Command Zero

AUTONOMOUS & AI-ASSISTED SOC

The Agentic SOC: Why Security Operations Must Reimagine Itself—And Fast

By Alfred Huger,
Cofounder & CPO at Command Zero

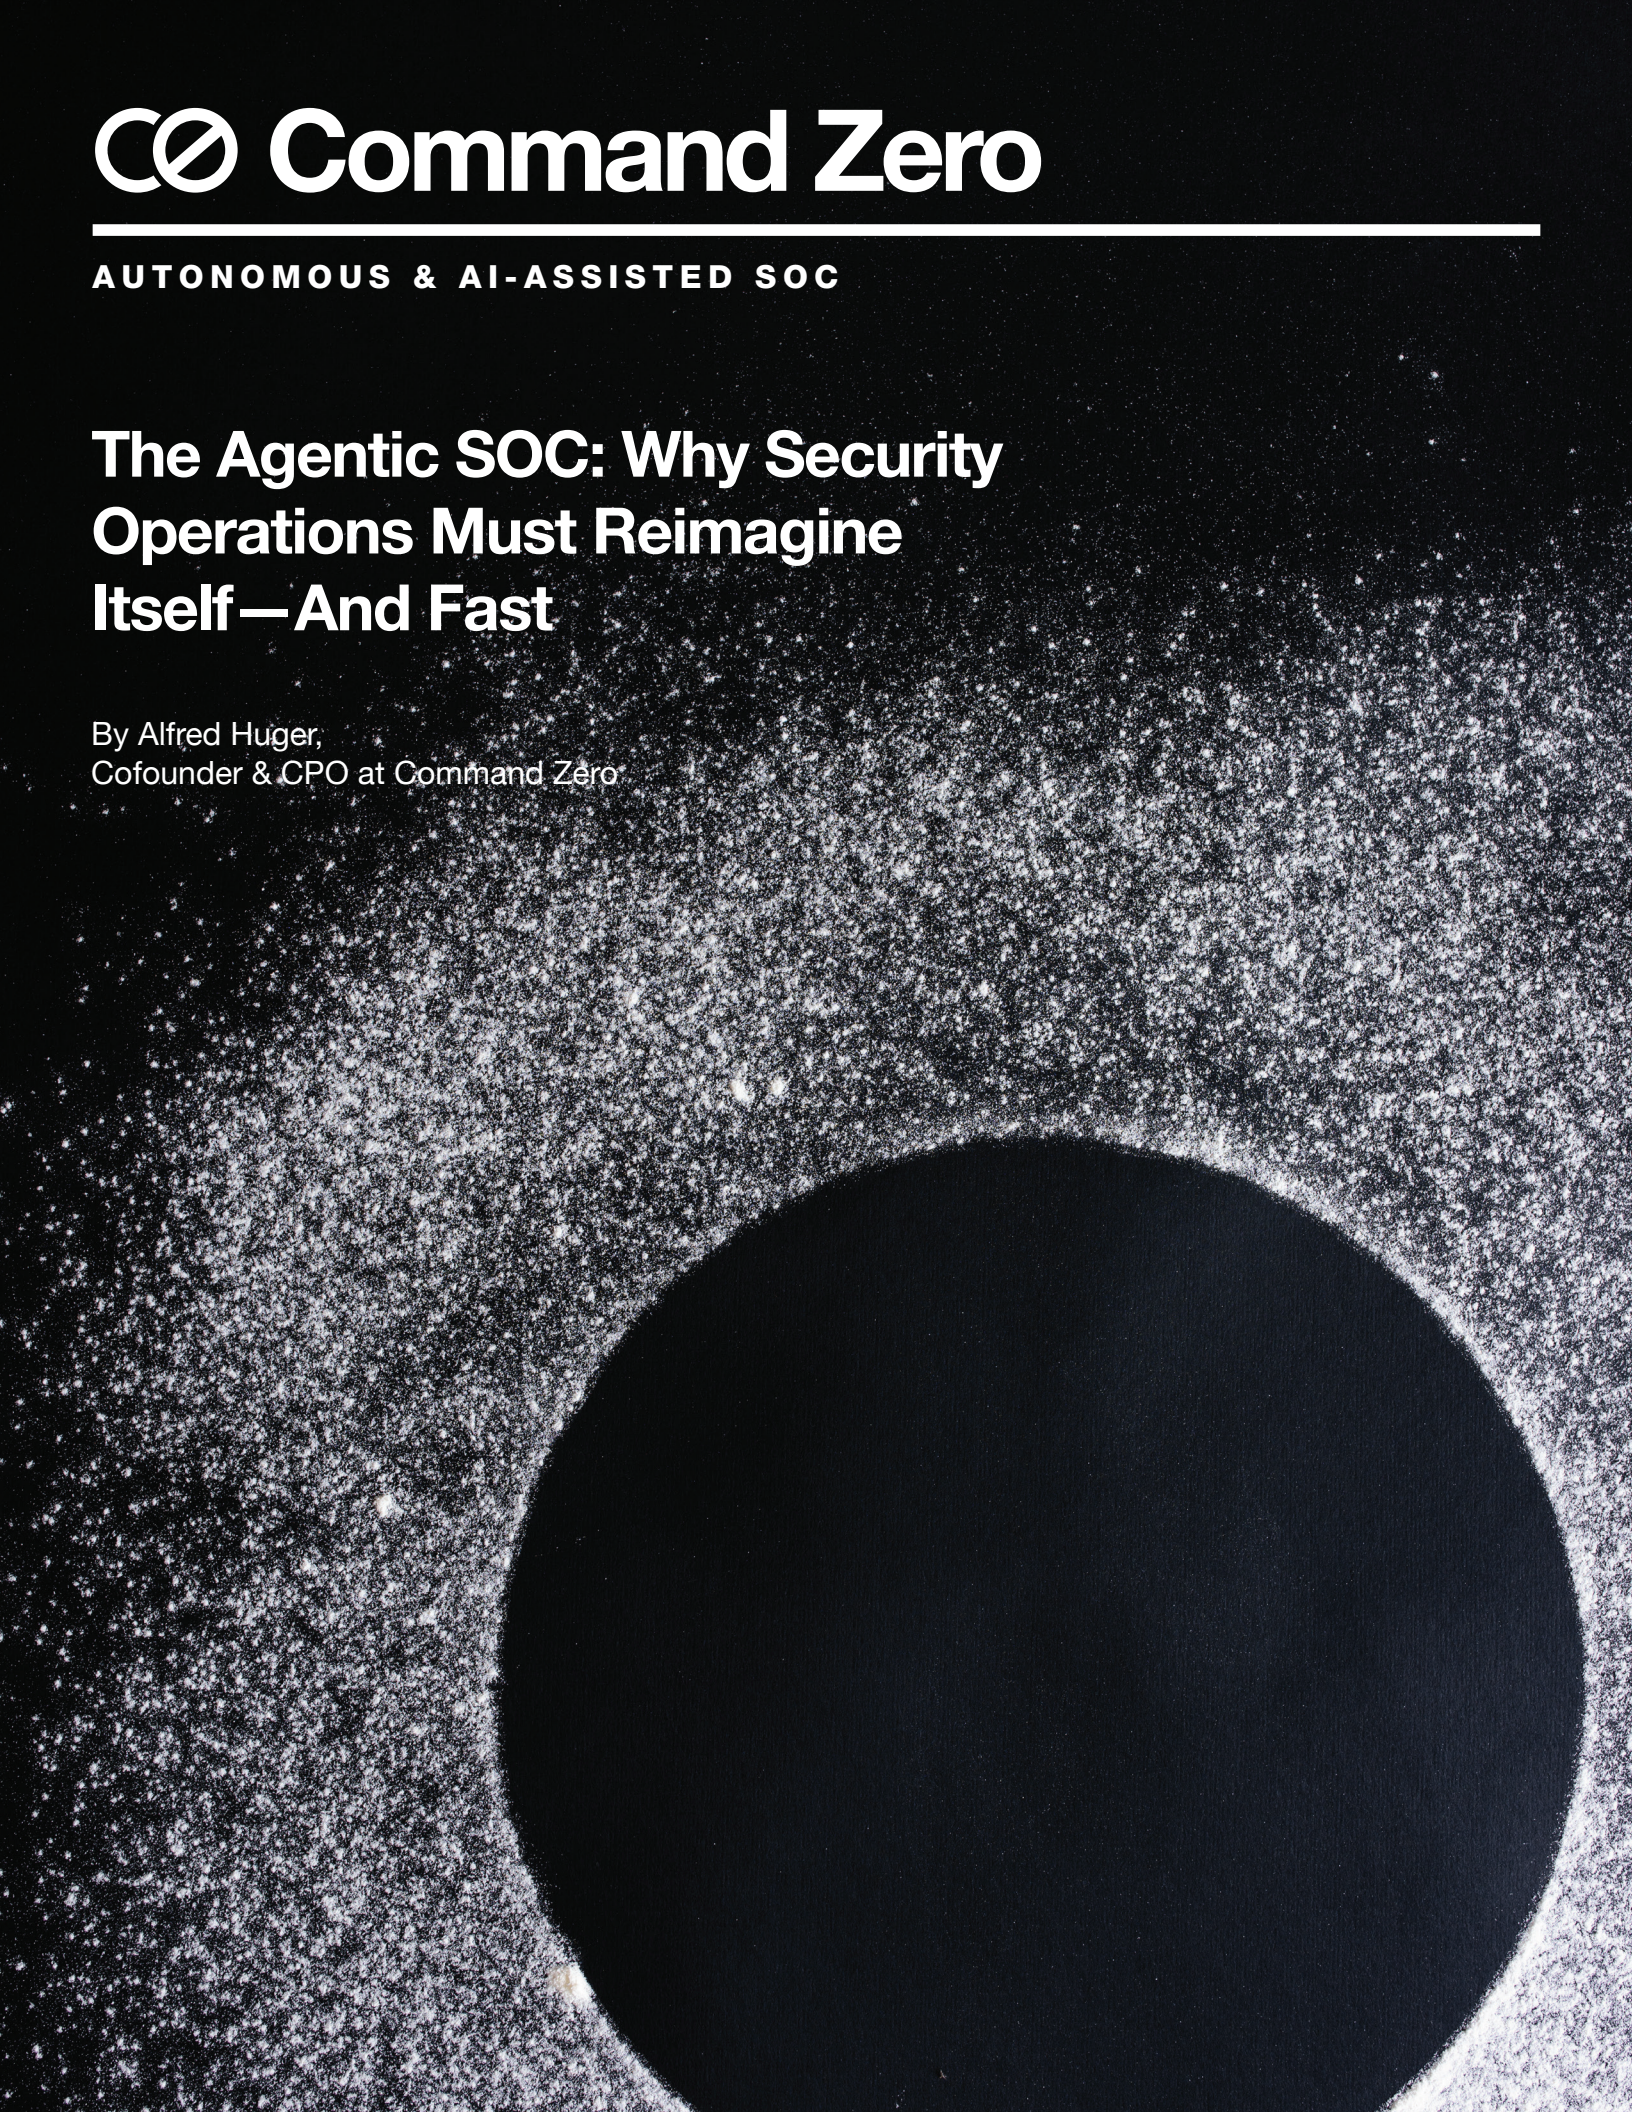

Foreword

Command Zero

After three decades building security software and leading product organizations through multiple successful acquisitions, I've witnessed many technological shifts in our industry. But what we're seeing today with AI agents in security operations represents something fundamentally different: a transformation that won't unfold over years—it's already happening at unprecedented speed.

At Command Zero, we're not simply adding AI features to existing workflows. We're pioneering a fundamental reimagination of how security operations centers function, and more importantly, we're doing it with principles that ensure these AI agents can be trusted, governed, and integrated into teams as reliable workforce members.

Alfred Huger,
Cofounder & CPO
at Command Zero

Why This Transformation Is Different—And Faster

The transition from traditional antivirus to EDR was a sea change that unfolded over several years. The shift to AI-powered security operations will happen far faster, and there are specific drivers forcing this acceleration.

Attackers are moving at machine speed. Adversaries are already adopting and will continue to adopt autonomous tool chaining for attacks. This means they're operating effectively at large language model speed and accuracy with unprecedented finesse. Because the SOC is the first point of contact for these adversaries, security teams need to respond with the same cadence—or get left behind very quickly. This reality becomes most evident in the SOC first.

Large language models are exceptionally well-suited for SOC work. Whether it's data extraction, data analysis, correlation, or summarization, these are all tasks that analysts perform daily as they work through their caseloads. LLMs play a substantial role in automating these functions while enhancing accuracy and consistency.

The market landscape is evolving rapidly. Competition in this space will be faster and more fierce than we've seen in previous technology shifts. Large language models make the competitive moat both more shallow and more narrow. This "rising tide raises all boats" reality applies equally to competitors and new entrants. While it's relatively easy to build a wrapper on top of an LLM, proving yourself as a capable product that delivers at scale remains a different conversation entirely—but that doesn't stop everyone from trying.

We believe this compression pushes us toward what I call "tier one compression." The ability to perform investigations on behalf of a tier one analyst will become a basic requirement—table stakes to compete in this space. However, this technology's applicability extends far beyond just triaging and managing tier one alerts. There's a much broader range of functions inside the SOC where AI agents can and should be deployed effectively.

The Non-Negotiable Principles for AI Agents in the SOC

Before diving into specific capabilities, it's critical to establish the principles that must guide how we deploy AI agents in security operations. These aren't aspirational guidelines—they're fundamental requirements for success.

Efficacy

Efficacy Still Matters Above All Else

At the end of the day, a large language model or any AI approach we bring to the SOC must be highly accurate, and it must maintain that accuracy as consistently as possible—in fact, more consistently than humanly possible. This is your "must be this tall to ride" threshold. Any investigation or work we perform on behalf of a SOC analyst must exceed the quality level they could achieve manually, and it needs to be highly precise every single time.

Agents Must Be Governable, Auditable, and Trainable

If we believe that AI agents represent the future workforce for the SOC, they must meet the same standards we apply to human employees. They must be:

- **Governable:** Operating within defined guardrails and constraints
- **Consistent:** Producing predictable, standardized outcomes
- **Auditable:** Providing transparent decision-making processes
- **Trainable:** Capable of learning and adapting based on feedback

If I'm going to give an agent in my environment an employee ID, I need to ensure I can hold them to the same standards and guardrails that I would apply to any team member. This principle applies whether the agent is investigating alerts or conducting outside-in investigations—the standard remains constant.

Speed and Affordability Are Essential

Solutions must be delivered at the speed required to matter and at a cost structure that's accessible to customers. If these solutions fall outside the grasp of a customer's affordability or if they're too slow to provide timely value, that fundamentally impacts both trust and adoption inside the SOC.

These principles guide everything we build at Command Zero. They're not afterthoughts—they're foundational to our approach.

Three Primary Use Cases: Where AI Agents Deliver Value Today

Security operations environments remain fundamentally event-driven. That's the operational reality, or what I call realpolitik. For most of our customers, this means working with alerts—inbound notifications that, if they can't be dismissed early, require investigation. In fact, they may need to be investigated to be dismissed. This represents the first critical role where AI agents deliver substantial value.

1. Alert-Based Investigations: Building Cross-Ecosystem Narratives

Our focus when ingesting and investigating alerts is to take that initial alert from ground truth—whether that's CrowdStrike Falcon, Microsoft Defender or Palo Alto Cortex at the endpoint or any other detection platform—and paint a comprehensive narrative about how that alert intersects with the rest of your security estate.

Our job isn't exclusively to validate whether the detection platform is right or wrong.
Instead, we determine what that alert means in the broader context of your environment and whether there's something you need to act on. When our AI agent ingests these alerts, we take a deep look into the alert itself, identify where else in your ecosystem the agent needs to investigate to build a complete picture, and then execute that investigation autonomously.

Common challenges with traditional alert investigation include:
• **Time constraints:** A single thorough investigation can consume an entire analyst's day
• **Context fragmentation:** Critical information scattered across multiple systems
• **Scope limitations:** Investigations often stop at the boundary of a single security tool
• **Inconsistent depth:** Time pressures force superficial analysis

Command Zero transforms alert investigations by providing cross-ecosystem visibility. If you have a compromised endpoint, our job is to tell you what that means to the rest of your estate—in Microsoft 365, SharePoint, EC2, Okta, or any other connected system. We're not simply summarizing alerts from an endpoint product or any other platform we extract alerts from (and we support several dozen platforms). We're investigating the full scope of potential impact.

This approach delivers investigations that are ready for review by tier one through tier N analysts. Everyone in the escalation path receives the data they need to work with, already compiled in a comprehensive dossier with high confidence verdicts. The investigation doesn't need to be restarted or reconstructed as it moves through your team—it evolves and expands as needed.

2

2. Bespoke Investigations: Beyond the Playbook

In simple terms, a bespoke investigation is one that doesn't start from an alert. These investigations tend to be the most complicated for team members because they're not driven by something that lends itself to a standard playbook. They may represent scenarios the analyst has never encountered before.

Examples of bespoke investigations include:

- Outside-in investigations based on intelligence or threat reports
- Insider threat scenarios (for example, investigating activity on an employee's last day)
- Proactive threat hunting based on emerging TTPs
- Compliance-driven reviews requiring comprehensive activity analysis

Our focus is enabling analysts to co-investigate with the AI agent, allowing it to handle data extraction, mapping, timelining, reporting, and navigation. This empowers analysts of any skill level to work through complex investigations while still delivering a conclusive verdict at the end—the same standard we apply to alert-based work.

The agent doesn't replace the analyst's judgment. Instead, it removes the drudgery of manual data collection and correlation, keeping analysts focused on high-value analysis and decision-making.

3. Functional Hunting: Democratizing Advanced Capabilities

If we believe that tier one will compress and that everyone effectively becomes tier two and above (or an orchestrator of agents), we need to empower the entire team to perform functions that were previously the exclusive domain of a small number of specialists.

Traditional hunting requires expertise in query languages like CQL or KQL, creating a significant barrier to entry. **We're not displacing those expert hunters,** but we are working to make their knowledge accessible to the whole team through what we call "functional hunting."

Functional hunting enables any team member to execute sophisticated queries through natural language or pre-built functions:

- Review all locked-out users across the environment
- Examine all Azure consent grants
- Identify clipboard activity theft in systems with EDR deployed
- Hunt for specific TTPs without knowing the underlying query syntax

This democratization allows organizations to multiply the effectiveness of their hunting program without requiring every team member to become a query language expert. The platform translates intent into action, executing the complex queries on behalf of the user.

How It Works: Transparent, Governable AI Investigations

When Command Zero completes an investigation—whether it originated from an alert, a bespoke inquiry, or a hunting exercise—we return a fully completed investigation package to the analyst. This package is designed to provide everything needed to make an informed decision.

Investigation Reporting: Tier One Through Tier N

For analysts working triage, we provide an initial report that walks them through the investigation path and our confidence assessment. Critically, **our summary speaks about more than just the original data source.** If the investigation began with an endpoint alert, we'll detail how it spread throughout the environment—into identity systems, cloud platforms, email, and beyond.

The report includes:
• **Proposed verdict with supporting evidence:** All data that drove the agent toward its recommendation
• **Business context:** Environmental information that shaped the investigation
• **Analyst annotations:** Comments attached to artifacts from previous investigations
• **False positive analysis:** Tests we ran to ensure the finding isn't a false positive
• **Discard rationale:** Everything we considered but ruled out, with explanations

This level of transparency is critical because **in the course of any investigation, an analyst must be able to stand behind their work. An AI agent must meet the same standard.**

If the analyst has sufficient information to reach a conclusion—and in most cases they do—they can accept the verdict, close the case, and move on. But they can also escalate it. When escalated, they invite another team member into the investigation, and that person receives access to more comprehensive reporting including IR-level details, complete log-level extraction, comprehensive timelines, and all supporting analysis.

The Investigation Console: Auditability and Extension

A key differentiator for Command Zero is that analysts can step directly into the investigation to see everything the agent did on their behalf. Our investigation console displays:
• The artifacts the agent chose to extract from the original alert
• The data sources it accessed to build the narrative
• The specific questions it decided to run against each data source

These questions are critical to our governance model.
They're built by Command Zero as high-impact, high-yield content that we ship daily or weekly, depending on the data source and the evolving threat landscape. We also enable customers to build their own questions by integrating detection content from platforms like Next-Gen SIEMs, importing that detection logic as questions available to both the agent and human analysts.

This approach is extremely important because it offers the agent tools that you control. These questions define what an agent can do, what inquiries it can make—and of course, you can expand that. But it provides governance, auditability, and consistency. You know with absolute certainty what your agents are capable of asking.

Plans and Playbooks: Codifying Investigation Logic

Whenever an agent completes an investigation, it can publish its investigative plan for review and modification. This means that for investigations of a certain type, if you want them handled a specific way—or if you wish to add constraints, broaden scope, ask additional questions, or enroll new data sources—that's entirely controlled by your team.

You can enshrine and codify these modifications so that your agent behaves exactly as you want every single time. This is what true governance looks like: not a black box making autonomous decisions, but a transparent, trainable system that operates according to your defined standards.

Cross-Ecosystem Investigation: The Microsoft 365 Example

A typical investigation might start with a hostname extracted from CrowdStrike Falcon. As the agent builds the narrative, it identifies an associated identity and expands scope into Microsoft Entra ID, pulling out identity information. From there, it investigates Exchange email, SharePoint access, and additional Entra ID activity.

This cross-ecosystem capability is where Command Zero truly excels. We can extract messaging headers, message traces, login data—all the elements necessary to paint a complete narrative when an event bleeds off the endpoint or conversely, starts in Microsoft 365 and moves to the endpoint. The principle remains the same in both directions.

This data isn't static—analysts can interact with it. They can review information from these automated questions, and if they want to extend the investigation, they can select any artifact and add it to their investigation, beginning to ask new questions about it. This allows investigators to pivot across all enabled data sources, including pulling content from Next-Gen SIEMs and executing new queries.

Once analysts extend the narrative, build additional timeline entries, and gather more evidence, they can direct the agent to reconsider its original reporting and verdict based on the new data. This isn't required—our assumption is that we provide the right data upfront every time—but when deeper investigation is warranted, the capability exists.

Collaborative Investigation

Complex investigations often require team collaboration. Command Zero enables analysts to invite team members into investigations, share notes with colleagues and with the AI agent. When leaving notes for co-workers, analysts can simultaneously instruct the agent on how to incorporate this new information when rebuilding reports or reconsidering verdicts.

Enrichment and Business Context: Making Every Investigation Smarter

As data streams into the platform, it's automatically annotated with business context and enrichment. An IP address, for example, can be enriched using any integrated product—Reversinglabs, IPinfo, IPdata, VirusTotal, CrowdStrike—the specific service doesn't matter. The key is that enrichment happens at the pipeline level, and once enriched, that data stays enriched for every investigation it appears in.

The same principle applies to notes about identities, annotations left by previous investigators, and outcomes of prior investigations. The more data in the system and the more investigations performed, the higher the accuracy rate of agent outcomes. Each investigation makes the next one smarter.

Watch Lists: Triggering Contextual Investigation Paths

Watch lists represent a powerful feature for contextual investigation. When Command Zero identifies an identity on a watch list during an alert investigation, it honors the original investigation but adds additional questions based on why that identity is being watched.

For example, if an administrative user triggers an endpoint alert, we execute the standard endpoint investigation—but because we know they're an administrator in SharePoint, we immediately push the investigation into Microsoft 365 to check for administrative activities. If the user is a researcher with GitHub access, regardless of where the alert originated, we investigate their GitHub activity for unusual commits or repository cloning.

Organizations can automatically populate watch lists with highly privileged users or new joiners to the company. They can also create custom watch lists for users on performance improvement plans, executives, frequent travelers, or any other cohort requiring special investigation attention.

Business Context: Passive Intelligence That Shapes Decisions

While watch lists actively trigger investigation paths, business context provides passive intelligence that shapes how the agent interprets findings.

This might include:
• Identifying a user as a red team member (contextualizing certain suspicious activities as authorized)
• Flagging an IP address as a sanctioned VPN head end (preventing unfamiliar travel alerts from being treated as true positives)
• Noting specific systems or accounts with expected unusual behavior patterns

This distinction is important: watch lists do things—they drive actions and expand investigation scope. Business context sits there—it's information that provides nuance to the agent's decision-making process.

Key Benefits for Security Operations Teams

By incorporating Command Zero into their workflow, security operations teams gain:

Accuracy: Thorough and precise investigations that exceed manual analysis quality, with consistent application of investigative methodology across all cases

Speed: Investigation time reduced from hours to minutes per case, with autonomous completion of data extraction and correlation
Confidence: Analysts empowered to close cases with certainty, backed by comprehensive evidence and transparent decision-making processes

Consistency: Standardized investigation processes across the entire team, regardless of individual analyst skill level

Scalability: Ability to handle significantly higher alert volumes without proportional staff increases
Knowledge Sharing: Investigative logic captured in questions and playbooks that serve as training resources and ensure best practices are applied universally

Benefits

Competitive Landscape: Transparency and Depth Matter
When comparing approaches to AI-powered security operations, several key differentiators emerge.

Beyond Tier One Focus: While some solutions concentrate exclusively on tier one triage with limited extensibility, Command Zero provides comprehensive support for tier one through tier N investigations. Our platform doesn't just tell you whether to escalate—it provides the complete investigation package needed at every level.

Transparent vs. Black Box Reporting: Early-stage solutions often deliver black box reporting—you receive a conclusion without visibility into how it was reached. Command Zero provides transparent, interactable data where analysts can see every question asked, every data source queried, and every piece of evidence considered. This transparency is essential for trust, auditability, and continuous improvement.

Ecosystem Breadth: Platforms that excel at data they own and control may struggle where they don't have native coverage. Command Zero extends investigations across your entire security estate—particularly into identity providers, SaaS platforms, and cloud environments where endpoint-centric solutions have limited visibility. Our ability to investigate Microsoft 365 environments, for example, provides capabilities that complement rather than compete with endpoint detection platforms.

Governance and Control: Solutions built as simple LLM wrappers lack the governance mechanisms required for enterprise deployment. Command Zero's question-based architecture, plan modification capabilities, and auditable decision trails provide the control necessary to deploy AI agents as trusted workforce members.

Example

Integration Example: Extending CrowdStrike Falcon's Investigation Capabilities

For organizations using CrowdStrike Falcon, Command Zero provides a powerful extension that leverages Falcon's strengths while addressing coverage gaps.

CrowdStrike excels at the data plane it owns—endpoints where it has direct instrumentation and control. Command Zero recognizes and respects that strength. Our role is to extend that visibility to where Falcon doesn't own the data and doesn't have control—particularly into Microsoft 365, Okta, Entra ID, AWS, and other environments critical to understanding attack scope and impact.

When a Falcon alert triggers, Command Zero doesn't just validate the endpoint finding. It investigates what that endpoint compromise means across your entire estate. **We can investigate inside Microsoft's environment like a hot knife through butter,** correlating endpoint activity with email access, file sharing, identity changes, and administrative actions—all the context needed to understand true impact and make informed containment decisions.

This complementary approach represents a primary use case we see with CrowdStrike customers: maintaining the strength of their endpoint detection while gaining the cross-platform investigation capabilities required for modern, distributed enterprise environments. Command Zero's holistic approach complements and improves existing security investments. Security Operations teams can cover all alerts at predictable quality, high speed and gain proactive capabilities to tackle threats at scale.

Path Forward

The Path Forward: Security Operations AGI

Command Zero's vision extends beyond solving specific investigation types. We're on a journey toward what I call "Security Operations AGI"—artificial general intelligence specifically designed for security operations. The question isn't whether we can build AI-powered security tools; it's what we build along that journey and how we ensure each step delivers tangible value while maintaining the principles of governance, transparency, and effectiveness.
The platform we've built today focuses on the use cases where AI agents provide clear, measurable value:

- Autonomous investigation of alerts from detection platforms
- Guided co-investigation for complex, unplaybooked scenarios
- Democratized hunting accessible to every team member

As the threat landscape continues to evolve and as attackers increasingly operate at machine speed, security operations teams need solutions that can match that pace while remaining under their control. **The future isn't about AI replacing analysts—it's about AI agents working as governable, auditable members of the team,** handling the data drudgery while keeping human expertise focused on high-value decision-making.

This transformation is already underway, and it's happening faster than previous technology shifts in our industry. Organizations that adapt quickly, adopting AI agents with the right governance and transparency principles, will maintain defensive parity with attackers operating at LLM speed. Those that don't, will find themselves perpetually behind.

Command Zero

Command Zero

AUTONOMOUS & AI-ASSISTED SOC