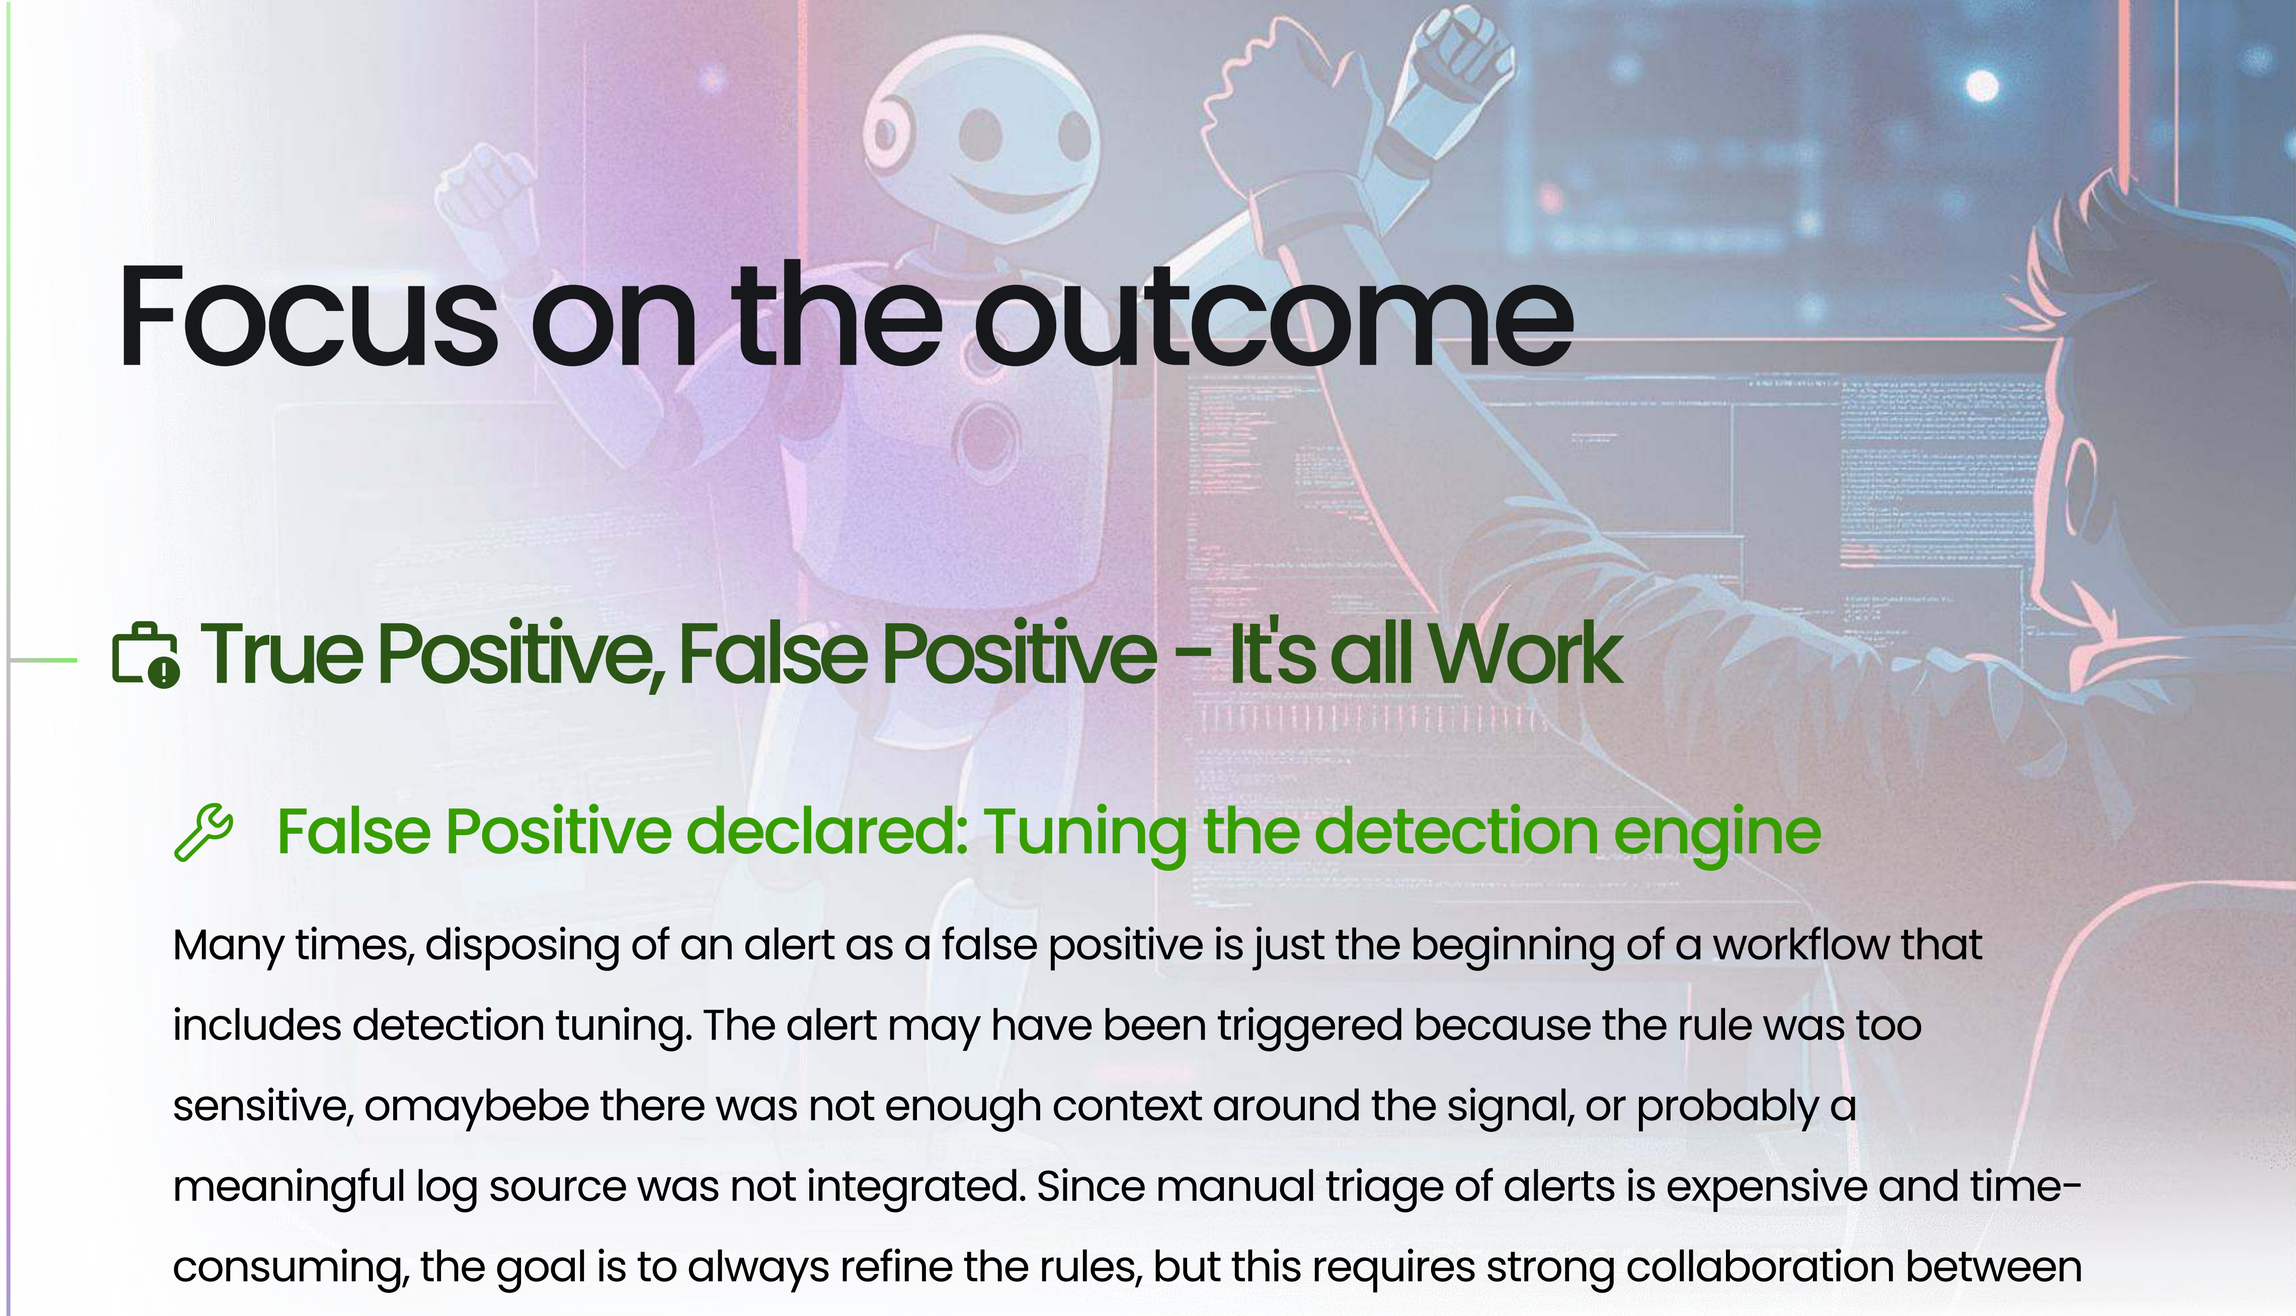

Focus on the outcome

True Positive, False Positive – It's all Work

False Positive declared: Tuning the detection engine

Many times, disposing of an alert as a false positive is just the beginning of a workflow that includes detection tuning. The alert may have been triggered because the rule was too sensitive, omaybebe there was not enough context around the signal, or probably a meaningful log source was not integrated. Since manual triage of alerts is expensive and time-consuming, the goal is to always refine the rules, but this requires strong collaboration between two busy teams.

True positive found: Kick-off incident response

If there is one thing SOC teams dread more than false positives, it's true positives! This chilling confirmation throws the security organization into high gear. The incident response protocols may span across not only security tasks like containment, recovery, and hunting but also often bleed into governance and risk compliance teams. The amount of context gathering and sharing is multiplied, along with the pain of doing so across disparate systems, on a tight timeline.

Patching exploited cracks

Every exploit shines a harsh spotlight on underlying vulnerabilities. Attackers rarely materialize from thin air; they exploit weaknesses in systems, people, and applications. When a true positive indicates a successful exploit, the Vulnerability Management team is immediately compelled into action. They need to urgently determine: Which vulnerability was exploited? Is it a known vulnerability? Is it patched? Are other systems vulnerable to the same exploit? The true positive acts as a critical, real-world validation of a vulnerability's risk. Too much to do, too little time.

Getting Straight to What Matters

Are you the right fit for an AI SOC?

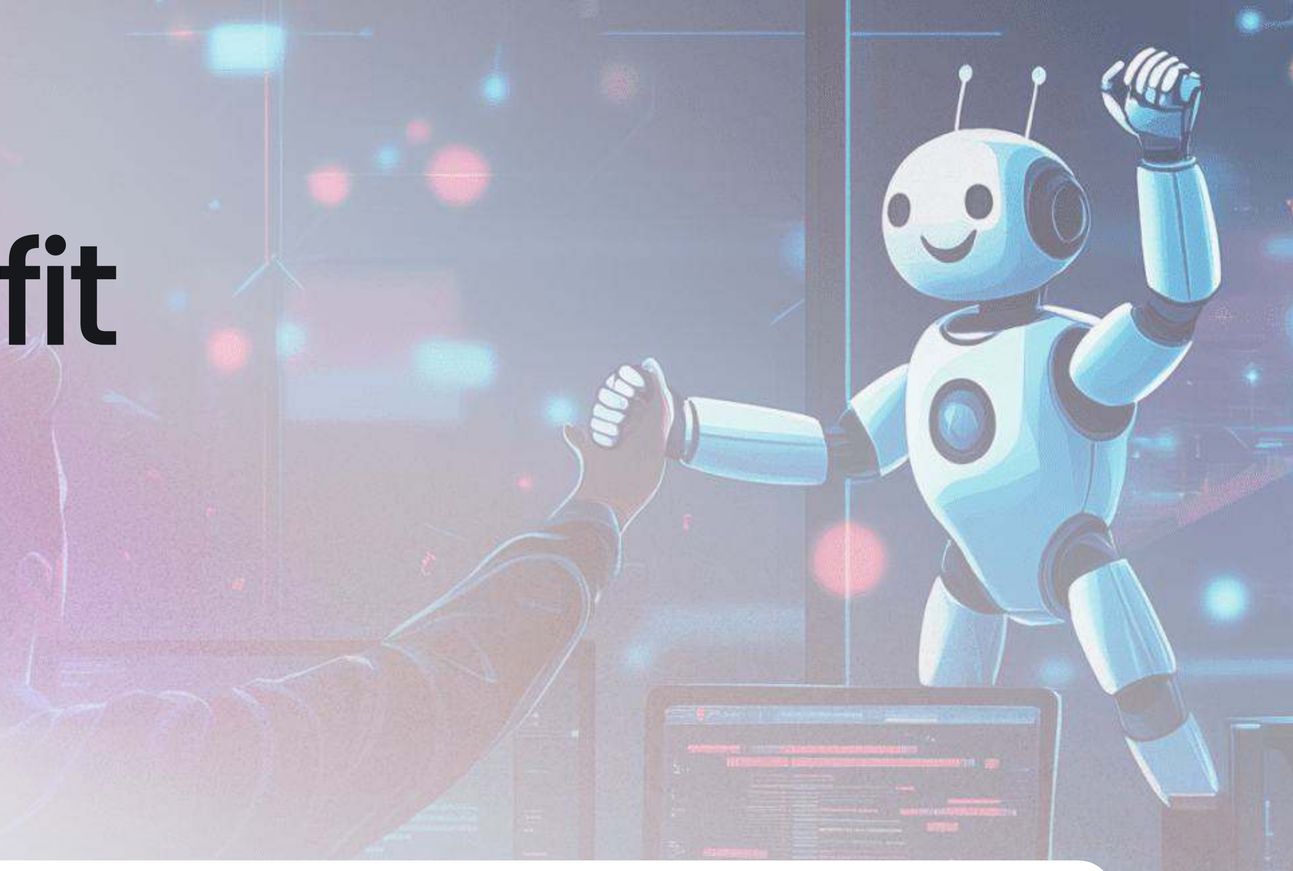

If your top priorities include...

Growth

You need to grow your SOC to manage a rising investigation workload but are struggling to hire and train analysts.

Faster response times

You're ready to drive down Mean Time to Identify (MTTI) and Mean Time to Contain (MTTC).

Automation

You want to automate, but SOAR turned out to translate into endless work manually building playbooks.

...Then you're our mission.

Our fully AI agents take over the heavy lifting within your SOC instantly. Though co-piloting is an option and can help your analysts teach our agents how best to meet your SOC's individual needs and standards, we offer full autonomy right from the start. The moment you enable Simbian's AI SOC agents, your workload plummets.

The numbers tell our story

12x

Coverage*

92% of alerts getting resolved allows you to do more.

7x

Insights

All analysts are guided by deep insights and in-depth guidance.

10x

Faster

From 30 minutes manual to 2 minutes triage. Down to just 10 to investigate!

How?

Our **patented technology** sees everything in your environment and finds the fastest route to connect all the dots.

Simbian has analyzed hundreds of thousands of results from real environments, and the results **speak for themselves.**

Ambuj Kumar, Founder & CEO, Simbian

*In a typical SOC environment (24x7 coverage with 3 tiers of analyst) when benchmarked against a diverse set of use-cases across EDR and SIEM.

Our agents, your tribal knowledge

We understand the complexity of the way you do things. Our agents ingest all of your documented standard operating procedures, but more importantly, learn from analyst feedback over time. So all that undocumented tribal knowledge is absorbed, documented, and honored every time our agents investigate through our patent-pending **Context Lake** technology.

Learn More ↗

Alert Investigation

All of your alerts from all your data sources are ingested into one place and investigated by our virtual SOC analysts.

Response

Want to get proactive and engage your SecOps program in deep-dive, threat-intel-driven hunts? Simbian automates this, too.

Threat Hunting

Suffering from security questionnaires pouring in from your customers? Our GRC Agent answers for you!

TrustedLLM™ Technology

Simbian's agentic technology leverages a powerful layer of hallucination and injection-resistant techniques within its Chain-of-Thought (CoT) and Retrieval Augmented Generation (RAG) reinforced inference systems.

Before Simbian was an AI-for-security company, it began with the goal of security for AI. While we evolved far beyond that original mission, our research and development efforts led to industry-leading LLM security and reliability technologies. As a customer of Simbian, you'll benefit from this reliable, secure operational framework which powers all of our agents.

Experience Safety of TrustedLLM Forever

Improved metrics

Drive your KPIs forward

Data-Driven decision making

In the ever-changing security landscape, timely and accurate data is crucial for strategic decision-making. Simbian's AI SOC provides comprehensive analytics and reporting, offering real-time insights into various performance metrics. By automating data collection and analysis, your team can focus on refining strategies and improving operational efficiency. Our system's robust analytics help pinpoint areas for improvement and drive KPIs forward with precision.

Increased threat detection accuracy

Simbian's AI-driven approach enhances threat detection accuracy by reducing false positives and negatives. Our intelligent agents continuously learn from new data and feedback, refining detection algorithms to differentiate between benign and malicious activities effectively. This improved accuracy not only elevates security posture but also contributes significantly to reduced Mean Time to Detect (MTTD) and Mean Time to Respond (MTTR) metrics.

Optimized resource allocation

Our AI SOC optimizes resource allocation, enabling your team to focus on high-priority tasks. By analyzing workload, alert types, and response needs, Simbian helps streamline operations and reduce unnecessary labor hours. This allows for better allocation of skilled resources, ensuring your team is always working on tasks that provide the highest return on investment.

Metrics that Matter

Retained learning

Incorporating experience for future Resilience

Continuous knowledge integration

Simbian's Context Lake technology ensures that all learned knowledge, both documented and experiential, is integrated into AI agents' operations. This process diversifies and deepens the learning repository, ensuring that even unanticipated threats are met with informed responses. Our system absorbs your team's insights, seamlessly incorporating them into its evolving operational blueprint.

Feedback-Driven evolution

Our agents are designed to retain and utilize analyst feedback, allowing them to evolve continuously and improve with each interaction. This ongoing dialogue between AI and human expertise enhances decision-making and refines response strategies, embedding a rich tapestry of tribal knowledge that bolsters security over time. Feedback becomes a pivotal component of agent learning, turning individual experiences into collective SOC wisdom.

Preserving institutional intelligence

As personnel change or shift roles, essential security practices and insights often risk being diluted or lost. Simbian prevents this by capturing and retaining institutional knowledge within our agents, ensuring that insights do not leave with employees. This retention of intelligence preserves continuity, enabling seamless transitions and sustained operational strength.

Knowledge That Grows

LLM choice & privacy

Balancing innovation with integrity

Strategic LLM selection

At Simbian, the choice of language model is crucial to providing a secure and efficient service. We carefully evaluate LLMs based on performance, adaptability, and security, ensuring that they meet our rigorous standards for operational excellence. This strategic approach allows us to deliver state-of-the-art solutions while maintaining the robustness necessary for critical security environments.

Privacy-First architecture

Privacy is foundational to our operations. Simbian's architecture is designed to protect sensitive data throughout the processing chain, employing state-of-the-art encryption and security protocols. By ensuring that data integrity and confidentiality remain uncompromised, we can deliver effective AI solutions with peace of mind regarding compliance and privacy standards.

Patching exploited cracks

Our commitment to privacy extends to how data is managed within the AI SOC ecosystem. All information processed by our AI agents is handled with the utmost care, aligned with both legal and ethical standards for data security. Simbian employs rigorous data segregation and protection practices, ensuring that your organization's information is safeguarded from unauthorized access or breaches.

Security in Every Byte

Tailored to your SOC

Tailored solutions for diverse needs

Customized integration

Every organization has its unique security requirements dictated by varying IT infrastructures, industry regulations, and operational nuances. Simbian's AI SOC is adept at integrating seamlessly with your existing systems, adapting to the specific context of your technological stack to ensure maximum compatibility and effectiveness without disrupting workflows.

Adaptive intelligence for specific needs

Our AI agents learn and evolve specifically for your business environment, factoring in industry-specific threats and internal policies to deliver bespoke solutions. This personalized adaptation ensures that our technology actively supports your operational objectives and security posture, seamlessly fitting into and enhancing your ecosystem with its tailored capabilities.

Dynamic response strategies

The dynamic nature of business environments necessitates agile response mechanisms. Simbian's AI SOC is built to adapt to any changes within your organization, recalibrating detection and response functions to match shifts in structure, policy, or focus. By doing so, we ensure that your security response remains agile and effective, no matter how your operational landscape evolves.

Tailored to You & Your Environment

Comprehensive coverage

All alerts, all tools, across all environments

Unified alert ecosystem

Simbian's AI SOC creates a unified alert ecosystem, integrating seamlessly with all tools and monitoring systems in your network. Our platform acts as a central hub where alerts are aggregated, analyzed, and triaged, ensuring no threat goes unnoticed and every alert is handled efficiently, freeing your team to focus on strategic security initiatives.

Tools-Agnostic approach

Our agents are designed to work with any security tool or system you already employ. Simbian offers a tool-agnostic approach, allowing for flexible integration across a broad spectrum of security products. This flexibility ensures that your AI SOC can leverage existing investments, maximizing their effectiveness without necessitating a complete system overhaul.

Holistic environmental surveillance

By providing visibility into all environments—whether cloud-based, on-premises, or hybrid—Simbian ensures comprehensive surveillance and response capabilities. Our AI agents can operate across varied IT landscapes, offering consistent and reliable coverage that adapts to the specific demands of each environment, thereby fortifying your security posture across the board.

Comprehensive and Cohesive

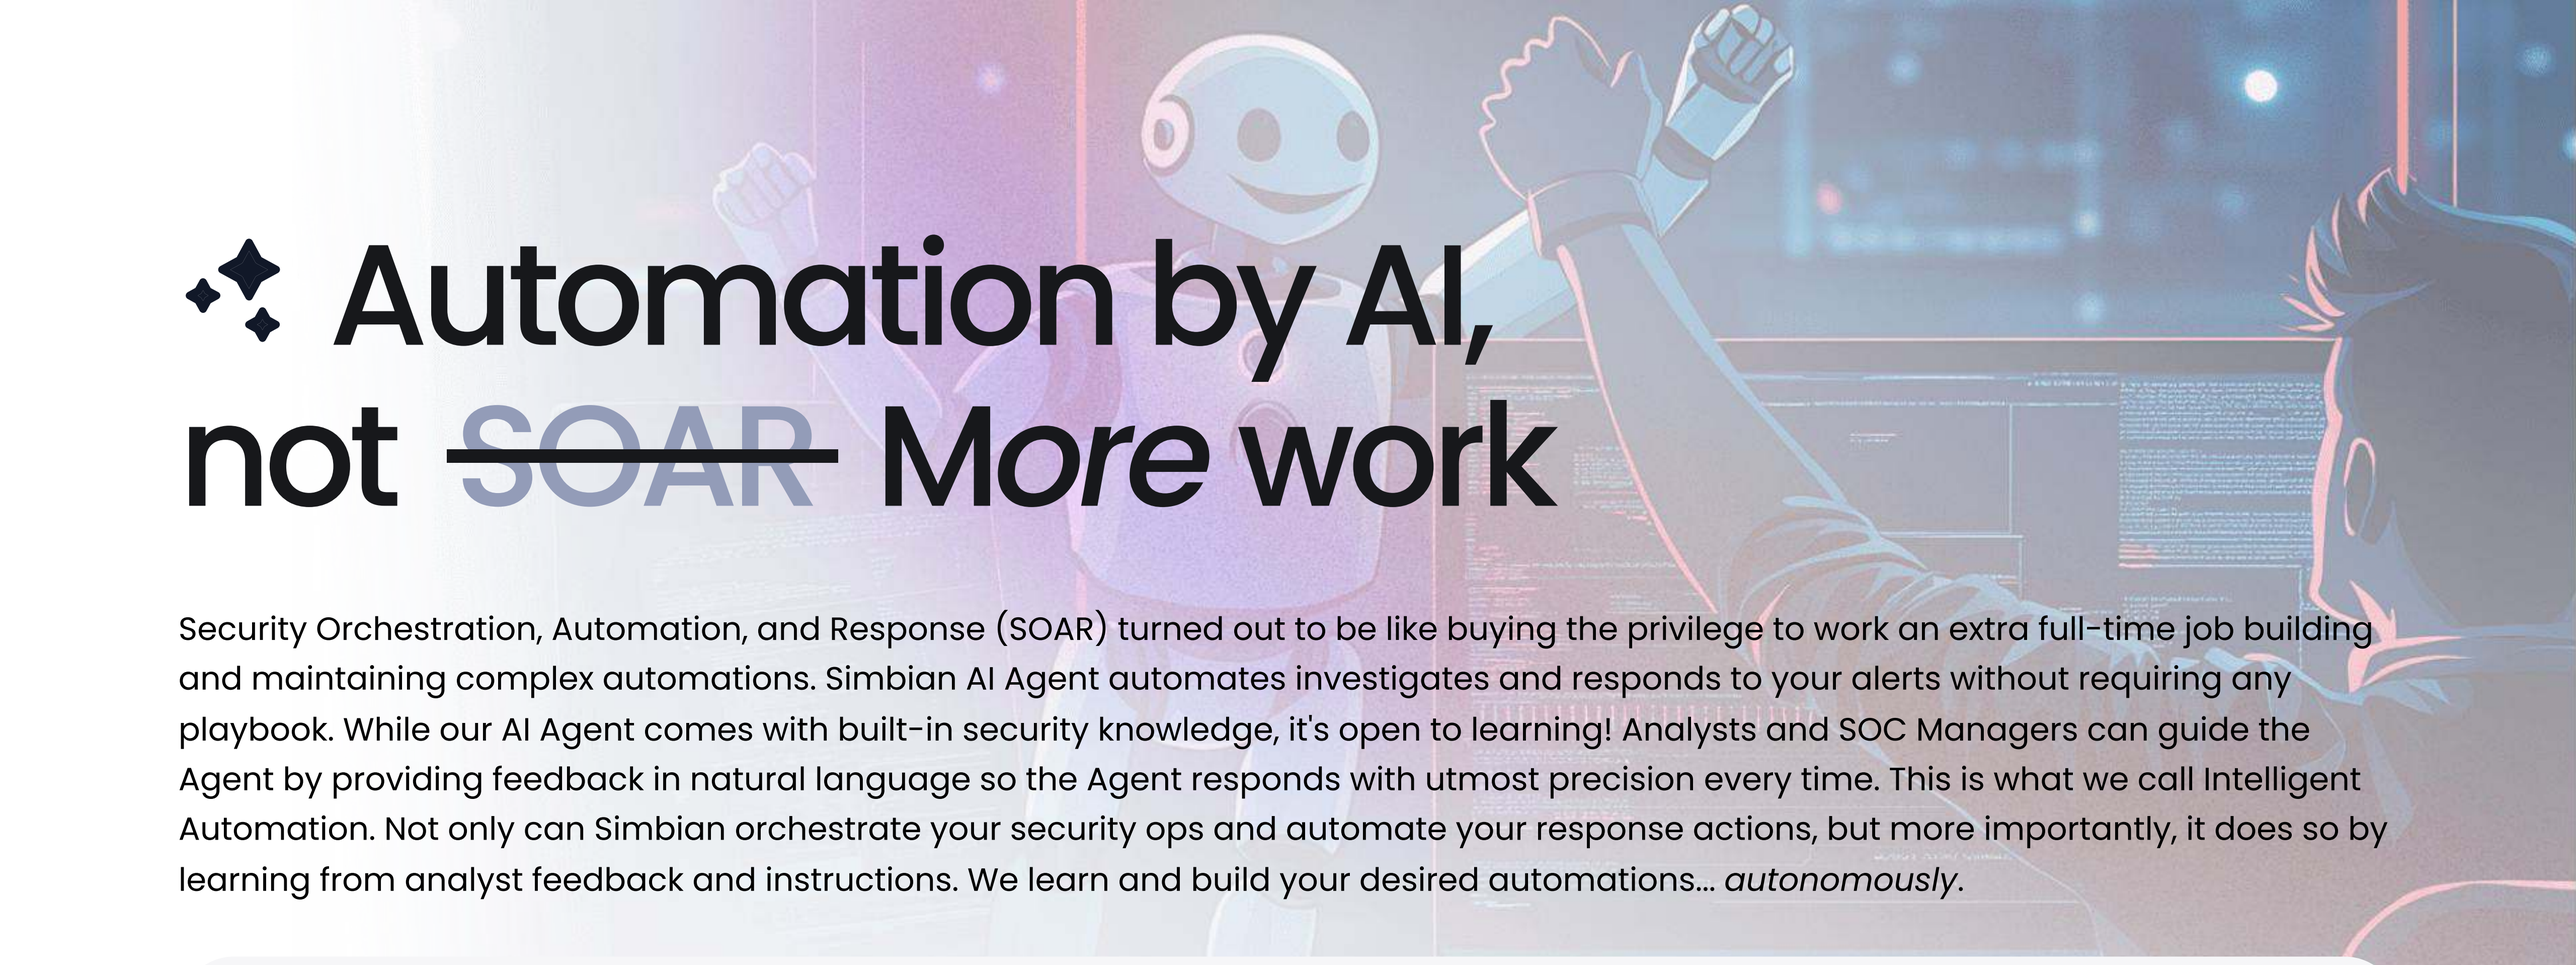

Automation by AI, not ~~SOAR~~ *More* work

Security Orchestration, Automation, and Response (SOAR) turned out to be like buying the privilege to work an extra full-time job building and maintaining complex automations. Simbian AI Agent automates investigates and responds to your alerts without requiring any playbook. While our AI Agent comes with built-in security knowledge, it's open to learning! Analysts and SOC Managers can guide the Agent by providing feedback in natural language so the Agent responds with utmost precision every time. This is what we call Intelligent Automation. Not only can Simbian orchestrate your security ops and automate your response actions, but more importantly, it does so by learning from analyst feedback and instructions. We learn and build your desired automations... *autonomously*.

Ready to start your AI journey with us?

Stable, non-usage-based pricing. On-prem deployment supported.
Schedule your Proof of Value deployment today, for free.

Explore Our Solutions ↗ Book a Demo

Simbian

Our AI models learn and improve over time, delivering increasingly accurate results and adapting to your ever-evolving risk surface and threat landscape.

media@simbian.ai | 809 Cuesta Dr Suite B # 104
Mountain View, CA 94040