Command Zero

AUTONOMOUS & AI-ASSISTED SOC

TECHNICAL BRIEF

Choosing the Right SOC Workloads for AI & Automation

A framework for rethinking what, when, and how we automate in security operations

By Eric Hulse | Director of Research | Command Zero

THE PROBLEM

More Automation Created More Pain

For years, the industry promised that automating SOC workloads would free analysts for meaningful work. The reality: alert fatigue is at all-time highs, analysts drown in playbook maintenance, and burnout is increasing. We automated the noise - not the pain.

The Five Automation Errors

1	**Wrong Unit of Work** We automated alerts instead of investigations; enrichment without synthesis
2	**Happy Path Design** Playbooks handle the 85% that doesn't matter, zero help with the 15% that does
3	**No Feedback Loop** Measured speed and throughput, never verdict accuracy; faster at being wrong
4	**Integration ≠ Automation** Most "automation" is API plumbing connecting tools, not making decision
5	**Automated Before Understanding** Codified tribal knowledge and ad hoc runbooks, encoding gaps as polic

THE FRAMEWORK

The Three-Layer Model: Where Autonomy Actually Belongs

The core insight: we over-invested in Layer 1 (data retrieval), skipped Layer 2 (pattern recognition) where AI's highest ROI lives, and risk automating Layer 3 (decision-making) that must stay human.

LAYER 1

Data Retrieval

OVER-INVESTED

Query, enrich, normalize. High volume, low cognition. The last fifteen years have been focused here.

LAYER 2

Pattern Recognition

HIGHEST AI ROI

Correlation, hypothesis formation, anomaly evaluation. Where analysts struggle most and AI adds most value.

LAYER 3

Decision-Making

STAYS HUMAN

Verdict, risk acceptance, response authorization. Judgment under uncertainty requires human accountability.

LAYER 1 Data Retrieval

Solved - But Not Sufficient

Alert triage, ticket creation, IOC enrichment, and reporting are rightly automated. But enriching data at machine speed doesn't help when an analyst still interprets 47 fields per alert. Layer 1 automated retrieval without relieving the cognitive burden downstream: working memory overload, context switching, decision fatigue, ambiguity stress, and emotional load still hit analysts on every investigation.

Layer 1 is solved. The real problem was never about getting data in - it was making sense of what came out.

LAYER 2 Pattern Recognition

The Highest AI ROI - And the Most Neglected Layer

Analysts spend 60% of investigation time assembling context across disparate tools. No automation currently helps them think - tools only deliver more raw material. Three domains define the Layer 2 gap:

- **Context Assembly** - Connecting alerts to related events, building timelines from 12+ data sources, recognizing when "separate" alerts are one campaign
- **Decision Support** - Evaluating which enrichment data matters for this specific alert, weighing conflicting indicators, generating and testing hypotheses about attacker intent
- **Knowledge Capture** - Distinguishing novel threats from known patterns, applying institutional knowledge, transferring expertise before senior analysts leave

The solution is an agentic investigation loop - not a linear pipeline. The system forms a hypothesis, gathers evidence, evaluates it, and iterates until it reaches a confident verdict or determines it needs human judgment. This is fundamentally different from SOAR playbooks that run fixed sequences.

The compounding advantage: every investigation enriches the knowledge layer. After 1,000 alerts, the system understands which patterns are benign noise and which are genuine signals - specific to your environment. Layer 1 playbooks can never achieve this.

The Experise Gap

Encode Analyst Cognition, Not Just Workflows

SOAR encoded mechanical steps: data flows, tool sequences, enrichment lookups, report generation. Analyst judgment was still required at every step. What actually needs encoding is analyst cognition:

- Which questions a senior analyst asks first
- How they evaluate ambiguity and incomplete evidence
- When they pivot their investigation approach
- What the absence of evidence means

The same password spray alert triggers three distinct mental stances in an experienced analyst: pattern matching ("is this noise I recognize?"), threat assessment ("is this a real attempt that failed?"), and blast radius evaluation ("did something actually happen?"). Each stance drives a different investigation path and outcome. Junior analysts don't know which questions to ask first. When you encode this mental model, junior analysts follow expert reasoning - not just runbook steps.

LAYER 3 Decision Making

What Stays Human - And Why It Must

Final verdicts, risk acceptance, and response authorization stay human. Not because AI can't process evidence - it can - but because accountability requires judgment, context, and authority that no model can carry. Blocking a user or isolating a server has real business impact. Someone must own that decision.

There is also a hidden cost to automating this layer: the expertise pipeline. Junior analysts develop pattern recognition from thousands of triage decisions. If automation removes that pipeline, today's juniors never become tomorrow's seniors. The three modes of analyst growth are: REPLACE (skills with no upstream value), SCAFFOLD (guide faster skill development - most powerful, least implemented), and CREATE (automate the safety net while analysts practice real casework).

IMPLEMENTATION

Right-Sizing Automation Across All Three Layers

Layer	Right Automation	Common Mistake	Success Metric
Layer 1 Data Retrieval	Fully automate data retrieval	Dumping all data on the analyst	Time to relevant information
Layer 2 Pattern Recognition	AI-assisted pattern recognition	Skipping this layer entirely	Verdict accuracy rate
Layer 3 Decision-Making	Human verdict, AI prep	Automating the decision itself	Override rate + quality

The implementation test: for every automation, ask which layer it serves. If you can't answer, you're probably making one of the mistakes above.

Your First 30 Days Action Per Layer

Layer 1

- Audit one playbook: Does it reduce cognitive load, speed up decisions, or help analysts grow? If not, redesign or retire it.

Layer 2

- Decompose one investigation: Map your most common alert type as a decision tree. That document is your Layer 2 automation spec.
- Make one automation transparent: Add the full reasoning chain to an opaque TP/FP verdict and watch analyst engagement change.

Layer 3

- Capture override context: Every analyst override - capture not just that they disagreed, but why. One sentence trains Layer 2 over time.
- Scaffold one junior analyst: Give a junior a Tier 2 case and track what they learn, not just close rate.