ANOMALI

One-pager

A PRACTICAL GUIDE

WHAT IS AN AGENTIC SOC PLATFORM?

Beyond Detecting. Decide. Act.

Gartner predicts that 40% of agentic AI projects will be cancelled in 2027, with their infrastructure and observability research showing that these failures cluster in auto-remediation and agent-led workflows. Over one-third (38%) of these projects fail because of data readiness problems. That data problem is becoming a security problem. AI is only as good as the data it operates on. Fragmented, ungoverned, or context-poor data amplifies noise, risk, and cost. Clean, organized, governed data is foundational to an agentic SOC.

WHY THE SOC NEEDS TO CHANGE

Most security operations centers already have visibility. What they lack is a way to turn that visibility into fast, governed decisions. Data volumes keep growing, adversaries move at machine speed, and teams are asked to show measurable outcomes on budgets and headcount that stay flat.

The deeper issue is structural. The legacy SIEM-centric SOC was built for log retention and alerting, not for decisions. Next-generation platforms promised AI and delivered more alerts instead of better outcomes. Analysts triage by instinct while ingest-based pricing strains budgets and hybrid, multi-cloud environments stretch teams past their limits.

Threat intelligence platforms were designed for research, not operational execution, and automation applied without context tends to add noise rather than reduce risk. The result is a SOC that stays reactive, costly, and inconsistent.

Security operations now need an intelligence-native architecture: one that applies context at ingest, builds an AI-ready data foundation, and drives governed response at machine speed with human oversight.

The Anomali Agentic SOC Platform brings together a high-fidelity security data lake, Managed Intelligence as a Service, powered by ThreatStream Next-Gen. and true agentic AI in a single operational experience. Instead of producing more alerts and dashboards, it turns telemetry and intelligence into guided, high-confidence decisions.

This guide explains what an agentic SOC platform is, how it works in practice, and how Anomali's core components operate together to support modern SOC and CTI teams.

WHAT IS AN AGENTIC SOC PLATFORM?

An agentic SOC platform is an operational security platform where AI-driven agents work alongside analysts, reasoning across security data and threat intelligence to support detection, investigation, and response.

Unlike traditional analytics or automation tools, it operates at the decision layer. Rather than generating alerts or running predefined playbooks, it continuously analyzes telemetry, enriched threat intelligence, and historical context to:

- Deliver intelligence-native detections at ingest
- Reduce alert noise and prioritize what matters
- Provide actionable investigative context
- Guide analysts toward effective response decisions
- Keep intelligence consistently operationalized across teams

Analysts stay in control. Agentic AI provides speed, context, and consistency, and every autonomous action runs inside governance the SOC defines. That is the difference between a platform built to detect and one built to decide.

THE SECURITY DECISION BRAIN: THREE CAPABILITIES, ONE GOVERNED SYSTEM

Most SOCs already run one or two of these capabilities. Few run all three as a single system, and that gap is where investigations stall and decisions drift. Telemetry, context, and action have to operate together, under one set of controls. Anomali calls that system the Security Decision Brain.

See Everything: Live security telemetry and signal across endpoint, network, cloud, identity, IT/OT, email, and SaaS, normalized, correlated, and queryable in real time.

Know What Matters: Threat graph, asset inventory, user baselines, topology, and incident history fused with your telemetry. Raw signal tells you what happened. Context tells you what it means.

Act With Confidence: Agentic AI brings agency to the SOC. Instead of static dashboards or chat-only copilots, agents reason over your data lake and intelligence context to guide investigations, recommend next actions and automate response workflows — fewer steps, faster decisions and consistent execution.

THE ANOMALI AGENTIC SOC PLATFORM

The Security Decision Brain runs on three integrated components. Each solves a specific operational problem, and the value compounds when they work together. Raw data without intelligence is noise. Intelligence without data is reporting.

1. UNIFIED SECURITY DATA LAKE

Every SOC operation depends on data. If telemetry is incomplete, delayed, or capped by indexing limits, investigations stall and decisions degrade.

The Anomali Unified Security Data Lake is a modern, AI-ready, hyperscale security data architecture purpose-built for intelligence-driven operations. Unlike legacy SIEM backends, it augments or replaces existing SIEM investments and removes ingestion bottlenecks and unpredictable cost. It provides:

- Centralized ingestion of security telemetry across cloud, endpoint, network, identity, IT/OT, and application sources
- Security-native normalization and correlation at ingest
- Always-searchable, investigation-ready data with no retention tradeoffs
- High-speed search and analytics across months or years of telemetry

Because telemetry is normalized and correlated as it enters the platform, everything downstream operates on clean, consistent context. Analysts and agentic AI pivot quickly across alerts, entities, and timelines with no blind spots. This is not cold storage. It is always-on and always-searchable, built for real-time and historical analysis.

2. MANAGED INTELLIGENCE AS A SERVICE (MIAAS), POWERED BY THREATSTREAM NEXT-GEN

Telemetry alone does not explain risk. Knowing that something happened is not the same as knowing who is behind it, why it matters, or what comes next.

WHAT IT IS	WHAT IT DOES
Unified Security Data Lake	Full-fidelity telemetry across cloud, endpoint, network, identity, IT/OT, and beyond. Always-on and always-searchable, with no legacy SIEM cost or performance ceiling.
Managed Intelligence as a Service (MIaaS), powered by ThreatStream Next-Gen	Real-time context that an agentic architecture depends on. Continuously enriches the data lake with threat actors, TTPs, and campaigns, operationalizing intelligence in minutes rather than business days.
Agentic AI	A stack-ranked decision queue, MCP-enabled agentic operations, and Tier 1 and Tier 2 triage agents that reflect your SOC's judgment rather than fixed vendor runbooks.

Managed Intelligence as a Service (MIaaS), powered by ThreatStream Next-Gen, is the intelligence layer of the agentic SOC and the safety layer the rest of the agentic architecture depends on. It continuously enriches your data lake with curated, confidence-scored threat intelligence:

- Threat actors, campaigns, and infrastructure
- Tactics, techniques, and procedures mapped to MITRE ATT&CK
- Behavioral and indicator-of-attack context
- Asset-based context tied to organizational risk

Instead of static intelligence, it's correlated with internal telemetry so context travels with every alert and investigation, and it operationalizes in minutes rather than business days. SOC and CTI teams prioritize on adversary intent and operational relevance instead of raw indicator volume.

3. AGENTIC AI

Agentic AI sits above the data and intelligence and points where decisions are made. Capabilities include:

- Automated enrichment and investigation of alerts
- Behavioral and IOA-based analytics focused on attacker intent
- Intelligence-ready guidance that travels with alerts and cases
- Context-aware response and mitigation recommendations
- Semantic search and natural-language interaction grounded in security context

Anomali's AI-driven agents reason across the Unified Security Data Lake and MIaaS intelligence to support investigation and response. They reflect your SOC's judgment rather than fixed vendor runbooks, and they reduce manual effort and inconsistency instead of replacing analysts. Analysts spend less time correlating data and more time applying judgment and acting.

HOW AGENTIC SOC OPERATIONS WORK

An agentic SOC platform is best understood through how it runs. Here is how detection, investigation, and response work together in the Anomali Agentic SOC Platform.

DETECTION: FROM SIGNAL TO MEANING

Detection begins with telemetry flowing into the Unified Security Data Lake. Because data is normalized and enriched in near real time, detections are built on complete, unmodified data rather than sampled or delayed logs.

MIaaS continuously enriches that telemetry with real-world threat intelligence, and agentic AI evaluates alerts against historical patterns, behavioral context, and adversary intelligence to suppress low-value signals and surface high-confidence threats.

INVESTIGATION: CONTEXT WITHOUT MANUAL CORRELATION

When an alert is escalated, agentic AI correlates automatically:

- Related entities and historical activity
- Threat actors, campaigns, and infrastructure
- Behavioral indicators aligned to attacker TTPs

Analysts pivot across indicators, timelines, and threat models using semantic search and investigation workbenches. Knowledge graphs bridge natural-language queries and structured security data, so multi-hop reasoning does not require complex query construction.

RESPONSE: GUIDED, GOVERNED ACTION

Agentic AI provides context-aware response recommendations shaped by the specific threat scenario and environment, informed by intelligence, telemetry, and prior outcomes. That reduces guesswork and variability. Because intelligence and context are embedded throughout, response is carried out consistently across SOC, CTI, and downstream tools such as SOAR, EDR, XDR, and ticketing systems. Response becomes consistent, intelligence-informed, governed, and measurable.

PROOF POINTS

Organizations adopt an agentic SOC platform to change operational outcomes, not to make incremental tooling improvements. Anomali targets:

METRIC	OUTCOME
30-50%	Total cost of ownership reduction
90%	Reduction in critical incidents
60-70%	Analyst time reclaimed from false-positive triage

WHY ANOMALI IS DIFFERENT

The security problem has become a data problem. Context, intelligence, and governed autonomy have to operate on one unified data layer.

- One consolidated data lake: a decade of curated intelligence combined with human and non-human identity topology.
- Operationalized managed intelligence: govern context-driven decisions across your entire security ecosystem.
- A unified platform with no concentration risk: vendor-neutral and architecture-agnostic, so it works with the tools you already run.

HOW THE PLATFORM WORKS AS A SYSTEM

The defining characteristic of an agentic SOC platform is how its components work together:

- The Unified Security Data Lake keeps telemetry complete, accessible, and normalized.
- MIaaS, powered by ThreatStream Next-Gen, turns intelligence into operational context.
- Agentic AI reasons across both layers to guide detection, investigation, and response.

This integrated approach reduces redundant tools and manual handoffs. Intelligence is no longer separate from operations, and AI is not bolted onto existing systems. The platform functions as a single operational plane where data, intelligence, and decisions stay connected.

OPERATIONALIZING AGENTIC AI ACROSS SECURITY OPERATIONS

An agentic SOC platform is a shift in how security operations are designed. Instead of optimizing individual tools, it enables better decisions and turns the SOC from a reactive alert center into an intelligence-native, AI-accelerated decision engine.

The Anomali Agentic SOC Platform brings together always-on telemetry, continuously enriched threat intelligence, and agentic AI so organizations can:

- Modernize without disruptive rip-and-replace
- Augment existing SIEM investments
- Reduce cost and complexity
- Prepare for AI-driven adversaries
- Run operations across hybrid, multi-cloud environments

For teams moving beyond alert-centric operations, the platform provides a practical path from information to action. Teams typically focus on:

- Faster detection and triage through intelligence-driven prioritization
- Shorter investigation cycles with full historical and adversary context
- Reduced alert fatigue and analyst burnout
- More consistent, repeatable security decisions across teams
- Stronger alignment between SOC and CTI operations

These outcomes come from the platform's ability to move from raw data to guided action rather than stopping at alerts or analytics.

SEE EVERYTHING.
KNOW WHAT MATTERS.
ACT WITH CONFIDENCE.

SEE ANOMALI IN ACTION

Request a Demo