ANOMALI

White Paper

GOVERNANCE-READY AGENTIC SOC STARTS WITH EVIDENCE

Why the fidelity of the underlying data decides whether autonomous security decisions can be audited, defended, and approved.

Research from 2026[1] recorded the fastest adversary breakout at 27 seconds, with the average eCrime intrusion reaching lateral movement in about 29 minutes. Decisions at that speed cannot wait for a full human review cycle at the moment they happen. That is the operational case for agentic security. It is also the governance problem a CISO has to solve first: an action taken in 27 seconds still has to be explainable to a board, an auditor, or a regulator weeks later, and the explanation has to hold up.

That explanation is only as good as the evidence behind it. An agent acting on fragmented, unenriched, unnormalized data produces decisions no one can fully reconstruct. Governance-ready agentic operations start earlier than most evaluations assume, at the fidelity of the data every decision runs on.

THE CORE REQUIREMENT

FIDELITY IS A GOVERNANCE REQUIREMENT, NOT A TECHNICAL PREFERENCE

The quality of the data underneath an agent determines whether its decisions can be governed at all. A decision is auditable only when each input to it can be traced: the originating signal, the context applied, the confidence behind the call, and the authority that permitted the action.

When those inputs live in separate tools under different schemas with different retention windows, the record is incomplete before the agent ever runs. Fidelity is what makes the record complete.

MIT's 2025 study of enterprise generative AI[2] found that the large majority of deployments produced no measurable return, and the common failure was rarely the model itself, but the data the model was given. Security is the least forgiving version of that problem. An agent handed raw telemetry inherits every gap in it and acts on those gaps faster than the analyst it was meant to support. For a CISO, there are direct consequences: you cannot approve action you cannot defend, and you cannot defend a decision whose inputs you cannot reconstruct.

[1] CrowdStrike 2026 Global Threat Report

[2] MIT 2025 GenAI study

DECISION TRACEABILITY

WHAT A DEFENSIBLE DECISION RECORD HAS TO CONTAIN

A defensible agentic decision is one you can reconstruct end to end after the fact. Every decision has to carry its full lineage, and missing any single element turns an audit into guesswork.

The elements below are what auditors, regulators, and board members are actually asking for when they ask the system did what it did.

DECISION ELEMENT	WHAT IT ANSWERS	WHY IT HAS TO BE IN THE RECORD
Signal and source	What the system saw, and where it came from	Establishes the factual basis of the action
Applied context	Why this signal mattered here (intel, identity, asset value)	Separates a real event from noise
Confidence basis	How sure the system was, and on what grounds	Shows the response was proportionate to the evidence
Authorizing policy	What the agent was permitted to do, and who reviewed it	Ties the action to governance, not improvisation
Impact calculation	What the action would affect if it misfired	Demonstrates the risk was quantified before, not after
Action and reversibility	What happened, and whether it can be undone	Establishes control and recoverability

This record is less expensive and laborious to produce when the data underneath it is normalized, enriched at ingestion, and fully retained. Where the schema is consistent and the history is complete, reconstruction is a query rather than a forensic project. The Open Cybersecurity Schema Framework (OCSF), backed by more than 200 organizations and 900 contributors under the Linux Foundation as of November 2024, gives that record a common structure that no single vendor owns.

POLICY ALIGNMENT

AUTONOMY CALIBRATED TO RISK AND AUTHORITY

Governance-ready autonomy means an agent's authority is bound by policy and enforced at every action, not asserted once in a settings page. The workable model is a graduated one: the independent action an agent may take is tied to the confidence of the decision, the reversibility of the action, and the risk to the business.

Generally available operation today tops out at machine-led triage with analysts supervising and deciding on escalations and novel threats. Agents investigate and resolve common, well-understood incident types; people keep judgment over the rest. Higher levels of independent action, where a human moves further out of the loop, belong on a roadmap and should be evaluated as projections rather than delivered capability. The governance principle holds at every level: each action is explainable, reversible, and calibrated to risk, and each new level is entered only when the evidence justifies extending the trust.

Policy alignment also means the autonomy level maps to obligations the CISO already carries. Every action is logged as audit evidence produced from the same operational record, in-region where residency requires it, so compliance reporting draws from the system of record instead of a parallel effort.

RISK REDUCTION OUTCOMES

WHAT EVIDENCE-GRADE DATA CHANGES

Evidence-grade data reduces risk in two ways that matter to the board: it shortens the time from signal to correct action, and it lowers the cost and reach of the incidents that do land. Both trace to the same cause: decisions made on complete, contextualized data rather than partial signals.

IBM's 2025 Cost of a Data Breach report found that organizations using AI and security automation extensively contained breaches roughly 80 days sooner and at about $1.9 million lower cost than organizations that did not, against an industry-average breach lifecycle of 241 days. The mechanism behind that figure is not software acting alone. It is analysts and agents making faster, better-supported decisions because the data reached them already enriched and scored.

Additionally, CrowdStrike's 2026 report found 82 percent of detections were malware-free, meaning the adversary is using legitimate credentials and native tools rather than files a scanner can flag. Catching that requires context the scanner does not have: identity behavior, intel correlation, and a historical baseline. Data that carries that context at the point of decision is what separates a contained intrusion from a breach lifecycle measured in months.

HOW TO MEASURE SUCCESS

A GOVERNANCE SCORECARD FOR AGENTIC OPERATIONS

The right measures track defensibility, not speed alone. A CISO should be able to report, at any time, what share of automated and agent-assisted decisions carry a complete evidence trail, how quickly a past decision can be reconstructed, and how much risk has moved off the analyst queue without moving out of view.

MEASURE	WHAT IT TELLS THE BOARD
Decision reconstruction time	How long to produce the full lineage of any past agent decision. Governance-ready looks like minutes from a query, not days of forensics.
Evidence completeness rate	Share of agent and automated actions with a complete, reconstructable record. Anything short of near-total is ungoverned action.
Confidence-calibrated action mix	Share of actions taken at each autonomy level, matched to their outcomes. Shows whether trust is extended only where evidence supports it.
Pre-analyst suppression rate	Share of false positives resolved before an analyst spends time on them, with the basis logged. Removes noise while keeping the record.
Mean time to prioritized remediation	Time from signal to the right action on the things that matter, rather than raw alert throughput.
Reversibility and rollback	Share of autonomous actions that are reversible, and time to reverse. The precondition for extending autonomy safely.
Audit-evidence generation time	Time to produce regulator- or board-ready evidence from operational data.

WHAT TO REQUIRE BEFORE APPROVING AUTONOMOUS ACTION

Five conditions make agentic operations something a CISO can approve and later defend:

- **Reconstructable decisions**. Every agent decision traceable to source, on demand.
- **Policy-bound autonomy**. Independent action calibrated to reversibility and risk, with human review wherever the stakes justify it.
- **A non-proprietary schema**. A common structure such as OCSF, so the evidence is defensible outside the tool that produced it.
- **Complete retention**. Full history kept, so a decision can still be reconstructed when the question arrives months later.
- **Audit evidence from the operational record**. Compliance output drawn from the same data the decisions ran on, in-region where residency requires it.

Autonomy is earned on the strength of the evidence behind it, and that evidence has to exist before the first agent acts.

See what a defensible agent decision looks like end to end. Book a demo today.

SEE ANOMALI IN ACTION

Request a Demo