ANOMALI

Guide

WHY DATA FIDELITY DECIDES WHAT YOUR AGENTIC SOC CAN DO

A guide for Head of SOC and CISO teams evaluating agentic security operations

An AI agent investigating an alert doesn't run one query and stop. It explores the data, forms a hypothesis, tests it, and looks again, often reaching back across years of telemetry in seconds. That behavior exposes the real constraint on agentic security operations, but it's not about the model quality or compute. What decides whether an agent reaches a correct conclusion is the data underneath it: whether that data is complete, consistent, enriched, and traceable to its source.

Most security teams have spent a decade getting good at detection. The harder problem now is decisioning, turning a signal into a defensible action fast enough to matter. Agentic AI can close that gap, but only when it reasons over high-fidelity data it can trust. This guide covers what data fidelity means in practice, why enrichment at ingest returns hours to analysts, and how Anomali's architecture is built to support a SOC that makes confident decisions.

THE DECISIONING GAP

The old SOC playbook treated alert volume as a measure of coverage. Point tools handled each function, analysts triaged by instinct, and threat intelligence sat in a separate system from operations. Detection improved, but the integration burden moved onto people. Correlation, investigation, and response still ran at human speed.

AI agents change where the stack converges. With agent-native APIs and standards like MCP, agents coordinate directly across security systems rather than through dashboards and manual handoffs. Security starts to look less like a sequence of tickets and more like continuous execution.

> **There is a catch.**
> An agent inherits the quality of the data it reads. Give it fragmented, sampled, or unenriched data and it will produce fast, confident, wrong answers. When that output drives an automated response, speed without fidelity becomes a liability rather than an advantage.

FIDELITY IS FOUR PROPERTIES, NOT A PILE OF DATA

Collecting more logs doesn't make a SOC smarter. A larger pile of low-quality data simply gives an agent more ways to be wrong. Fidelity describes whether the data an agent reasons over is actually fit for a decision, and it comes down to four properties.

COMPLETENESS

Full-fidelity telemetry across endpoint, network, cloud, identity, IT/OT, email, and SaaS, retained long enough to matter. Sampled logs and 30-to-90-day retention windows create blind spots an agent cannot reason around. When nothing ages out, an investigation can reach back years without hitting a wall.

CONSISTENCY

One schema, applied at ingest. Anomali normalizes telemetry to OCSF and de-duplicates it, so a query returns the same entity the same way regardless of which tool produced the event. Fragmented schemas force agents and analysts to reconcile formats before they can reason, and that reconciliation step is where errors enter.

ENRICHMENT QUALITY

Raw signal tells you what happened. Context tells you what it means. Anomali correlates curated, confidence-scored intelligence with your telemetry as it lands, including threat actors, campaigns, infrastructure, and TTPs mapped to MITRE ATT&CK, so who, why, and what-next travel with the alert instead of waiting in a separate report.

PROVENANCE

Every enriched record carries its origin and its confidence score. That traceability is what lets a decision be explained after the fact and reversed if it was wrong. For a CISO, provenance is the line between an autonomous action you can defend in an audit and one you cannot.

These properties compound. Completeness gives normalization something worth normalizing, consistency gives enrichment a clean substrate, and enrichment is only trustworthy when its provenance holds up.

HOW INTELLIGENCE AT INGEST RETURNS TIME TO ANALYSTS

The most expensive minutes in a SOC are the ones analysts spend rebuilding context that should have been there already. When threat intelligence lives in a platform separate from operations, every investigation starts with manual correlation: pull the indicator, check the feed, map it to the actor, judge whether it matters. Repeat across the queue and it becomes the workday.

A threat intelligence platform that enriches telemetry at ingest changes the starting point. Because the Intelligent Unification Layer correlates intelligence with events as they arrive, context is already attached when an analyst opens the alert. Intelligence that used to live in reports is operationalized as detection and enforced controls, often within minutes of arriving rather than after a manual review cycle. Anomali reports that analysts reclaim a large share of the time previously lost to false-positive triage; teams should validate that against their own baseline before treating it as a target.

The time returned matters most where skilled analysts are scarce. Hours move away from correlation, which the platform handles well, toward investigation and response decisions, which stay with people.

WHERE ANOMALI IS DIFFERENT

For a team comparing approaches, a few architectural choices set Anomali apart.

YOU ENTER WITHOUT DISPLACEMENT

Deploy on top of your existing SIEM and infrastructure, managed or bring-your-own on Databricks, Snowflake, Sentinel, or S3. Enrich your data without a rip-and-replace project. Every deployment can expand into the full Agentic SOC Platform on your own timeline.

TWO ASSETS NO COMPETING DATA LAKE CARRIES

The threat graph is a decade of operationally curated intelligence that enriches every query, detection, and investigation natively rather than as a bolt-on. The identity and entitlement topology is a quantitative model of every human and non-human identity in your environment. Together they form the substrate that makes identity-grounded governance and blast-radius calculation possible.

INTELLIGENCE THAT IS BUILT TO ACT

Anomali operates one of the industry's largest threat intelligence repositories, fused with your telemetry and scored for confidence rather than delivered as raw indicator volume. Through Trusted Circles, peers in your sector can share IOCs, TTPs, and campaigns bi-directionally, so an early warning to one member becomes automatic detection for the others, with TLP gating enforced at ingestion.

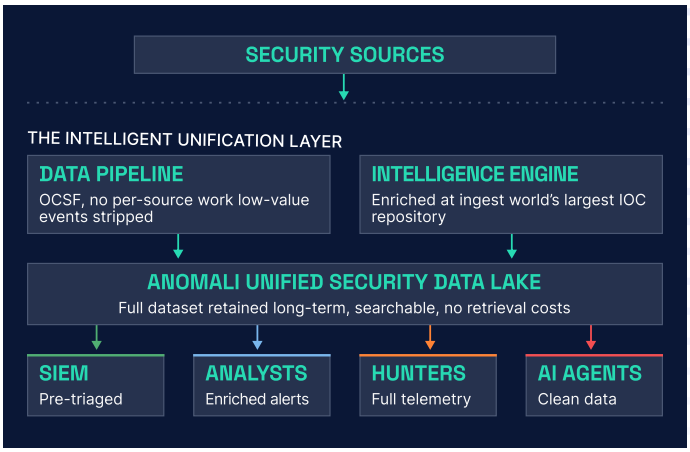

AUTONOMY A CISO CAN SIGN OFF ON

Automated action is where most SOC leaders get nervous, and with good reason. An agent that blocks the wrong thing at machine speed can cause the incident it was meant to prevent.

Anomali structures autonomy as a graded path rather than an on/off switch. Generally available capabilities today run from analyst copilot through assisted triage, with the analyst supervising and making the call. Higher levels of autonomous operation are on the roadmap. At every level, the platform calculates blast radius before an autonomous action executes, and every decision is logged, explainable, and reversible. Compliance evidence is produced from operational data as the system runs, rather than assembled in a separate exercise before an audit.

WHAT CHANGES FOR YOUR TEAM

For the Head of SOC, investigations that used to consume a shift resolve faster because context and history are already in place, and alert fatigue drops as low-confidence signals are suppressed before they reach the queue. For the CISO, the payoff is a defensible autonomy story, retention and analytics costs that track real usage rather than peak provisioning, and a modernization path that augments the current stack instead of replacing it.