seceon
SECURITY FOR EONS

Seceon SERA AI AutoSOC

Agentic AI SOC: Evidence-Grounded Investigation

Autonomous alert investigation, deterministic verdicts and guarded SOAR response, native to the Seceon OTM Platform.

Product Overview

Seceon SERA AI AutoSOC is the agentic investigation layer of the Seceon OTM Platform. It investigates Major and Critical alerts end to end, reaches a scored verdict, and hands confirmed threats to guarded SOAR response.

Related alerts are clustered into one case. Domain committees for network, identity, endpoint, and email/cloud then set a prosecutor agent against a defender agent for up to 12 investigative turns, using 14 native tools that query the customer's own OpenSearch and Cassandra stores. An adjudicator and chief analyst issue the verdict.

Models advise; code decides. Verdicts are scored deterministically from an Observed, Inferred, Speculative evidence hierarchy, and actions run only through SOAR playbooks with hardcoded guardrails, on premises, in the cloud or fully air-gapped



Highlights

- **Adversarial multi-agent investigation:** prosecutor, defender, adjudicator and chief analyst committees per domain
- **Deterministic, evidence-based verdicts:** the same evidence produces the same score on every run
- **No silent closures:** cases with missing data stay open as degraded, name the gap and recommend the fix
- **Guarded response:** 100+ SOAR playbooks, sub-90-second automated response and 70% of incidents resolved without analysts
- **Sovereign by design:** full feature parity on premises and air-gapped, with local inference and no required outbound traffic
- **Built-in AI safety:** aiTRiSM360 redaction, prompt-injection defense and a tamper evident audit ledger

Deployment Coverage

Deployment Capability	Description
On-Premises	Customer-owned appliances, local model gateway, and locally hosted models
Air-Gapped	Full feature parity, no required outbound traffic, and data-diode readiness
Sovereign / Private Cloud	Supports OpenStack, VMware, Nutanix, OpenShift, and private Kubernetes
Public and Government Cloud	Supports AWS, Azure, Google Cloud, and OCI, including AWS GovCloud and Azure Government
Hybrid / MSSP	Distributed CCE collectors, central APE, and multi-tier multi-tenancy
SIEM Overlay	Runs alongside Splunk, Sentinel, or QRadar and returns adjudicated incidents

Investigation & Verdict Pipeline

Stage	What it does
Normalization (SEF)	Standardizes logs, flows, identity, cloud, and OT data into one event format, supported by 3,000+ connectors
Pre-Correlation (APE)	Correlates events at ingestion so alerts arrive with related activity attached
Case Clustering	Groups Major and Critical alerts by entity, context, and time
Evidence Gathering	Retrieves indicators from local OpenSearch and Cassandra without requiring model calls
Adaptive Fan-out	Investigation lead activates only the domain committees required for the case
Adversarial Debate	Prosecutor and defender agents debate evidence across up to 12 query turns
Adjudication	Adjudicator evaluates the debate; only the chief analyst issues the final verdict
Deterministic Scoring	Calculates scores in code using Observed, Inferred, and Speculative evidence categories
Exposure Weighting	Incorporates host vulnerability posture, including CISA KEV, into the verdict
Gap Handling	Flags missing data, marks the case as degraded, and recommends corrective action
Guarded Response	Executes SOAR playbooks for containment within hardcoded guardrails
Closed-Loop Verification	Re-queries affected entities and uses aiBAS360 to re-test remediated controls

Evidence & Grounding

- Semantic data layer, not RAG: agents run precise queries against structured SEF data
- 14 native tools: events, flows, prevalence, prior closures, assets, auth history, process trees, TI and vulns
- Fresh evidence read from current platform data at investigation time
- Analyst dispositions and tenant-scoped memory improve verdicts without changing model weights
- 100+ threat intelligence sources, with attribution to adversary profiles such as LockBit

Response, Guardrails & Verification

- 100+ prebuilt SOAR playbooks for firewalls, identity, endpoints, email, and cloud.
- Host isolation, process termination, C2 blocking, account disablement, and tenant-wide email removal.
- System and machine accounts are never disabled; domain controllers are exempt from automatic isolation.
- Medical devices and SCADA lines are routed for human approval.
- No-code playbook editor with natural-language drafting and optional approval gates.
- Checkpointed workflows resume after a restart without repeating destructive actions.

AI Safety & Platform

- Stateless agents; memory and rules are stored in a tenant-scoped store outside the agent runtime.
- Schema-bound JSON output; models never calculate scores or execute actions.
- aiTRiSM360 redaction and prompt-injection defense applied before every model call.
- 27 MITRE ATLAS-mapped rules covering prompt injection, jailbreaks, and MCP tool poisoning.
- Detection, correlation, and SOAR continue operating even if SERA AI is unreachable or switched off.
- SHA-256 tamper-evident ledger; every query, index, and cache key is tenant-scoped.



Primary Use Cases

Use Case	How SERA AI AutoSOC Helps
Cut alert fatigue	93.9% fewer false alarms in 15 days at a 209-tenant MSSP partner
Autonomous L1/L2 triage	Agents investigate and score cases; routine benign cases need no model call
Faster containment	Guarded SOAR playbooks with sub-90-second automated response
Ransomware before encryption	Multi-stage threat models, canary files and shadow-copy deletion alerts
Identity and non-human identity attacks	Kerberoasting, DCSync and token abuse with UEBA and identity context
Exposure-aware prioritization	Alerts on hosts with exploitable vulnerabilities upgraded automatically
Sovereign and classified SOCs	Local inference and air-gapped parity for defense and government
AI SOC on an existing SIEM	Overlay on Splunk, Sentinel or QRadar without rip-and-replace
MSSP scale	Per-tenant isolation, tuning APIs and multi-tier multi-tenancy

Why Seceon SERA AI AutoSOC

Differentiator	Benefit
Debate, not a single prompt	Prosecutor and defender agents test both explanations before a verdict
Code scores, models advise	Repeatable verdicts that a model cannot grade for itself
No silent closures	Missing data yields a degraded case, never a false benign
OTM Platform native	Runs on CCE, APE and SEF with pre-correlated context, not API stitching
Sovereign AI	Full parity on premises and air-gapped with local models
Safe automation	Guardrails protect domain controllers, system accounts and OT assets
Predictable AI cost	No AI usage fees on local models; benign closures use no model calls

Deployment Snapshot

Item	Detail
Release baseline	Seceon OTM Platform 12.3.3 (GA) with SERA AutoSOC
Form factor	Native OTM module; virtual appliances as OVA, AMI and VHD
Operating system	Seceon Hardened Linux 9.x, FIPS 140-3 validated crypto, TLS 1.3
Isolation boundaries	On-premises tool harness · signed client · stateless model gateway
Data stores	Customer OpenSearch and Cassandra; raw telemetry stays in APE
High availability	Standby-aware clustering, replicated stores, redundant collectors
Time to value	Installation in about 5 hours; full functionality within 2 weeks

Inference options

Option	Where models run	Data path	AI cost
Local models	On premises or air-gapped	Inside customer boundary	No AI usage fees
Hosted models	Stateless HTTPS gateway	Redacted prompts only	Prepaid credits with quotas
AI switched off	None	Detection, SOAR continue	None

Ideal Buyers

- Enterprise SOCs facing alert volume and analyst shortages that want autonomous triage
- MSSPs and MSPs scaling managed detection and response without adding analysts per tenant
- Government, defense, critical infrastructure, and regulated industries requiring local AI
- Teams on Splunk, Sentinel or QRadar that want an AI SOC layer without replacing the SIEM

One-Line Pitch

Seceon SERA AI AutoSOC investigates every critical alert with adversarial AI agents, scores verdicts deterministically from evidence, and executes a guarded response on-premises, in the cloud, or fully air-gapped.

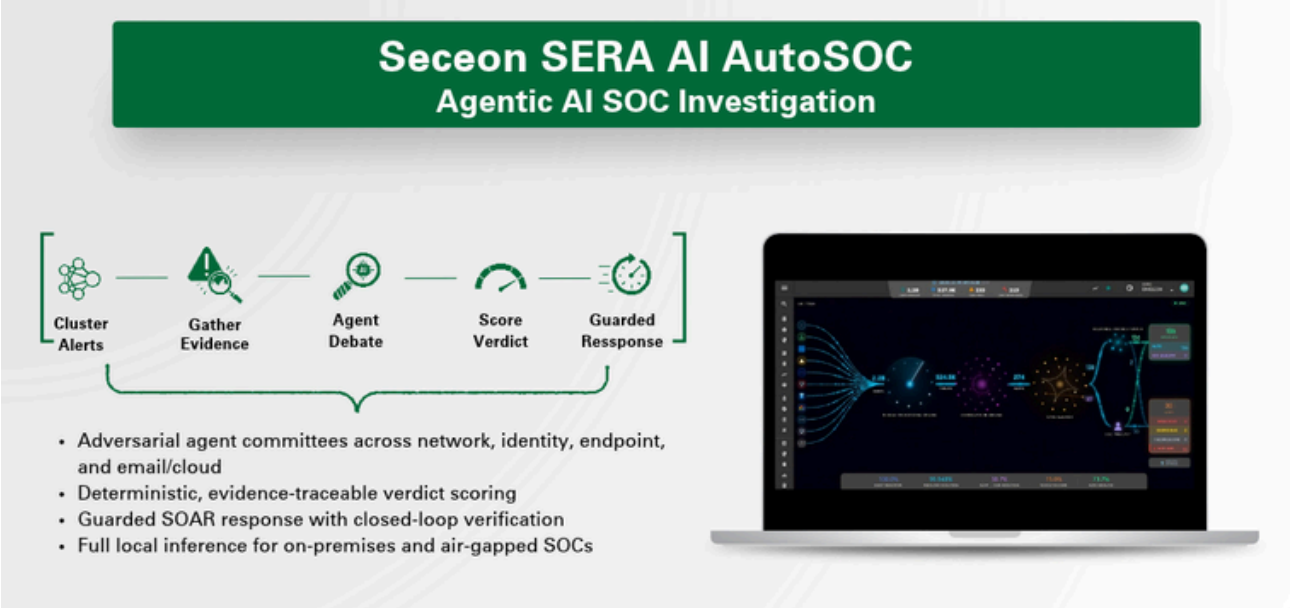

Seceon SERA AI AutoSOC investigates critical alerts with adversarial AI agents, scores verdicts from evidence, and executes guarded response, so analysts focus on decisions instead of triage.

About Seceon

Seceon simplifies cybersecurity for MSPs, MSSPs, and IT teams by reducing risk and complexity. Our AI and ML-powered aiSIEM and aiXDR platform provides comprehensive threat detection and response. By integrating data from logs, identity management, networks, endpoints, clouds, and applications, we deliver real-time, accurate alerts and automated remediation. Trusted by over 850 partners, Seceon supports more than 9,800 customers with efficient, high-margin security services and continuous compliance.

Learn more about Seceon

Schedule a Demo www.seceon.com/contact-us/



© 2026 Seceon Inc. All rights reserved. aiSIEM and aiXDR are trademarks and registered trademarks of Seceon Inc. in the United States and other countries. All other brand names, product names, or trademarks belong to their respective owners.