ANOMALI

Overview

ANOMALI CORPORATE OVERVIEW

Beyond detecting. Decide and act.

See everything. Know what matters. Act with confidence.

WHO WE ARE

Anomali began in 2013 as ThreatStream, a threat intelligence company in Redwood City, California. The product collected, curated, and scored threat intelligence to help security teams distinguish real risk from noise. Rebranded in 2016, Anomali spent a decade building widely used threat intelligence recognized as a category leader by Forrester and Frost & Sullivan. That decade created a continuously curated threat graph spanning industries, actors, campaigns, and techniques—now operationalized directly into detection, investigation, and response via AI agents under human supervision.

THE PROBLEM WE SOLVE

Most mature SOCs no longer struggle to see threats—they struggle to decide what to do fast enough. Analysts get thousands of alerts a day across disconnected tools, threat intelligence lives in separate systems, and triage becomes reactive instead of anticipatory. AI has intensified the problem: attackers operate at machine speed, producing phishing, malware variants, and infrastructure faster than human investigation can keep up. Meanwhile, AI is only as good as its data—scattered telemetry in inconsistent formats increases noise and degrades decisions. An AI agent without trustworthy intelligence can confidently be wrong at machine speed, which is worse than slow response.

WHAT THE PLATFORM DOES

Anomali brings telemetry, intelligence, and governed AI together on a single data layer organized around three security priorities.

See everything: a unified security data lake holds full-fidelity telemetry from endpoint, network, cloud, identity, IT/OT, email, and SaaS, normalized to a common schema and searchable against years of history in seconds. Storage and compute are decoupled, so there are no retention trade-offs or performance ceilings—you pay for analysis when it runs.

Know what matters: Managed Intelligence as a Service, powered by ThreatStream Next-Gen, enriches every event at ingestion with threat context, fusing a large curated threat repository with your incident history and baselines. Detections include confidence and explain who, why, and what's likely next.

Act with confidence: embedded agents investigate, prioritize, and respond, acting on high-confidence cases first. Before autonomous action, the platform calculates blast radius across identities and entitlements, ensuring actions are explainable, reversible, and logged. Organizations can adopt these capabilities in stages, with changes to deployment—but not the foundation.

WHAT MAKES ANOMALI DIFFERENT

Our differentiators are built into the architecture: complete platform coverage with no concentration risk. It pairs a decade of operationally curated threat intelligence—matched for breadth, fidelity, and velocity—with MIaaS that operationalizes intelligence in minutes by enriching every signal at ingestion. Identity-grounded governance performs blast-radius math before autonomous action. The unified data lake includes native Threat Graph and identity topology, embedding curated, scored intelligence directly into queries with context competitors can't replicate. Trusted Circles enable private, bi-directional threat sharing with automatic TLP-gating for instant early warning.

WHERE THIS IS HEADING

Anomali built its reputation as a leader in threat intelligence. The platform now extends that foundation into the decisions and actions that intelligence is supposed to drive, on one governed data layer where context, intelligence, and bounded autonomy work together. The goal is a security operation program that allows teams to make decisions at machine speed to scales its outcomes without scaling its cost or its headcount.

PROOF POINTS

The following reflect customer-reported outcomes, anchored by a large global financial-services deployment.

A **30-50%** reduction in total cost of ownership versus legacy SIEM

A **90%** reduction in critical incidents.

Search times reduced from days to seconds across years of telemetry.

Roughly

60 - 70%

of analyst time reclaimed from chasing false positives.

Notably, those results came at the most conservative autonomy setting, with a human reviewing every decision and agentic capabilities not yet engaged.

SEE ANOMALI IN ACTION

Request a Demo