Dropzone AI

From Reactive to Proactive:

The Future of AI in the SOC

Steve Zalewski – Former CISO, Levi Strauss

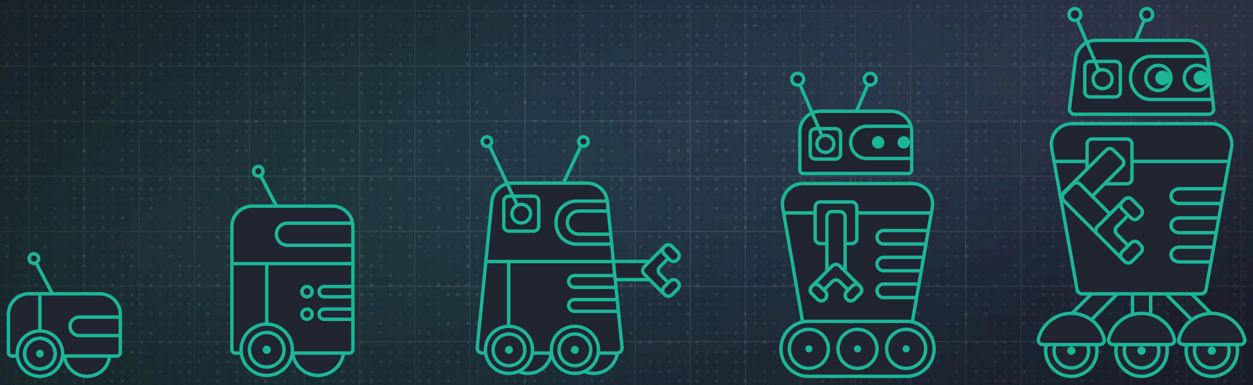

Contents

About the Author

Steve Zalewski is currently a Cybersecurity Advisor at S3 Consulting, Inc and a co-host of the Defense in Depth podcast. Prior to this, he was CISO at Levi Strauss & Co. Steve has multiple patents in data protection and multi-processor operating system design and holds CISSP, CISM and CRISC security certifications.

The New Battlefield

In the past year, the balance of power in cybersecurity has shifted. Attackers no longer rely solely on human ingenuity or manual tools; they are increasingly deploying artificial intelligence to accelerate reconnaissance, automate phishing, create deepfakes, and orchestrate multi-pronged intrusion campaigns at machine speed. This has created a new reality for security operations centers (SOCs): the threat landscape is now driven by AI, and defenders must adapt accordingly.

Traditional SOC models were designed for a different era in which human analysts could keep pace with the flow of alerts, and static defenses could be hardened over time. Today, those models are under strain. Alert volumes are exploding, false positives consume precious time, and experienced analysts are burning out. The result is a widening gap between the speed of attack and the speed of defense.

Closing that gap is no longer optional. If adversaries are leveraging AI to mount faster, more adaptive campaigns, SOCs **must** respond in kind to reduce workload and fundamentally reframe their defensive posture. This isn’t about cutting costs; it’s about survival.

The good news is that defensive AI is also maturing. With the right approach, SOCs can deploy AI to handle the high-volume, repetitive investigations that slow response times, allowing human analysts to focus on critical decision-making and strategic resilience. Solutions like Dropzone AI’s AI SOC Analyst, which can investigate alerts autonomously, integrate with existing tools, and produce decision-ready reports within minutes, are already demonstrating how defensive AI can change the tempo of security operations without adding headcount or complexity.

This white paper uses a simple but powerful lens, the analogy of raising a child, to explain the stages of AI maturity, tie them directly to SOC operations, and show why adopting AI is now a strategic necessity to counter AI-powered attackers.

The AI Maturity Model — From Child to Adult

Artificial intelligence doesn’t spring into existence fully formed. Like a child, it develops or matures in stages, growing in capability, judgment, and independence over time. Understanding these stages is essential for SOC leaders deciding where AI fits in their operations, how much to trust it, and what role it should play in defending against AI-enabled threats.

Below, we outline the five stages of AI maturity using the child-to-adult analogy, with direct links to SOC functions and the types of work AI can handle at each stage. At every level, value can be realized without committing to the full journey; organizations can stop at the stage that aligns with their operational readiness and risk tolerance.

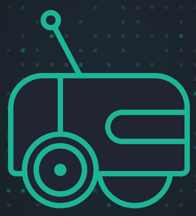

STAGE 1 (AGES 3-5)

Machine Learning & Correlation

Like a preschooler learning to speak and recognize familiar patterns, AI at this stage can identify relationships in data and report simple truths, but it’s not ready for complex reasoning. In SOC operations, this translates to automating the drudgery of collecting, correlating, and normalizing security data.

SOC Application:

01 Aggregating alerts from SIEMs, EDRs, and cloud tools.

02 Deduplicating noisy signals to reduce analyst workload.

03 Highlighting basic anomalies for further review.

Dropzone AI Example:

Automatically ingests alerts from security tools, deduplicates them, correlates related events, and enriches IOCs, freeing analysts from manual tasks and enabling faster triage.

STAGE 2 TEENAGER YEARS

Generative AI

Now capable of following instructions and producing detailed responses, AI at this stage can summarize, explain, and format information in useful ways, but still requires oversight to ensure accuracy. In the SOC, this means taking structured or semi-structured data from investigations and producing reports that are clear, consistent, and ready for analyst review.

SOC Application:

01 Drafting investigation summaries for common alert types.

02 Populating ticketing systems with clear, standardized findings.

03 Creating templated reports that can be quickly validated and escalated.

Dropzone AI Example:

Generates decision-ready investigation reports that human analysts can review in minutes, enabling Tier 1 analysts to process more alerts per shift and focus on incidents that require judgment and escalation.

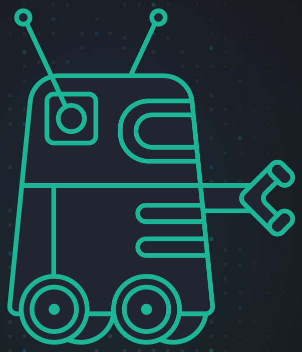

STAGE 3 COLLEGE YEARS

Agentic Workflows

At this stage, AI moves beyond static tasks to manage multi-step workflows that feed the output of one process back into the next. Like a college student who can break down complex assignments into smaller projects, AI here can orchestrate several investigative steps in sequence, adjusting based on intermediate findings.

Agentic workflows accelerate investigations, particularly for complex alerts that would normally be escalated to Tier 2 analysts. They also ensure that investigative logic is applied consistently, regardless of human bandwidth or experience levels.

SOC Application:

01 Automating end-to-end alert triage, from enrichment to cross-tool correlation.

02 Recursively querying multiple data sources to verify suspicious activity.

03 Running concurrent investigations without analyst intervention.

Dropzone AI Example:

Dropzone AI's SOC analyst springs into action when an alert arrives, dynamically generating investigative queries, pulling evidence from across integrated tools, and adjusting its path as each new clue emerges. Within minutes, it delivers a detailed, decision-ready report, enabling security teams to act quickly and confidently. Human-in-the-loop design allows analysts to verify findings and refine the AI's approach.

STAGE 4 TRUSTED YOUNG PROFESSIONAL

AI Agents

Now functioning as a trusted virtual assistant, AI can execute sophisticated tasks from start to finish, checking in with humans only when necessary. In SOC terms, an AI agent at this stage can fully handle Tier 1 workflows, triaging alerts, gathering evidence, and even performing containment actions. At the same time, analysts focus on high-priority incidents and strategic initiatives.

Tier 1 and Tier 2 staffing needs can be dramatically reduced, shifting human effort to proactive threat hunting, incident response planning, and other strategic work. Human analysts transition from investigators to verifiers and strategists.

SOC Application:

01 Investigating and resolving routine alerts autonomously.

02 When confidence is high, perform automated containment (e.g., quarantining endpoints, disabling compromised accounts).

03 Maintaining context memory to ensure investigations align with organizational policies and priorities.

Dropzone AI Example:

Dropzone AI works like a tireless teammate, investigating every alert the moment it appears and scaling effortlessly when volumes spike. It builds each case with rich, organization-specific context from change tickets to travel schedules, and can even interview users, just like a human analyst would. The result is a clear, decision-ready report that gives human analysts everything they need to dismiss, escalate, or initiate the incident response process.

STAGE 5 INDEPENDENT EXPERT

AI Worker

This is the AI "adult," a semi-autonomous entity capable of solving problems in novel ways without being bound to predefined processes. In the SOC, this could mean an AI that follows known investigative playbooks, proposes new investigative methods, adapts instantly to emerging attack patterns, and develops optimized workflows humans might never consider.

Stage 5 raises important governance and trust considerations. While the AI's independence can unlock new defensive capabilities, human oversight must ensure its actions align with business risk appetite and compliance requirements. At this stage, the SOC shifts from managing alerts to managing the AI itself.

SOC Application:

01 Continuously improving processes for investigation, improving detections, or initiating threat hunts without explicit human instruction.

02 Proactively identifying new attack patterns across historical and real-time data.

03 Designing and executing entirely new investigation tactics for emerging threats.

Dropzone AI Example (Vision Stage):

In the future, Dropzone's AI SOC analyst will be able to work with an AI detection engineering agent and AI threat hunting agent to automatically act on investigation findings. Dropzone AI will be able to draft and test new detection and investigation workflows independently, adapting cyber threat intelligence and policies as threats evolve. This evolution will make it an even greater force multiplier for the SOC, scaling defensive capabilities at a pace that keeps ahead of attacker innovation.

Trust but Verify

The Human-in-the-Loop Journey

As AI progresses through its maturity stages, its role in the SOC shifts from helpful assistant to trusted operator. What changes even more is the human analyst's role from performing most of the investigative work to overseeing and guiding the AI's output. This evolution is not about removing humans from the loop but redefining where and how their expertise adds the most value.

Human oversight is constant in the early stages of AI maturity (**Stages 1–2**). Analysts verify every output to ensure accuracy, learn how the AI works, and identify where it needs guidance. At this point, AI is an apprentice capable of handling basic tasks but depends on supervision.

By **Stage 3**, with agentic workflows in place, the human role shifts toward quality control. Analysts review investigations, selectively focusing on high-impact alerts or cases where the AI flags suspicious activity without a confident conclusion of it being malicious. This approach balances trust with vigilance, allowing the SOC to scale investigative capacity without scaling headcount.

In **Stage 4**, as AI agents become reliable enough to handle full investigative cycles, humans primarily act as governors of autonomy. They define the thresholds for automated containment, decide when human review is mandatory, and update policies as business risks evolve. Here, the analyst's job becomes one of orchestration rather than execution.

Stage 5 introduces a new governance challenge. An AI worker capable of novel problem-solving requires strategic oversight, ensuring innovative approaches align with compliance requirements, organizational priorities, and acceptable risk levels. In this stage, humans become managers of AI talent, as they would manage any high-performing specialist, setting objectives and reviewing performance.

Dropzone AI in Practice:

- ☑ **Human-in-the-loop by design:** Analysts can review and validate conclusions and feed corrections back into the system to improve accuracy over time.
- ☑ **Configurable oversight:** Teams can give the AI instructions to guide its investigations and make the findings fit into established workflows and procedures. They can also set automated action thresholds starting with notifications and escalating to containment as confidence builds.
- ☑ **Context memory:** AI investigations are informed by organization-specific knowledge and environment details, ensuring alignment with internal policies without constant manual intervention.

The goal is not to eliminate human judgment but to apply it where it matters most: strategic decisions, complex risk assessments, and oversight of AI actions. "Trust but verify" is the operational philosophy that allows SOCs to embrace AI's speed and scalability without surrendering control.

As we'll see in the next section, this trust model is what makes it possible to shift from static defenses where attackers have the advantage to dynamic, adaptive defenses that can match the pace and unpredictability of AI-powered adversaries.

From Static Defense to Dynamic Resiliency

SOC strategies have been rooted in a static defense mindset for decades: build the walls high, fortify the perimeter, and patch any visible cracks. The underlying assumption was that attacks could be prevented outright if the defensive surface could be hardened enough. In reality, this model has always favored attackers. They only need to find one weak point, and they can take as much time as they need to find it.

In the era of AI-powered offense, the asymmetry becomes untenable. Adversaries are now probing defenses continuously at machine speed, using polymorphic tactics to evade detection and exploit vulnerabilities the moment they appear. Against a static defense, the outcome is inevitable eventually; something will give.

The alternative is to shift the SOC's mindset from prevention at all costs to resiliency by design. A resilient SOC accepts that some incursions will succeed but is engineered to detect malicious activity quickly, disrupt the attack chain before significant damage occurs, and adapt defenses on the fly. This approach turns the "attackers only need to be right once" mantra on its head: in a resilient model, defenders only need to be right once during the attacker's progression to stop the breach from becoming a business-impacting incident.

What Dynamic Resiliency Looks Like in Practice:

01 **Continuous monitoring of user and entity behavior** to identify deviations from established norms in near real time.

02 **Adaptive investigative workflows** that change as new evidence emerges, rather than following a fixed sequence.

03 **Layered detection and containment** capabilities that engage at different points in the attack chain, reducing reliance on a single point of prevention.

04 **Rapid reconfiguration of defenses** to close gaps discovered during attempted intrusions.

The Role of AI in Dynamic Defense

AI is the technology that is enabling this shift. Only machine-speed investigation and decision-making can match the velocity of machine-speed attacks. By detecting anomalies, triaging alerts, and initiating containment actions in minutes or seconds, AI allows defenders to operate in the same temporal space as their adversaries.

Dropzone AI in Practice:

- ☑ Performs real-time alert ingestion and triage across SIEM, EDR, and cloud security tools, ensuring no signal is overlooked, even when your team is busy.
- ☑ Gathers and analyzes historical access patterns, business context, and even conducts user interviews, enabling faster detection of insider threats or compromised accounts.
- ☑ Executes auto-containment actions such as quarantining endpoints or disabling accounts at machine speed, buying time for human analysts to focus on higher-order response planning.
- ☑ Operates 24/7, providing continuous, adaptive coverage without the fatigue or bandwidth limits of human-only teams.

By replacing static fortifications with adaptive, AI-driven defenses, SOCs can change the terms of engagement. Instead of reacting too late to fast-moving threats, they can meet AI-powered attackers on equal footing, shifting from a position of disadvantage to one of strategic parity.

This shift requires more than technology alone; it demands reframing the SOC’s mission. Rather than aiming for perfect prevention, the goal becomes rapid disruption and sustained resilience. In the next section, we’ll explore why making this shift is an operational improvement and a strategic imperative for any organization facing AI-enabled threats.

The Strategic Case for AI in the SOC

For many organizations, discussions about AI in the SOC still begin with a familiar refrain: "Will this help us reduce costs?" While efficiency is a welcome outcome, focusing solely on cost savings misses the point. The real question is: Can our SOC keep pace with AI-powered attackers without it?

Adversaries have already answered that question for themselves. Criminal groups and state-sponsored actors are weaponizing AI to automate reconnaissance, craft personalized phishing lures, develop evasive malware, and orchestrate multi-stage attacks in parallel. Their operational tempo has shifted from human speed to machine speed, and that gap is widening daily.

This is not a theoretical risk; it is a present operational reality. If the SOC's investigation and response cycles are measured in hours while the attacker's are measured in seconds, no amount of human skill can close the gap. The only viable path forward is to match adversary velocity with AI-driven defense.

Why This Is a Strategic Imperative, Not a Budget Exercise:

01 **Defensive parity:** AI levels the playing field, enabling SOCs to investigate, decide, and act at the same speed as AI-driven offense.

02 **Resiliency over perfection:** The goal shifts from preventing every attack to detecting, disrupting, and ejecting intruders before critical damage occurs.

03 **Capacity without compromise:** AI can scale investigative capacity instantly during alert spikes, maintaining quality without overextending human analysts.

04 **Risk-driven priorities:** AI frees human analysts to focus on complex, high-impact threats where judgment, context, and business awareness are essential.

Operational Proof Points with Dropzone AI:

☑ Eliminates the time that alerts sit in a queue without investigation. Dropzone AI starts alerts immediately after they hit your SOC queue.

☑ Reduces mean time to detect (MTTD) and mean time to respond (MTTR) by conducting autonomous, decision-ready investigations in under 10 minutes.

☑ Filters out false positives at scale, allowing analysts to spend more time on genuine incidents without hiring additional staff.

☑ Maintains investigation quality during surge conditions, ensuring critical alerts are not lost in the noise.

☑ Learns from each case via context memory, aligning investigative behavior with the organization's unique environment and priorities.

The Cost of Inaction

Failing to adopt AI in the SOC is not just a missed opportunity, but an acceptance of strategic disadvantage. Without machine-speed capabilities, defenders will continue to face an ever-growing backlog of alerts, delayed containment, and increased exposure to business-disrupting incidents. Worse, they will cede the initiative to attackers, reacting to moves that have already unfolded rather than shaping the engagement.

In short, AI in the SOC is no longer a "nice-to-have" future technology. It is a present-day requirement for organizations seeking to maintain operational security in the face of AI-augmented threats. The longer adoption is delayed, the more ground is lost, and in cybersecurity, lost ground is rarely regained.

The good news is that SOCs do not need to leap directly to full autonomy to gain strategic benefits. By following a phased adoption roadmap, organizations can capture value at every stage of AI maturity, building trust and capability in parallel. In the next section, we'll outline a practical path forward, showing how to begin leveraging AI in the SOC today while preparing for the threats and opportunities of tomorrow.

Adoption Roadmap for SOC Leaders

Adopting AI in the SOC doesn't have to be a single, high-risk leap into full autonomy. In fact, the most effective and sustainable approach is to build capability in phases, aligning AI maturity with your team's operational readiness and trust level. Each stage delivers tangible value on its own, so you can realize benefits immediately while preparing for more advanced capabilities over time.

This phased roadmap mirrors the AI maturity model introduced earlier, with each step strengthening your SOC's ability to operate at machine speed against machine-speed threats.

PHASE 1 (STAGE 1–2)

Efficiency Through Automation

Objective: Eliminate repetitive manual work, reduce noise, and speed up initial triage.

Deploy AI for data correlation and alert deduplication across SIEM, EDR, and cloud tools.

Use generative AI to **auto-generate investigation summaries** and populate ticketing systems for Tier 1 workflows.

Build trust through **selective human verification of AI** output at this stage.

Dropzone AI Fit:

- ☑ Ingests alerts from 70+ integrated tools, applies threat intelligence, and reduces false positives.
- ☑ Produces clear, decision-ready investigation reports in minutes, enabling faster validation.

PHASE 2 (STAGE 3)

Scalable Triage and Recursive Investigation

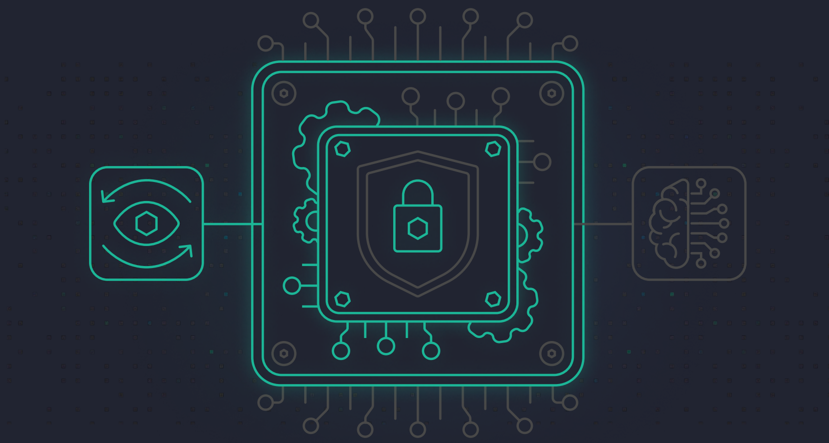

Objective: Increase investigative capacity without increasing headcount.

Deploy **agentic workflows** to conduct multi-step investigations and pull evidence from multiple tools without manual prompts.

Continue **selective human review**, focusing analyst attention on high-impact or uncertain cases.

Begin **trust-building through feedback loops**, refining AI performance based on analyst input.

Dropzone AI Fit:

☑ Dynamically generates investigative queries, pivots across data sources, and adapts its workflow as new evidence emerges.

☑ Maintains accuracy through human-in-the-loop oversight, learning from each review.

PHASE 3 (STAGE 4)

Trusted Virtual Assistance

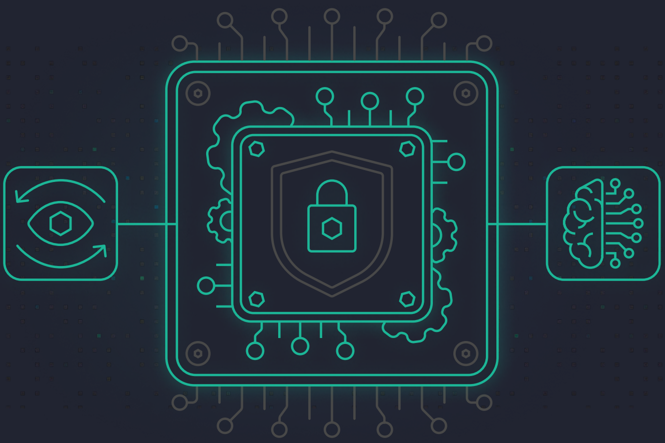

Objective: Free human analysts to focus on complex, high-priority threats.

Allow AI agents to **fully handle Tier 1 and Tier 2 investigations** from start to finish.

Configure **automated containment actions** for routine, high-confidence scenarios (e.g., quarantining endpoints, disabling compromised accounts).

Provide AI with direction to align investigations with **workflows and organizational priorities.**

Dropzone AI Fit:

☑ Executes configurable auto-remediation actions at machine speed.

☑ Takes direction through custom strategies and instructions to better support workflows and procedures.

PHASE 4 (STAGE 5)

Adaptive, Proactive Defense

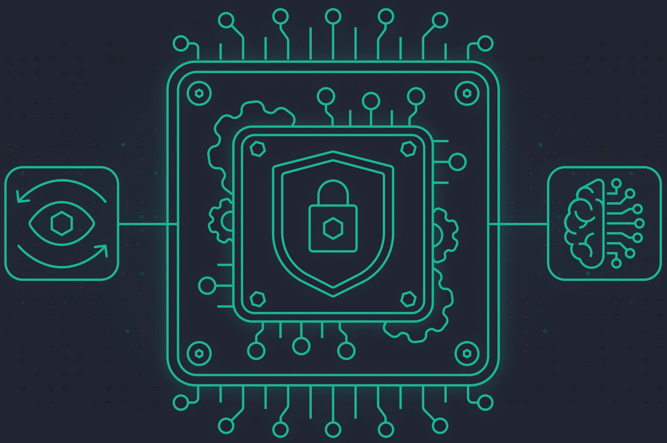

Objective: Leverage AI as a semi-autonomous partner capable of novel problem-solving.

Enable AI workers to design new investigation strategies for emerging threats.

Incorporate AI-generated process improvements into the SOC's operational playbook.

Establish **strategic oversight frameworks** to ensure alignment with compliance and risk tolerance.

Dropzone AI Vision:

☑ Future capabilities to tune detection rules, automate threat hunts, operationalize cyber threat intelligence, and prioritize misconfigurations and vulnerabilities for remediation.

Key Takeaways for SOC Leaders

☑ **You control the pace.** The roadmap allows you to capture immediate gains without overcommitting to full autonomy before you're ready.

☑ **Value exists at every stage.** Even initial deployments can reduce alert fatigue, speed investigations, and improve analyst morale.

☑ **Trust builds over time.** Human-in-the-loop oversight ensures that as the AI matures, so does your confidence in its judgment.

By following this phased approach, you can leverage AI in your SOC today, using it to close the speed gap with attackers while preparing your team for tomorrow's AI-driven defense strategies.

In the final section, we'll return to the child analogy to underscore the stakes: your AI, like your adversary's, is growing quickly. The question is whether yours will mature in time to meet them on equal footing.

Raising AI for Defense

Artificial intelligence in the SOC is not a static tool; it's a growing capability. Like a child, it learns in stages, gains independence over time, and becomes more capable as it matures. Whether in the toddler stage of simple pattern recognition or the young professional stage of trusted autonomy, each level of maturity offers real, measurable value to security operations.

But here's the reality: your adversaries are raising their own "AI children" too. They train them to probe your defenses, adapt their tactics, and exploit vulnerabilities at a pace no human team can match. The outcome will depend on which side's AI grows faster and how effectively it's guided.

The SOCs that thrive in this new environment will be those that:

01 **Adopt early** to capture immediate efficiency and visibility gains.

02 **Build trust incrementally**, using human-in-the-loop oversight to refine AI accuracy and decision-making.

03 **Evolve their strategy** from static prevention to dynamic resiliency, matching attacker speed with adaptive defenses.

04 **Empower their people** by freeing analysts from repetitive work, enabling them to focus on complex judgment calls and proactive threat management.

With AI-powered attackers already in the field, this is no longer a "someday" conversation. It is a strategic necessity that requires leadership, foresight, and a willingness to change entrenched operational models.

Dropzone AI's AI SOC Analyst is designed to meet you at your current stage of maturity, delivering value today while preparing your SOC for the AI-driven challenges of tomorrow. Whether you start by automating basic alert triage or move toward full investigative autonomy, you'll build the capabilities your team needs to keep pace and win against machine-speed adversaries.

Your AI will grow. The question is whether it will grow fast enough, and with the right guidance, to stand toe-to-toe with the adversary. Now is the time to start raising it.

Reinforcements have arrived

Dropzone AI's agents make SOCs fast, scalable, accurate, and proactive. With Dropzone AI autonomously handling routine Tier 1 alert triage, organizations can spend less time on reactive security and more time on proactive security. The Dropzone AI SOC Analyst replicates the techniques of elite analysts and is trusted by more than 100 enterprises and MSSPs, including CBTS, Pipe, UiPath, Zapier. Learn more at www.dropzone.ai.

Request a Demo

Dropzone AI