BEYOND THE PILOT

AI in Security Operations

JOE SCHUMACHER

Chapter 2
How AI Reframes Security Operations

A security alert fires. Someone investigates. They pull data from your SIEM, check cloud logs (e.g., AWS CloudTrail, Azure Log Analytics), cross-reference threat intelligence, and search user directories. They stitch fragments from multiple systems into a coherent story. This takes hours. Your leadership asks the natural question:"Can AI speed this up?"

Yes, but the answer depends entirely on one foundational reality: the quality and accessibility of your data. Many organizations expect AI to solve security speed problems when the real bottleneck is incomplete data visibility. If your event data scatters across disconnected systems, an AI system inherits that fragmentation. If your log retention fluctuates with storage costs, historical analysis becomes unreliable. If you have blind spots in your monitoring, an AI system amplifies them at scale.

This chapter reinforces what you already understand about security operations and reframes where AI genuinely helps. More importantly, it establishes the constraints that determine what's realistic in your environment.

The Three Core Loops: Where AI Fits

Security operations rest on three interconnected loops: detection, investigation, and response. Understanding where AI accelerates and where it cannot, is essential before you choose any architecture.

Detection: Generating Signals

Detection generates alerts based on rules, baselines, or behavioral anomalies. Your SIEM processes hundreds to thousands of alerts daily. Most of the alerts appear to be noise, with a handful being threats. We are hitting a point where human analysts cannot triage this volume effectively.

AI creates genuine leverage here. An AI system scores alerts in seconds based on your environment, threat patterns, and historical data. It can enrich each alert with context such as user role, device posture, network segment, and threat intelligence, then surface the most likely threats first. Analysts then verify top-priority alerts instead of sifting through all of them.

What stays human: The decision to act. After AI raises an alert, an analyst determines whether it's an actual threat by performing a lighter manual analysis. The analyst understands the intent and context of the alert in ways that AI can struggle with. They know which services legitimately generate suspicious behavior.

Investigation: Building Understanding

An alert is triaged as real. Now an analyst gathers context: What happened before? After? Is this user's behavior abnormal? Does this correlate with other signals?

AI accelerates the investigation significantly. An AI system searches historical logs in seconds, correlates events across infrastructure, and surfaces patterns humans would miss. It suggests related events and identifies similar past incidents.

What stays human: The determination of significance. AI can show that five events correlate. Humans determine whether that correlation represents a real attack, misconfiguration, or coincidence.

Response: Acting on Judgment

An analyst has investigated and determined that an alert is real. Now comes judgment: Is this critical? Should we isolate systems? Revoke credentials? These decisions depend on your environment, threat model, and business context.

AI informs response by providing a detection rule to block the threat or suggesting actions based on similar incidents and known playbooks. It provides the context needed for a human to take quicker actions.

What stays human: The decision itself. Response scope requires human judgment and accountability. An AI system shouldn't make these decisions autonomously without understanding the full risks.

What's Feasible

Data architecture determines what AI patterns are even possible. This depends on whether you operate traditional on-premises, cloud-centric, or hybrid security operations.

Traditional on-premises models centralize data in a SIEM. Network, system, and identity data flows into a single point. Detection rules run on the SIEM. Analysts query the SIEM. An AI system integrates tightly with your SIEM: it receives alerts, queries historical data, and enriches results.

Cloud-centric models distribute data across cloud-native stores. AWS generates CloudTrail logs in S3 buckets. Azure generates logs in Log Analytics. SaaS applications expose logs via APIs. Detection rules run in cloud SIEMs or serverless functions.
An AI system must handle distributed data ingestion, query multiple backends in parallel, and synthesize results.

Hybrid organizations operate simultaneously. On-premises infrastructure feeds your SIEM. Cloud workloads generate distributed logs. An AI architecture must bridge both models.

An AI system designed for traditional SIEM architecture will not work well in cloud-centric environments. An AI system intended only for cloud-native environments will struggle with legacy infrastructure. Understanding your operational model before selecting solutions is critical.

What AI Can Do

Security operations teams are not homogeneous. Different roles need fundamentally different support from AI.

Tier 1 SOC analysts perform triage. They benefit from AI-pre-scored alerts and from routing high-priority work to senior analysts. They suffer when AI makes autonomous decisions without human verification.

Tier 2 and 3 analysts perform investigations and have technical depth. They need AI that accelerates their research, such as log search, correlation, and threat intelligence synthesis, without replacing their judgment. Experienced analysts treat AI as a tool; they verify outputs and catch AI errors. Junior analysts often trust AI output without verification. This matters architecturally: systems designed for experienced analysts fail when junior staff operate them without oversight.

Threat hunters search for novel threats that evade detection. They need AI that helps them search large datasets, identify patterns across sources, and test hypotheses at scale.

Detection engineers build and tune detection rules. They need AI that helps them test rules and identify false-positive patterns. They are skeptical of automation that replaces their craft because they know detection is a discipline and an iterative process.

A generic "AI for security operations" solution ignores these differences and solves no one's problem well. Worse, it often creates new problems: junior analysts rely on AI too heavily; senior analysts are frustrated by systems designed for their juniors; experienced hunters find tools that don't match their workflow. Your architecture must account for how different roles actually work and what expertise each role brings.

The Human Judgment Boundary

Understanding where humans make decisions is critical because this is where AI must fit, not replace.

AI can accelerate triage: Score alerts, rank by priority, surface the top candidates for human review. Time savings are significant when analysts verify results.

AI can accelerate investigation: Search historical logs, correlate events, and synthesize threat intelligence. Analysts then determine what findings actually matter. Time savings are moderate to significant, depending on your data volume.

AI can inform responses by summarizing findings, suggesting actions based on similar incidents, and providing context for faster decision-making. Humans make containment, escalation, and notification decisions.

Where AI adds complexity causes failures:

When humans are removed from judgment, if an AI system auto-suppresses or auto-prioritizes alerts without verification, real threats disappear silently. Alert fatigue drops, but so does detection. By the time you discover the AI system is wrong, threats have passed unnoticed.

When automation replaces expertise, if an AI system presents conclusions without justification, analysts cannot effectively verify them. If they unquestioningly trust the system and it fails, they stop learning, and expertise erodes as operations become dependent on flawless AI performance.

When the response is automated without oversight, an AI system can automatically isolate systems or revoke credentials based on its own judgment, causing operational disruption, compliance issues, or missing critical context that a human would catch.

The principle: AI informs and accelerates. Humans judge and decide. When this boundary blurs, projects fail.

Key Takeaways

This chapter explains that AI is a powerful engine, but it only runs as well as the data you feed it and the team that steers it.

- **The Data Bottleneck**: AI cannot fix a mess. If your data is scattered across different systems or has "blind spots," the AI will only make those problems bigger and faster.
- **The Three Loops**: AI provides the most value in Detection (sorting through noise) and Investigation (finding connections). However, the Response (making the final call) should almost always stay with a human.
- **One Size Does Not Fit All**: Your setup, whether you are all in the cloud, all on-premises, or a mix of both, changes which AI tools will actually work for you.
- **The Junior Analyst Risk**: Less experienced team members might trust the AI too much. You must design your system so that humans still verify the AI's work, otherwise, your team's skills will "atrophy."

Strategic Questions

Before choosing an AI path, use these questions to gauge how much work your foundation needs.

- When an alert happens today, how many different screens does an analyst have to jump between to get the full story?
- Are we choosing an AI tool that fits our specific environment (Cloud vs. Hybrid), or are we trying to force a "square peg" into a "round hole"?
- If the AI makes a mistake and shuts down a critical business system, who is accountable, and does the AI have the power to override the machine?
- Does our current team have the seniority to double-check the AI's logic, or are we just hoping the AI is always right?

What Comes Next?

Understanding your foundation is the first step toward choosing the right "blueprint" for your AI system. You wouldn't build a house without choosing an architectural style first, and the same applies here. There isn't just one way to "do AI" in security; there are specific patterns designed to solve specific problems.

In Chapter 3, we move from theory into design. We will introduce three key architectural patterns: **MCP**, **RAG**, and **Agents**. We'll define these terms in plain English and help you identify which one matches the specific bottlenecks we identified in your operational state.

About the Author

Joe brings a relentless operational mindset to the world of Artificial Intelligence, backed by over twenty years of experience at every level of the cybersecurity landscape. From the high-pressure environment of incident command to the strategic oversight of a virtual CISO, Joe has spent his career navigating the friction between emerging threats and organizational constraints.

As the founder of Focused Hunts, Joe bridges the gap between proactive defense and reactive response. His consultancy integrates advanced technologies with deep practitioner expertise, serving as a trusted advisor for organizations seeking to modernize their security posture without losing technical rigor.

In this volume, Joe addresses the critical intersection of architecture, operations, and governance. He moves beyond the "AI hype" to provide a repeatable framework for building agentic workflows and RAG systems that are actually sustainable in a production environment. His approach focuses on the "amplification model," using AI not as a replacement for human expertise but as a catalyst for it.

Joe holds the GIAC Certified Forensics Analyst (GCFA) and Certified Information Systems Security Professional (CISSP) certifications. He remains committed to the idea that, while technology accelerates at a dizzying pace, the core of security remains a human journey that requires constant discernment, discipline, and a focus on operational reality.

https://www.focusedhunts.com/

Take a break from work with this short interactive hunt game:
https://game.focusedhunts.net/

Books available at Amazon
https://www.amazon.com/dp/B0GKHZNMP4