TRIAD SECURE™

WHITEPAPER

The Cost of an Alert

The hidden economics of security operations — why detection keeps improving while no one measures what it costs, and how to make investigation cost visible.

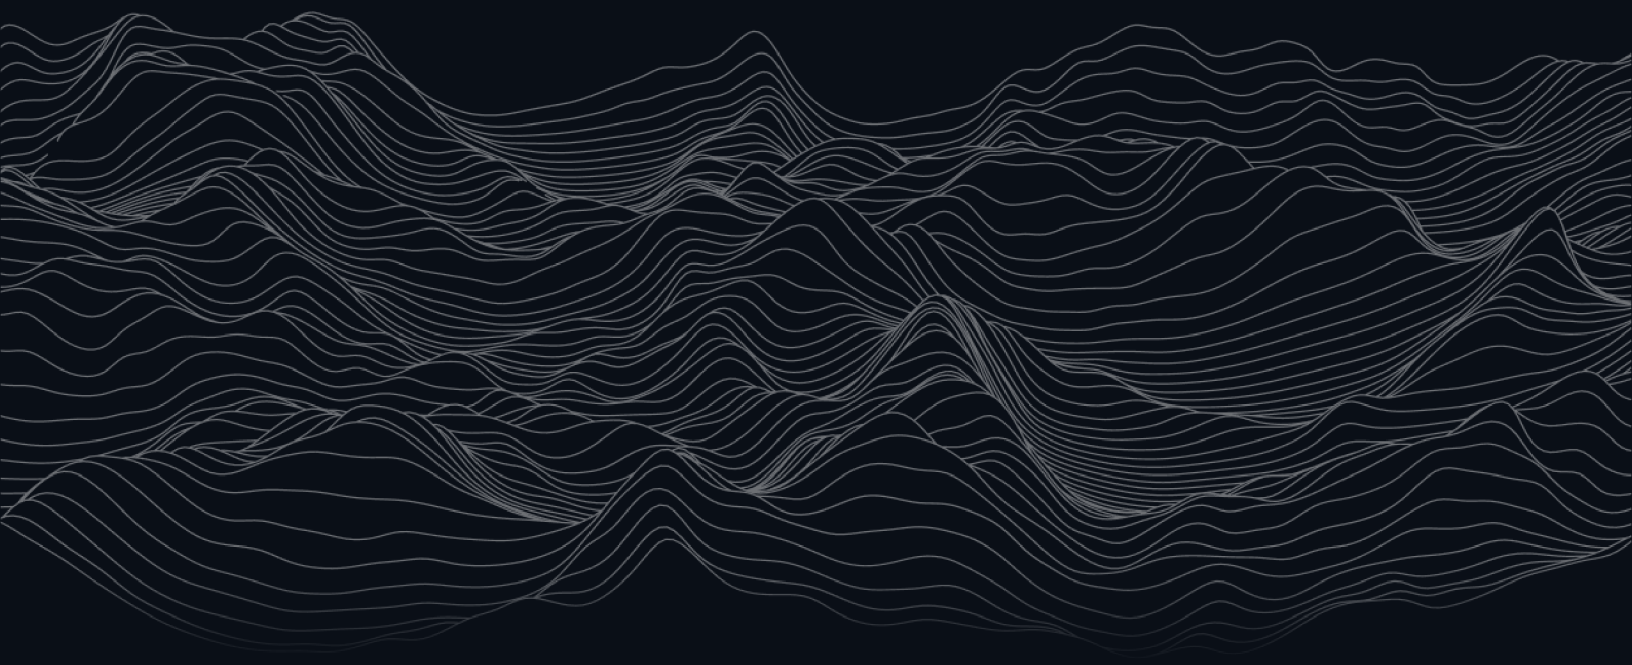

Executive Edition · Version 1.0 · August 2026
Triad Secure · Security Decisions That Persist · triadsecure.com

THE PROBLEM

Detection Improved. Nobody Measured What It Costs.

Security operations centers have never been better at detecting threats — or worse at knowing what detection costs them.

115M

alerts analyzed in a four-year production-SOC study published at USENIX Security

0.01%

of those alerts were associated with confirmed attacks

49%

were benign triggers — correct detections of legitimate activity that still required human effort to close

A decade of investment in SIEM, EDR, XDR, cloud, and identity platforms bought visibility. It also bought work: every alert generates a chain of triage, context gathering, correlation, documentation, escalation, and handoff that consumes the scarcest resource in the SOC — analyst attention. Nearly every one of those 115 million alerts made a claim on it.

Most SOCs measure none of that effort. They track detection and response speed — MTTD, MTTR, alert volume, closure rates — which describe security *events* but not security *work*. This paper makes the case for a second set of measurements: the operational economics of investigation. It shows where hidden costs accumulate, what to measure instead, where AI genuinely helps, and how to establish a defensible baseline.

> **Key takeaways.** (1) Traditional SOC metrics measure security events, not security work. (2) The dominant cost of an alert is the human investigation it triggers, not the tool that generated it. (3) Recurring false positives behave like debt — you pay interest in analyst hours until the root cause is fixed. (4) Headcount is not capacity; operating-model design determines how much analyst effort is actually usable. (5) AI creates value on bounded, verifiable tasks — measurement, not model access, is the prerequisite.

THE MEASUREMENT GAP

We Measure Security Events, Not Security Work

Two SOCs can produce identical metrics while carrying radically different costs.

Consider two SOCs that each investigate 500 alerts a week. Conventional reporting says they perform identically. In reality, one resolves alerts through integrated workflows with automated enrichment, while the other forces analysts to manually reconstruct evidence across a SIEM console, endpoint telemetry, identity systems, ticketing, and threat intelligence — rebuilding context at every step. Same output, radically different cost. Traditional metrics cannot tell these two organizations apart.

Organizations cannot optimize dimensions they fail to observe.

The blind spot has consequences. When leadership sees only closure rates and response times, it optimizes for speed — and speed can be achieved by shallower investigation, weaker documentation, premature escalation, or unsustainable analyst workload. The metrics improve while the operation degrades.

Alert fatigue is usually framed as a human-factors problem: analysts get desensitized. It is more useful to treat it as an economic condition — a persistent imbalance between the investigative demand alerts create and the finite processing capacity of the team. Controlled experiments bear this out — analysts working under an 86% false-alarm rate showed measurably lower precision and slower task completion than analysts at 50%. False positives are not neutral workload; they degrade the performance of the whole human-machine system.

> **The practical consequence.** Every efficiency claim made to a board today — about consolidation, automation, or AI — is being made without a denominator. You cannot show that investigation cost fell if you never established what it was.

ALERT ECONOMICS

Every Alert Is a Claim on Analyst Attention

When investigative demand exceeds operational capacity, the cost arrives in a predictable order.

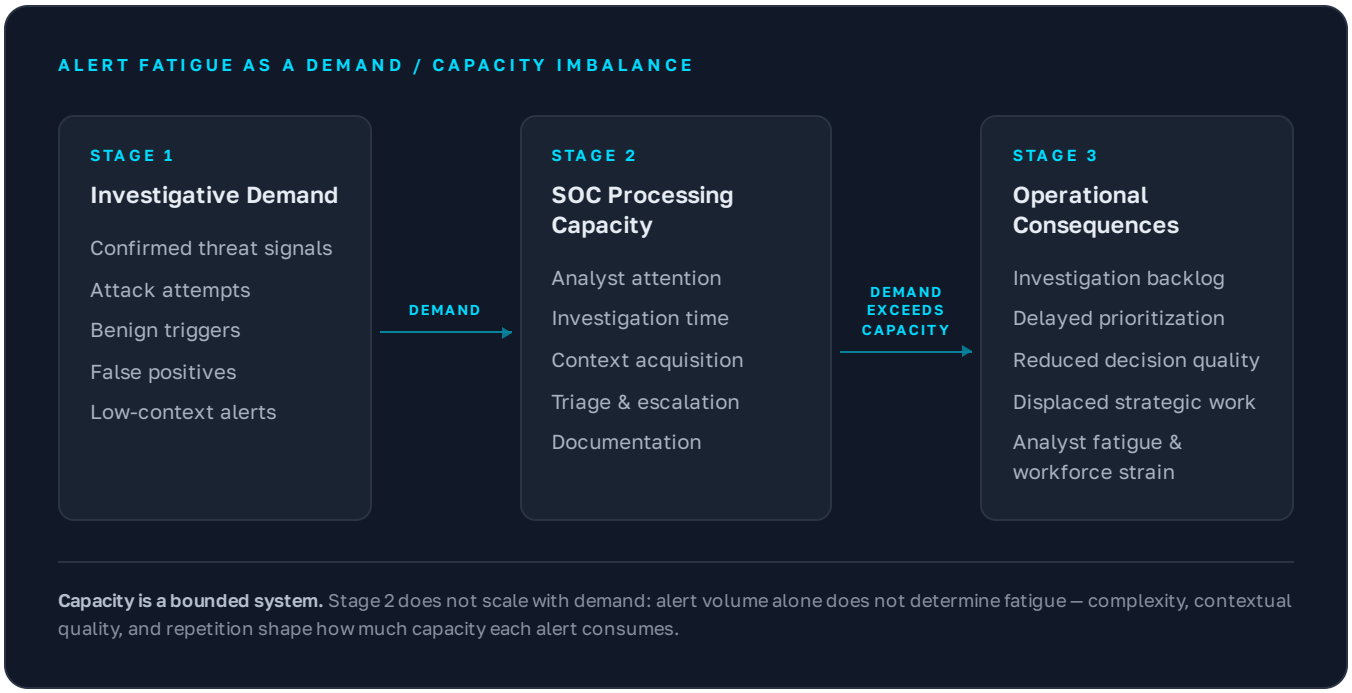

When the imbalance persists, backlogs grow first, then prioritization slips, then decision quality erodes, then strategic work is displaced — and finally the workforce itself strains: burnout, attrition, and the loss of the institutional knowledge that made investigations fast in the first place. This is why adding detection technology without measuring its operational demand can make a SOC worse off: each new telemetry source improves visibility while adding claims against a capacity pool that did not grow.

> **The measurable version.** Track investigative demand — alerts entering investigation, by class — against usable capacity, meaning analyst-hours actually available for analysis. Reported weekly as a ratio, it predicts every Stage 3 consequence before any of them appear in MTTR.

COST VISIBILITY

Tool Cost Is Visible. Investigation Cost Is Not.

Security budgeting captures the cost of acquiring capability — not the larger cost of operating it.

When that labor is unmeasured, it functions as a hidden subsidy. Analysts compensate for weak integrations by retrieving evidence manually, for missing context by rebuilding asset and identity information, for inconsistent documentation by repeating one another's work. A platform can look inexpensive precisely because the cost of operating it is buried in analyst workflows. Two organizations running identical stacks can carry radically different investigation costs — the difference lives in workflow design, context availability, and knowledge preservation, not in the tools.

> **The question to ask at renewal.** Not "what does this platform cost?" but "what does an investigation originating from this platform cost, and has that number moved since we bought it?" Only the second question is about value.

THE UNDERLYING CONDITION

Security teams don't lack signals. They lack a visible data plane.

An alert is only as decidable as the environment around it. When assets, identities, network paths, cloud posture, and prior decisions live in separate systems, every investigation begins by reassembling the environment by hand — and that reassembly is the cost.

Triad Secure is an orchestration platform. It makes the entire data plane visible and correlates it against live events, so each alert arrives with the surrounding context required to decide — and so the decision, once made, persists into the next one.

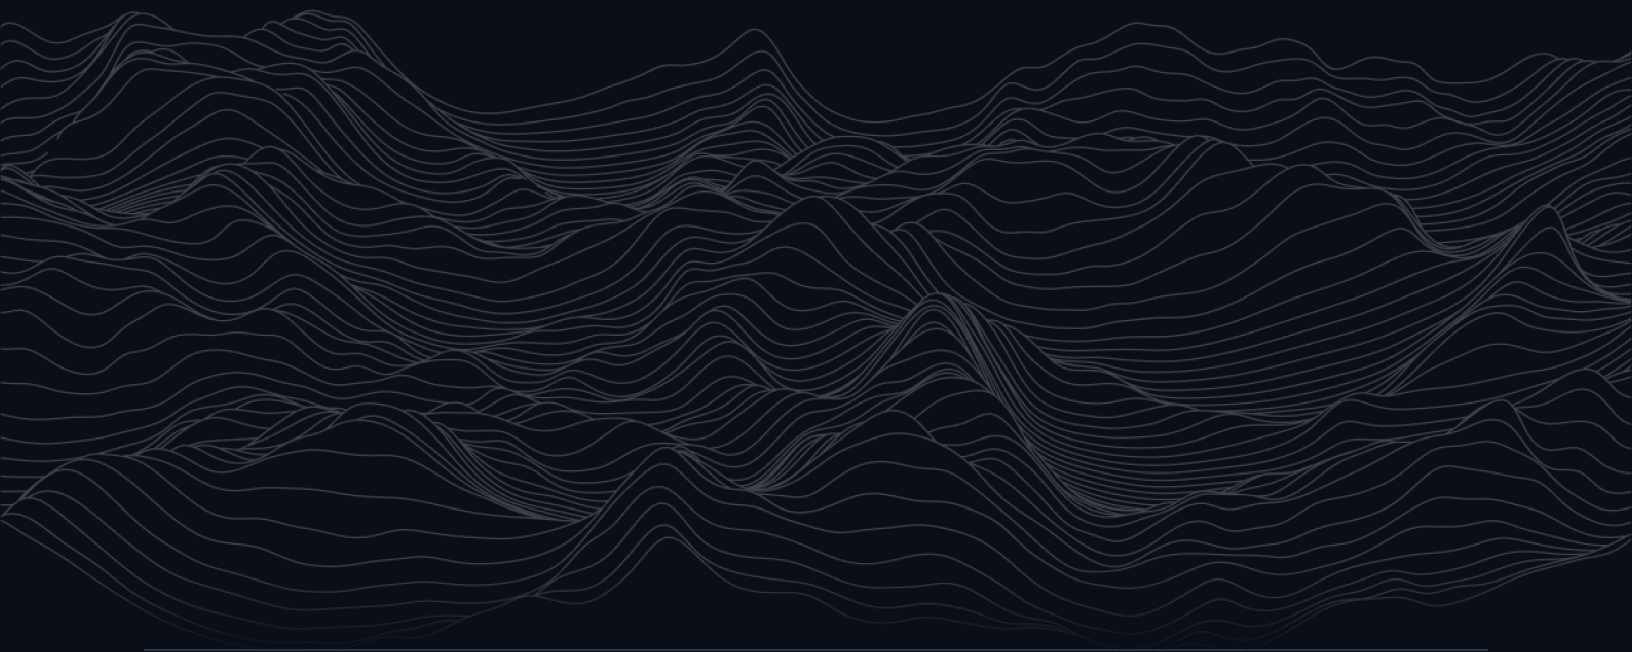

OPERATIONAL DEBT

False Positives Are Operational Debt

Closing an alert resolves the task. It does not resolve the condition that generated the alert.

When an overly sensitive rule, a missing piece of business context, or a duplicate-alert pattern goes unremediated, the organization pays for the same deficiency again and again. This explains why a SOC can look busy and productive while becoming no more efficient — the queue keeps moving, but the operation keeps servicing the same debt. It also warns against the cheap fix: aggressive suppression reduces workload but can silence the rare, consequential signal. The goal is to eliminate avoidable recurring effort, not sensitivity.

> **Where to start paying down.** Rank recurring alert patterns by total analyst-hours consumed over 90 days — not by volume. In most environments the top ten patterns account for the majority of avoidable investigative effort.

ANALYST CAPACITY

Headcount Is Not Capacity

Scheduled labor is consumed by friction before it becomes net-new analytical work.

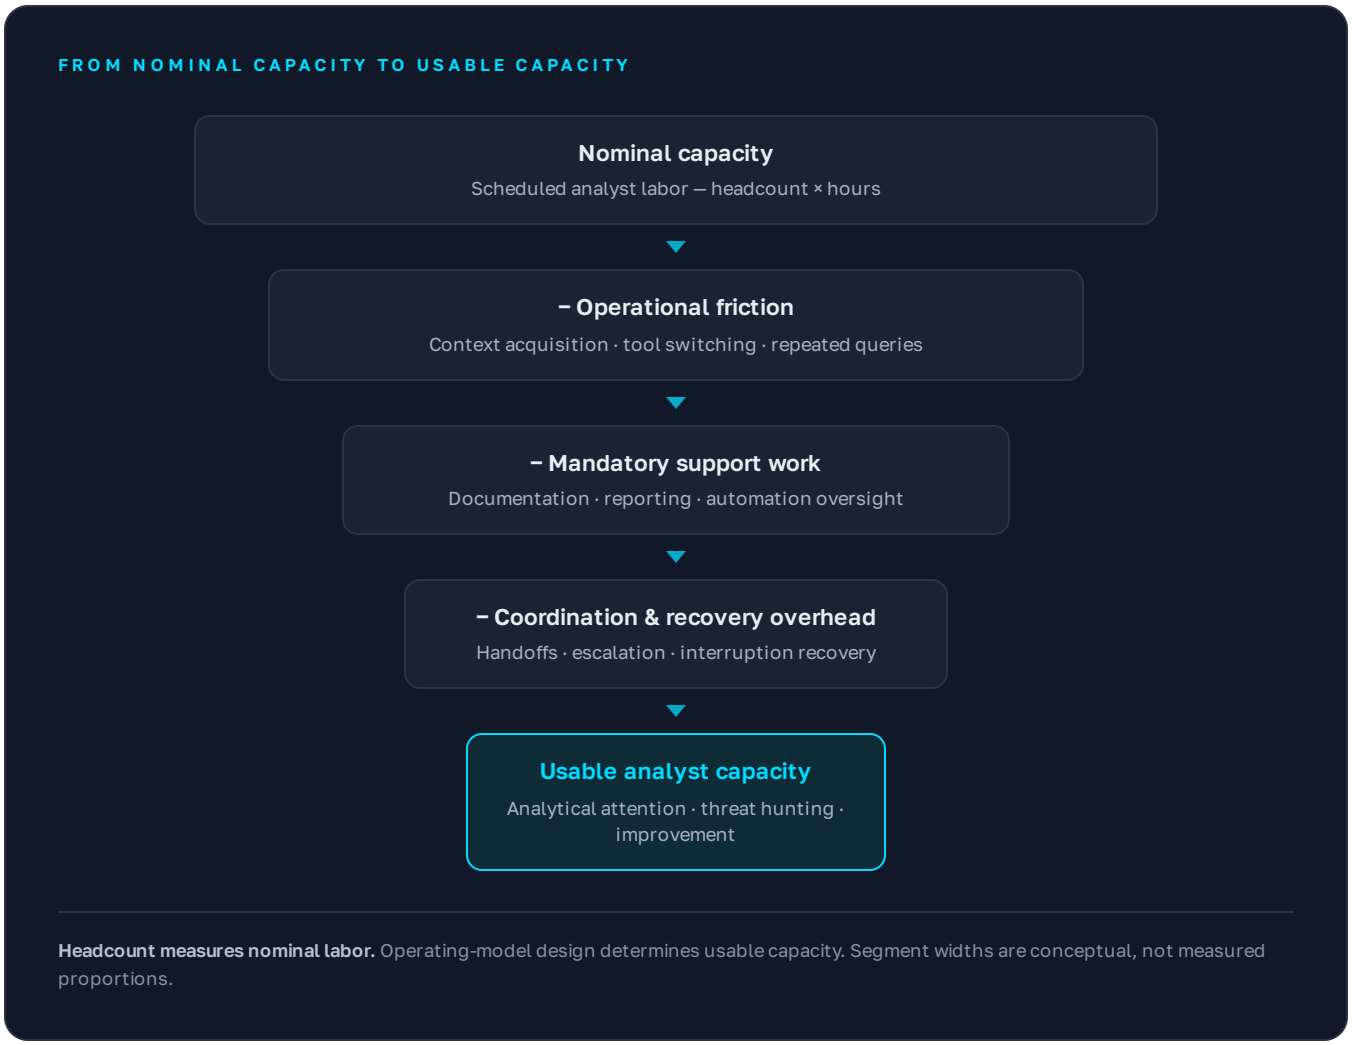

Handoffs are a particularly expensive drain: an investigation that transfers without its evidence, reasoning, and open questions attached forces the recipient to rebuild knowledge the organization already paid to create. This is why hiring alone rarely fixes an overloaded SOC — more analysts inherit the same friction. The differentiated question is how much capacity is recoverable and redirectable toward threat hunting, detection engineering, and improvement work.

> **The metric most SOCs are missing.** Usable capacity as a percentage of nominal. Until that ratio exists, every staffing conversation is a negotiation about headcount rather than about work.

THE ARITHMETIC

What This Costs: A Worked Example

Illustrative figures, not benchmarks — the point is that the arithmetic is simple once the inputs are measured, and most SOCs have never measured them.

INPUT / RESULT	VALUE
Team	**10 analysts, ~$150K fully loaded each**
Investigated alerts	**1,000 per week**
Median active investigation time	**18 minutes**
Direct cost per investigated alert	**~$22**
Annual direct investigation cost	**~$1.1M**
Cost of establishing benignity (~49% of alerts)	**~$550K per year**
Recoverable if one-third of recurring benign workload is remediated	**~$180K · ~2,500 analyst-hours · ~1.2 FTE redirected**

Note what the model reveals: roughly half of this SOC's investigation budget is spent proving that nothing happened — and a meaningful fraction of that spend recurs against known, fixable causes. That is the case for measuring before buying.

> **Why estimates beat silence.** An explicit, reviewable estimate with disclosed assumptions outranks a false exact figure — and both outrank the status quo, in which the largest line item in security operations is not written down anywhere.

MEASUREMENT

What to Measure Instead

No single KPI can represent SOC performance. Keep MTTD and MTTR — and add the operational dimensions they cannot see.

SECURITY-OPERATIONS EFFICIENCY — AN EMERGENT PROPERTY OF SIX DOMAINS

DOMAIN 1
Technical Effectiveness
- Detection coverage
- Alert quality & duplication
- Telemetry completeness

DOMAIN 2
Investigative Efficiency
- Active investigation time
- Context-acquisition time
- Tool transitions · rework

DOMAIN 3
Decision Quality
- Evidence sufficiency
- Disposition accuracy
- Reporting quality

DOMAIN 4
Investigation Continuity
- Handoff completeness
- Context preservation
- Decision reuse

DOMAIN 5
Analyst Capacity
- Usable capacity
- Workload distribution
- Proactive-work allocation

DOMAIN 6
Organizational Value
- Total investigation cost
- Capacity recovered
- Improvement vs. baseline

Report the domains together, segmented by alert class. A blended average conceals exactly the variation that matters.

The essential additions to a conventional dashboard: cost per alert by disposition, active investigation and context-acquisition time, rework and repeated-query rates, handoff completeness, usable analyst capacity, and improvement against your own baseline.

> **A reporting rule.** Never report a domain in isolation, and never report a blended average across alert classes. Gains in one domain are frequently paid for by losses in another — segmentation is what makes the trade visible.

AUTOMATION & AI

Where AI Fits — and Where It Does Not

The most consequential efficiency lever available to the SOC, and the most commonly misapplied.

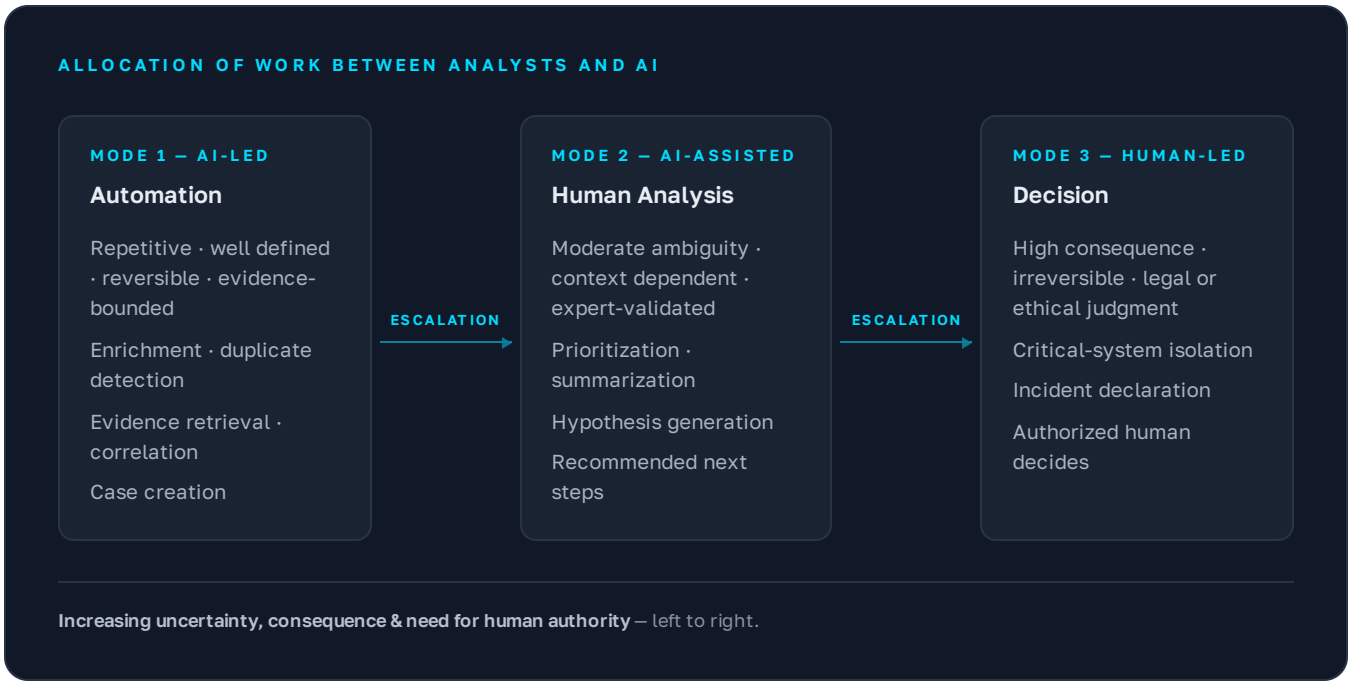

AI creates real value on bounded, verifiable work. It assists — but does not replace — analysts on judgment-heavy synthesis, and consequential decisions must remain under accountable human authority, with evidence the human can actually inspect. A "human in the loop" who cannot see the evidence, challenge the recommendation, or stop execution is a checkbox, not a control. And one sequencing rule matters more than any model choice: operational readiness precedes AI readiness. Deployed onto an unmeasured process, AI makes bad workflows faster, less visible, and harder to challenge. Measure first.

> **The readiness test.** Before deploying AI into an investigative workflow, confirm you can state that workflow's current cost, duration, and rework rate. If you cannot, you will not be able to tell whether the AI helped.

SELF-ASSESSMENT

Where Does Your SOC Stand?

Ten questions. Answer honestly — most organizations cannot answer yes to more than two or three.

1	Can you state your median active investigation time by alert type?	Yes / No
2	Do you know what share of alerts close as benign triggers vs. technical false positives vs. attack attempts?	Yes / No
3	Can you put a dollar figure on reaching a benign disposition?	Yes / No
4	Do you measure context-acquisition time separately from analysis time?	Yes / No
5	Do handoffs transfer investigative state — evidence, reasoning, uncertainty — or just the alert?	Yes / No
6	Do you track how often analysts repeat queries or rebuild timelines someone already produced?	Yes / No
7	Do you know how much senior-analyst time is consumed by routine consultation?	Yes / No
8	Is automation measured by net workload reduction after validation and exceptions — not actions executed?	Yes / No
9	Is any analyst capacity protected for proactive, capacity-creating work?	Yes / No
10	Could you demonstrate, against a baseline, that last year's investments reduced investigation cost?	Yes / No

Seven or more "no" answers is typical. It is not a failing grade — it is the visibility gap this paper describes, and it is closable.

THE PATH FORWARD

Establish the Baseline.

You do not need a year of instrumentation to start. The Triad Secure Operational Assessment establishes a credible, defensible baseline in weeks — in two stages.

STAGE 01

Rapid Modeling

Model your investigation cost, false-positive burden, and recoverable capacity from operational inputs you already have — alert volume, disposition mix, investigation time, team size — with every assumption disclosed.

STAGE 02

Operational Validation

Validate the model against case records, workflow data, and analyst interviews. You get the concentration points of avoidable work and a prioritized intervention plan, with the measurement to prove what changed.

To schedule an assessment or request the full Research Edition — including the complete measurement architecture, thirteen operational benchmarks, and the peer-reviewed evidence base — visit triadsecure.com.

WHAT THE ASSESSMENT PRODUCES

OUTPUT 1

Investigation Cost Baseline

- Cost per alert by disposition
- Annual investigation spend
- Every assumption disclosed

OUTPUT 2

Map of Avoidable Work

- Recurring-pattern concentration
- Context and handoff gaps
- Recoverable capacity estimate

OUTPUT 3

Prioritized Intervention Plan

- Sequenced by hours recovered
- Owner and effort per item
- Measurement to prove the change

SOURCES

This edition summarizes a longer research paper synthesizing recent peer-reviewed work, including: Yang et al. (2024), *33rd USENIX Security Symposium* — the 115M-alert production-SOC study · Layman & Roden (2023), *Human Factors and Ergonomics Society* — false-alarm rate and analyst performance · Bridges et al. (2023), *Computers & Security* — SOAR tools in use · Agyepong et al. (2023), *Computers & Security* — SOC analyst performance measurement · Ofte (2024), *International Journal of Information Security* — situation awareness in critical-infrastructure SOCs · Chhetri et al. (2024), *ACM Transactions on Internet Technology* — human–AI teaming against alert fatigue.