securonix

SOLUTION BRIEF

The Securonix Agentic Mesh

Building a Governed, AI-Powered SOC

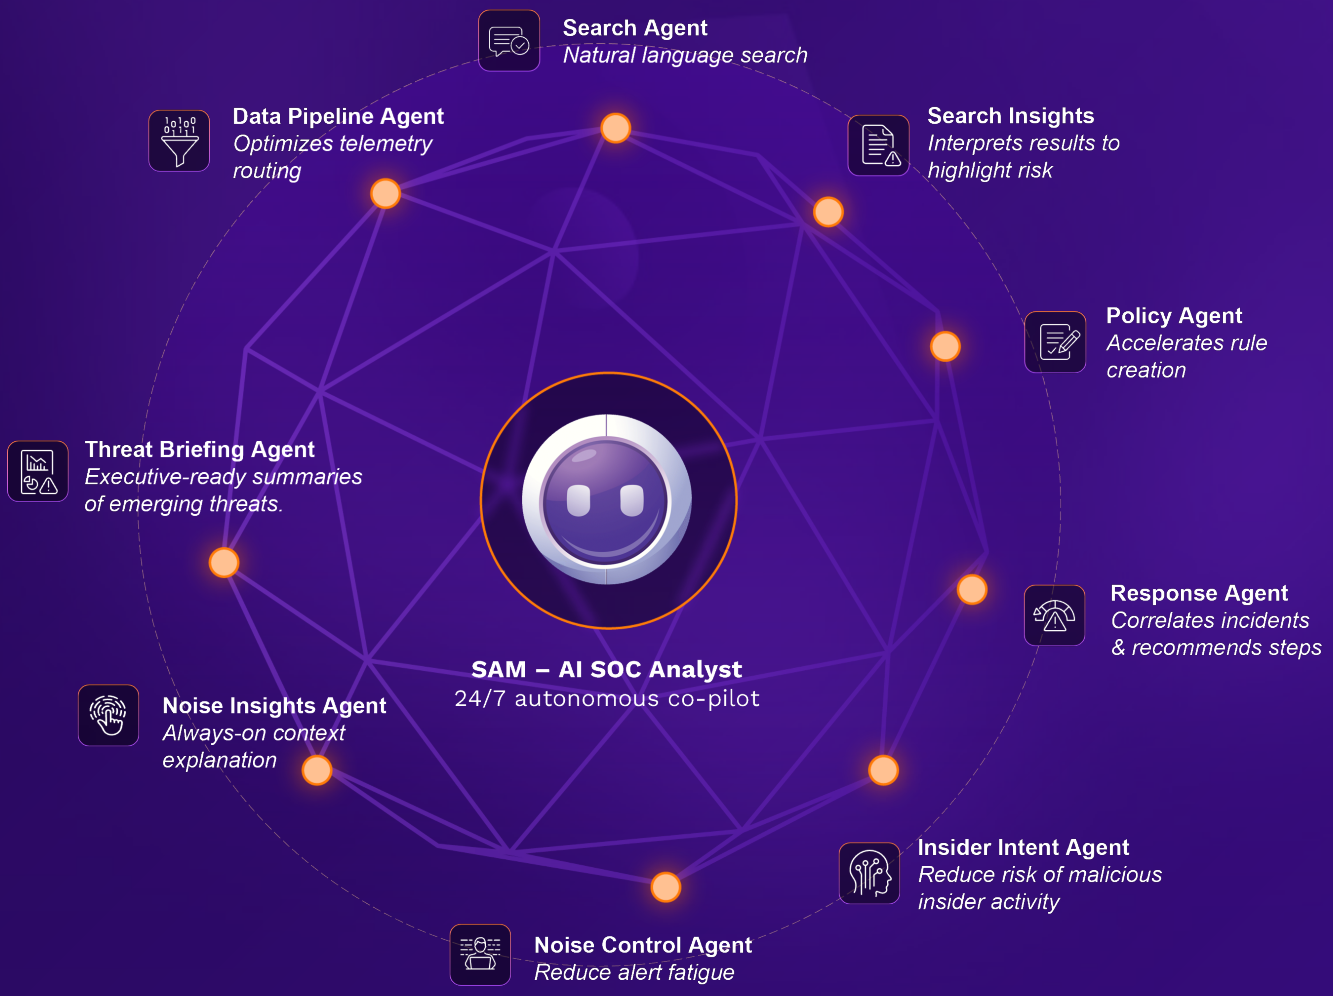

Transforming Security Operations with Trusted, Explainable AI

Modern security operations centers are reaching a breaking point. Every day, analysts face escalating volumes of alerts, fragmented data across multiple clouds, and mounting regulatory oversight. Threat actors are using artificial intelligence to increase speed and sophistication, while security teams still rely on disconnected tools that demand manual investigation and repetitive triage.

At the same time, boards and executives are demanding measurable results. They want faster response, lower cost, and proof that automation is operating safely within compliance boundaries. Yet most AI copilots and automation tools still function as isolated assistants that cannot explain their reasoning or prove accountability.

The result is a trust gap. Security leaders hesitate to scale AI because they cannot verify its actions. Analysts spend more time validating automation than responding to threats. And organizations struggle to balance innovation with compliance.

Securonix Agentic Mesh closes this gap. It delivers the **industry's first productivity-based AI model**, enabling organizations to align cost directly with measurable analyst work performed by **SAM, the AI SOC Analyst**. Instead of paying for data ingestion or seats, customers pay for AI productivity, tracked in real time, explained transparently, and governed by human oversight.

From Automation to Governed Autonomy

The Agentic Mesh represents a new operating model for security. It connects a network of intelligent, explainable AI agents, each designed to think, act, and learn under the supervision of human analysts. Together, they form a collaborative fabric of governed autonomy that accelerates triage, investigation, and response while ensuring every action remains auditable and explainable.

Built on AWS AgentCore, the framework unifies AI agents, data pipelines, and analyst workflows within a single, policy-controlled layer. Every AI-driven action is validated by Agentic Guardrails to meet internal governance, regulatory, and audit requirements, empowering SOCs to scale automation confidently and responsibly.

The Challenge Security Teams Face

1. **Overwhelmed SOCs and Talent Shortages**
   Most organizations cannot hire fast enough to keep up with alert volumes and new attack vectors. Analysts are stretched thin, leading to fatigue and missed threats.
2. **Fragmented Tools and Data Silos**
   Traditional SIEM and SOAR tools operate in isolation, forcing analysts to pivot between dashboards and lose valuable context.
3. **Opaque AI Decisions**
   Early automation and GenAI copilots lack explainability. They act fast but cannot justify their decisions, introducing compliance and governance risks.
4. **Rising Regulatory Pressure**
   Frameworks such as SEC, GDPR, and DORA now require demonstrable oversight and auditability of AI-driven processes.
5. **Demand Proof**
   Boards and regulators no longer ask if an organization uses AI, but how it governs and measures AI in action.

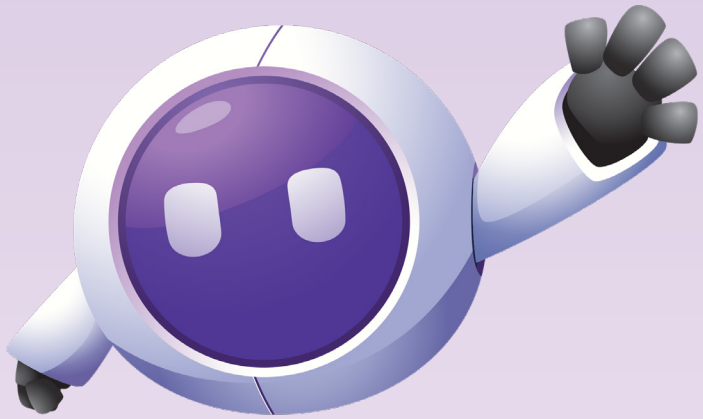

SAM: The AI SOC Analyst

At the center of the Agentic Mesh is SAM, the world's first productivity-based AI SOC Analyst. SAM acts as a trusted teammate, automating Tier 1 and Tier 2 investigations, correlating alerts, and preparing recommendations that analysts can review, refine, and approve via the Agent Desk.

SAM explains its reasoning step-by-step, blending institutional knowledge with continuous learning. Analysts can challenge findings, pause actions, or request deeper insights, maintaining full control while achieving measurable gains in speed and accuracy.

Measured outcomes include:

- **Up to 60%** reduction in false positives
- **50% faster** investigations
- **3x** increase in analyst productivity
- **Verifiable ROI** through tracked AI contribution

AI Agents Working in Concert

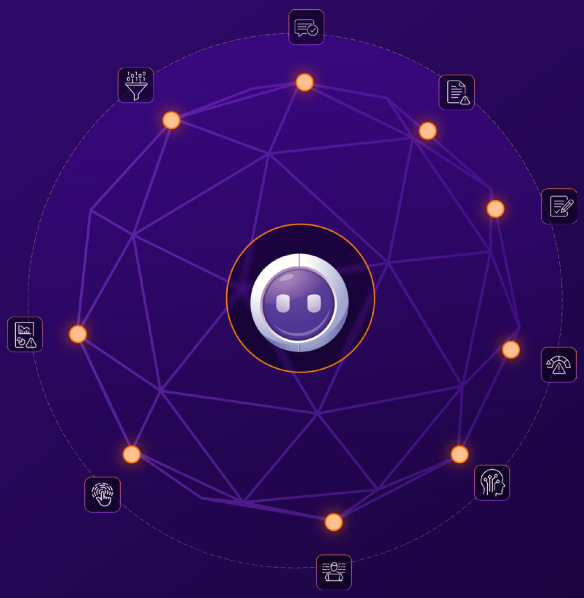

Each modular agent within the Securonix Agentic Mesh contributes to faster, safer, and more accountable SecOps:

SAM - AI SOC Analyst

Autonomous digital teammate that automates Tier 1 and Tier 2 investigations. It triages alerts, correlates context across SIEM, UEBA, SOAR, and TIP, and produces explainable summaries for analyst validation.

Search Agent

Performs natural-language search and data exploration. Translates plain-language questions into optimized queries across security datasets to speed-up threat hunting and root cause analysis.

Search Insights Agent

Transforms raw search results into contextual narrative summaries. Highlights key entities, anomalies, and risks while generating findings for analysts or executives.

Noise Control Agent

Reduces alert fatigue by filtering duplicate, redundant, and low-confidence alerts. Learns from analyst feedback to refine accuracy and ensure focus on high-value incidents.

Response Agent

Automates triage and containment actions with explainable reasoning. Integrates directly with SOAR workflows and ensures every recommendation aligns with organizational policy.

Data Pipeline Agent

Optimizes data ingestion, retention, and cost through AI-driven tiering and DPM Flex pricing. Aligns data value to cost and provides real-time visibility into usage and efficiency.

Policy Agent

Automates the creation and enforcement of AI policy rules. Defines risk thresholds, manages approval of workflows, and ensures AI actions comply with internal and regulatory requirements.

Insider Intent Agent

Detects early signs of insider risk by analyzing behavioral drift, psycholinguistic signals, and identity activity. Correlates patterns to predict and prevent data loss or fraud.

Investigation Agent

Correlates evidence, enriches findings, and produces explainable incident reports. Integrates internal and external intelligence for a comprehensive investigative context.

Threat Intel Agent

Aggregates and normalizes multiple intelligence feeds to enhance detection accuracy and enrich alert context. Continuously updates indicators and signatures.

Productivity-Based Pricing: Aligning Value to Outcomes

Traditional pricing models tie cost to data volume or seats, creating friction between efficiency and spend. The Securonix productivity-based AI model shifts this paradigm. Customers receive a baseline SAM capacity, with every minute of AI-driven work tracked in a transparent productivity dashboard. When AI value grows, capacity scales proportionally, ensuring customers pay only for verified results.

This pricing innovation gives organizations financial flexibility while reinforcing accountability and measurable ROI across the SOC.

Built for Leaders Who Demand Accountability

In a market crowded with AI promises, Securonix stands apart by delivering measurable, explainable, and trusted results. The Agentic Mesh gives organizations the confidence to modernize their SOC without sacrificing control or compliance. It combines governed autonomy, collaborative intelligence, and scalable performance to create an operational framework where every AI-driven decision is transparent, auditable, and aligned with business outcomes.

Governed AI Architecture

Securonix is the only platform that operationalizes AI within an explainable, policy-controlled framework. Every action is validated against human and regulatory standards.

Open and Scalable Foundation

Unlike closed ecosystems that limit data interoperability, the Agentic Mesh is built on open, cloud-native infrastructure integrated with AWS and Snowflake for performance and flexibility.

Human-in-the-Loop Design

Every process keeps analysts in command. Automation supports decisions but never bypasses oversight.

Proven Performance

Securonix customers have achieved a 193% ROI verified by the Forrester TEI study, with up to 70% faster MTTR and 90% fewer false positives.

Recognized Leadership

Gartner has named Securonix a 6x Leader in the Magic Quadrant for SIEM, with strengths in open data architecture and AI-driven governance.

AI Workforce Enablement

Through the Agentic AI Academy, Securonix builds analyst fluency in governed autonomy, closing the talent gap that limits AI adoption across enterprises.

The Future of Security Operations

With the Securonix Agentic Mesh, organizations move beyond reactive automation to **governed, outcome-driven autonomy**. For analysts, it means less noise and more impact. For SOC leaders, it means visibility and verifiable ROI. For executives and boards, it means trust in AI that is accountable, measurable, and compliant by design.

That is what it means to be **Breach Ready. Board Ready. AI Powered.**

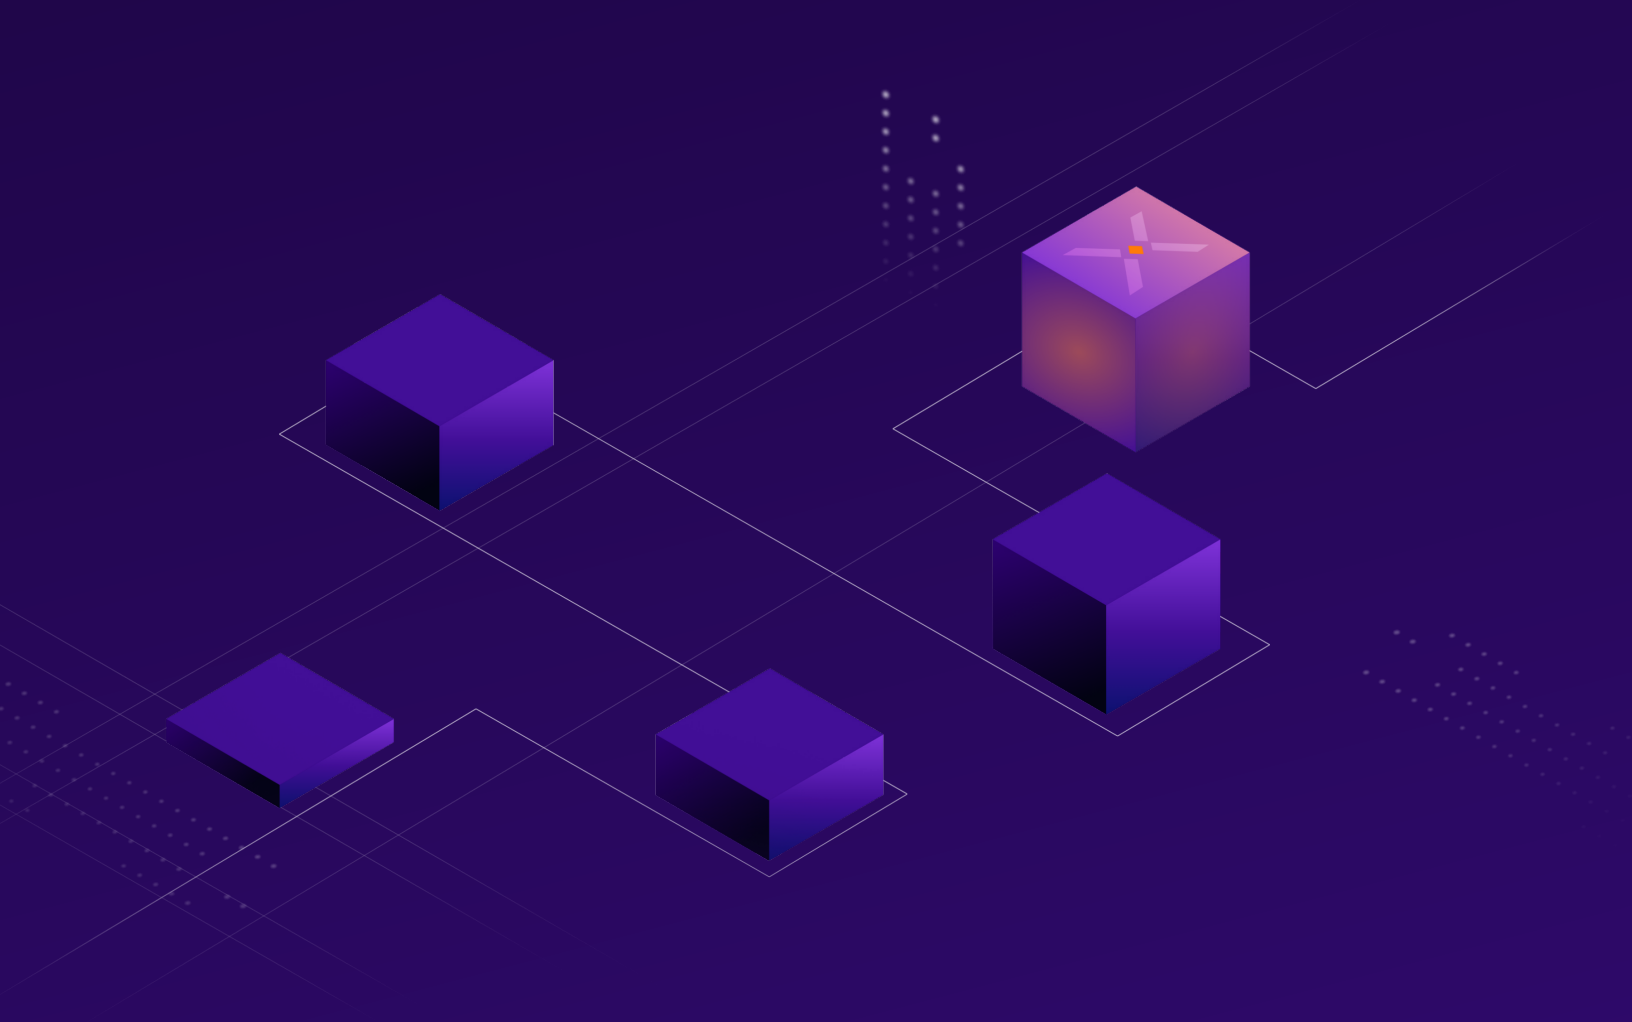

About Securonix

Securonix is building the future of security operations, not by layering on more tools, but by simplifying and unifying the ones that matter. Our vision is a security platform that works the way your team does: modular, intelligent, and outcome-driven. We believe security should be built for speed, designed for scale, and proven by results. That's why we measure success not in alerts processed, but in risks mitigated, time saved, costs reduced and trust earned. For more information visit **securonix.com**