Bricklayer AI

The New SOC Operating Model

How Security Teams Scale Investigations, Decisions, and Responses When Operational Capacity Is No Longer the Constraint

Executive Summary

Scaling security has always been a tough puzzle. The catch is whether you fixed the actual root cause.

For the past decade, the industry has poured time, effort, and other resources into threat detection and response (TDR). More data sources, more correlation rules, more SIEM tuning, more EDR coverage. That investment was necessary, and it paid off: security teams now have more visibility into their environments than at any point in history.

But visibility didn't scale the SOC. It gave the SOC more work.

Every new data source generates more signals to evaluate. Every new detection capability produces more alerts to investigate. The teams responsible for turning that signal into decisions didn't grow at the same rate. The nature of their work is contextual, adaptive, and judgment-intensive in ways that resisted the kind of automation that worked well for repetitive tasks.

The result is the situation most security leaders recognize immediately: better tools, more data, but a team that is the limiting factor.

This operating model was built for a world where human execution is the only available engine. Where operational capacity is genuinely scarce. Where every investigation requires an analyst to drive it forward, step by step, from alert to conclusion.

AI agents introduce a new reality. Not AI as a productivity layer on top of existing workflows, but AI as a genuine contributor to investigative work: gathering evidence, assembling context, coordinating across systems, and progressing investigations at a pace that doesn't require human involvement at every step.

When investigative work is governed appropriately, the resource constraint changes. The challenge is no longer completing enough work; it becomes directing a security operation capable of completing significantly more, with the right controls, the right context, and human oversight applied where it matters most.

But there is something larger at stake than investigation throughput. Today's security organizations don't just have a capacity problem. They have a fragmentation problem. Alerts, investigations, threat intelligence, vulnerability findings, and threat hunting are all happening simultaneously, across disconnected tools, teams, and workflows. Context gets created and then lost. Decisions happen in isolation. Work that should inform other actions doesn't always do so.

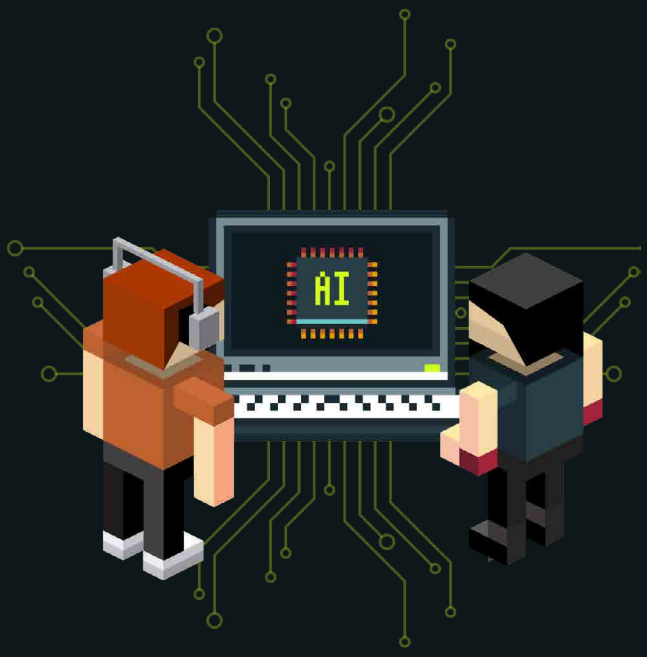

The agentic SOC isn't just about doing more, it's about connecting systems and building a security operation that functions as a coherent system.

Chapter 1

The SOC Scaling Problem Has Changed

The SOC faces a simple reality: there will always be more security work than people available to perform it.

Over the past decade, organizations responded by investing heavily in technologies designed to improve visibility and control. SIEM platforms expanded the ability to collect and analyze security data at scale. Endpoint security solutions improved detection fidelity. Cloud security and identity tools provided visibility into environments that had previously been opaque. Threat intelligence platforms aggregated external context that analysts could use to enrich their findings.

These investments solved real problems. Most organizations today have significantly more information about their environments than they did five years ago. Security teams can identify suspicious behavior, investigate activity, and respond to threats with a level of visibility that was previously impossible outside the most well-resourced programs.

But greater visibility created a new operational challenge the industry hasn't reckoned with.

Every new data source creates additional information to review. Every new detection capability creates additional signals to evaluate. Every new technology platform introduces new context that analysts need to understand before they can make sense of what they're seeing. The SOC didn't just get better visibility into threats, it got better visibility into everything. This means far more to process, evaluate, and act on.

The challenge for many security organizations is no longer finding information, it is using that information to act quickly enough to matter.

A single alert may require an analyst to review identity logs, endpoint telemetry, network events, threat intelligence feeds, asset context, historical activity, and open vulnerability findings before determining whether the organization faces a meaningful risk. That review is not optional. Skipping steps is how real incidents get missed. But performing every step thoroughly, for every alert, at the volume modern environments generate is something no team has the capacity to sustain.

This work requires judgment. It requires experience. It requires time. And it requires the analyst to hold a significant amount of context in working memory, manually connecting threads across systems that were not designed to talk to each other.

This is why experienced security leaders often describe a frustrating paradox: their teams have better tools than ever, yet investigations feel harder and more time-consuming than they should be. The visibility problem got solved. The investigation problem did not.

Chapter 2

What More Analysts and More Tools Haven't Solved

When security organizations face increasing workload, the responses are predictable. Add more people. Add more technology. Automate more processes. Each response makes intuitive sense, and each has delivered real value in specific ways.

But none of them change the fundamental operating model.

Hiring more analysts increases capacity, but it also introduces compounding constraints. Experienced security professionals are genuinely difficult to find and retain, and competition for them is intense across every industry. New analysts require months of training and contextual immersion before they can independently manage complex investigations. Senior analysts often become the bottleneck at both ends, managing their own investigations while simultaneously bringing junior team members up to speed. Adding people helps, but the return diminishes as team size grows and coordination overhead increases.

Technology investments create a parallel challenge. The average enterprise security team operates across a sprawling ecosystem of tools: SIEM, EDR, identity, cloud security, vulnerability management, threat intelligence, case management, and more. Each platform was procured to solve a specific problem, and each solves it reasonably well in isolation. But analysts working across their ecosystems spend a significant portion of their time on investigative overhead. They navigate between systems, re-enter context, translate findings from one tool's format into another's, and manually assemble a picture that no single platform can provide on its own.

Buying a better SIEM or a more capable EDR does not reduce that context-switching burden. It often increases it, because more capable tools generate more information that requires evaluation.

Automation has been the most durable response, and it has delivered genuine value. Organizations have automated enrichment workflows, alert triage, ticket creation, notifications, and specific response actions. These improvements have helped teams operate more efficiently and reduced the volume of purely mechanical work.

But automation has a specific ceiling, and most mature security organizations have hit it.

Automation excels when the workflow is predictable. If an alert fires, enrich the indicator, check the reputation score, open a ticket, notify the right team. That sequence can be fully automated and reliably executed at scale. Real investigations are rarely that predictable. They begin with incomplete information and change shape as new evidence surfaces. An analyst may begin investigating what appears to be a phishing attempt and discover it's connected to a compromised service account. They may find that an endpoint flagged for suspicious process activity is running an unpatched vulnerability being actively exploited. The investigation doesn't follow a fixed path because the threat doesn't follow a fixed path.

An automated workflow can execute a predefined sequence. It cannot decide what that sequence should be based on what it finds along the way. It cannot recognize when new evidence requires revisiting an earlier conclusion.

The result is that many organizations have successfully automated the parts of security operations that don't require much judgment, while leaving the parts that do almost entirely dependent on human execution. More efficient humans, certainly. But still humans doing the work that matters most.

Automation excels when workflows are predictable. Real investigations rarely are.

Chapter 3

Automation Improved Efficiency. It Didn't Change the Operating Model.

Efficiency and operating model change are often conflated in conversations about security automation. Confusing them leads organizations to expect scale from tools designed to improve productivity.

Efficiency improvements are real and valuable. Reducing the time an analyst spends on mechanical tasks returns hours that can be redirected toward higher-value work. For individual analysts and individual investigations, that improvement is meaningful.

But efficiency improvements operate within the existing operating model. They make the current approach faster. They don't change what the current approach assumes.

The current approach assumes that a human analyst will be the primary actor in every significant investigation. They determine what questions need answering. They identify what evidence is required. They evaluate findings, draw conclusions, decide on next steps, and coordinate response. The tools around them have gotten more powerful. The workflows have gotten more streamlined. But the investigation still belongs to the analyst from start to finish.

This assumption has shaped security operations for decades, and it made sense for most of that time. Human judgment was the only mechanism capable of navigating the complexity and ambiguity of a real security investigation.

That is no longer completely true.

AI agents capable of performing adaptive, contextual investigative work introduce a new opportunity. They're not a replacement for human judgment, but rather a genuine extension of operational capacity that allows security teams to do more investigative work than human execution alone could support.

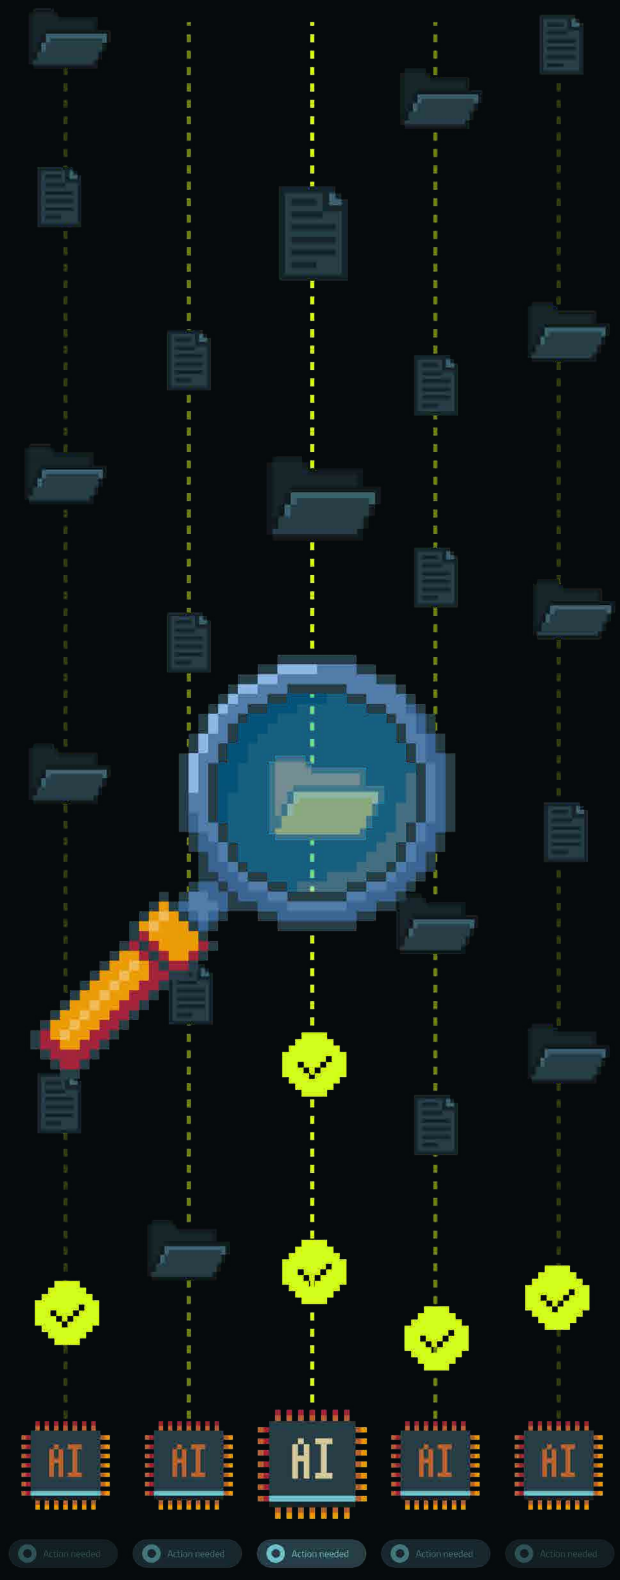

The question shifts from how to make analysts more efficient within the existing operating model, to what it looks like when the operating model itself changes.

Chapter 4

The Copilot Phase — Real Value, Defined Limits

AI copilots for security operations represent a meaningful step forward and provide genuine value. Analysts spend a significant portion of their time collecting information, searching documentation, summarizing findings, and drafting investigation notes. These activities are necessary, but they are not where experienced analysts create the most value. Copilots reduce that burden by making information faster to access, easier to interpret, and quicker to communicate.

For many organizations, this represents the first practical application of AI within daily security operations. Analysts can query their environment in plain language. They can get summaries of complex alert chains. They can access technical explanations without context-switching to external documentation. They can produce investigation reports in a fraction of the time.

These improvements matter. A faster analyst is more productive. A more productive analyst can do more.

But copilots have a defined ceiling. Understanding it matters for how organizations approach what comes next.

> A copilot improves the way an analyst performs work.
> It does not change the underlying dependency on the analyst to do it.

Analysts still drive the investigation. They still determine what questions need to be answered. They still identify what evidence is relevant. They still evaluate what the findings mean and decide what happens next. The copilot helps with some of the mechanics: finding information faster, summarizing more efficiently, drafting more quickly. But the human remains the operational engine behind every investigation.

This means copilots improve productivity without changing scale. If a team has the capacity to actively drive 40 investigations per day, a copilot might help those analysts do their 40 investigations faster and with less friction. That is genuinely valuable.

But the organization still cannot drive 400 investigations. The capacity ceiling remains essentially fixed.

Productivity and scale are different outcomes. They require different investments. Copilots remain a valuable part of the analyst experience and an important part of a larger system, but they are not the answer to the scaling problem. Organizations that treat them as such will find the underlying constraint remains where it was.

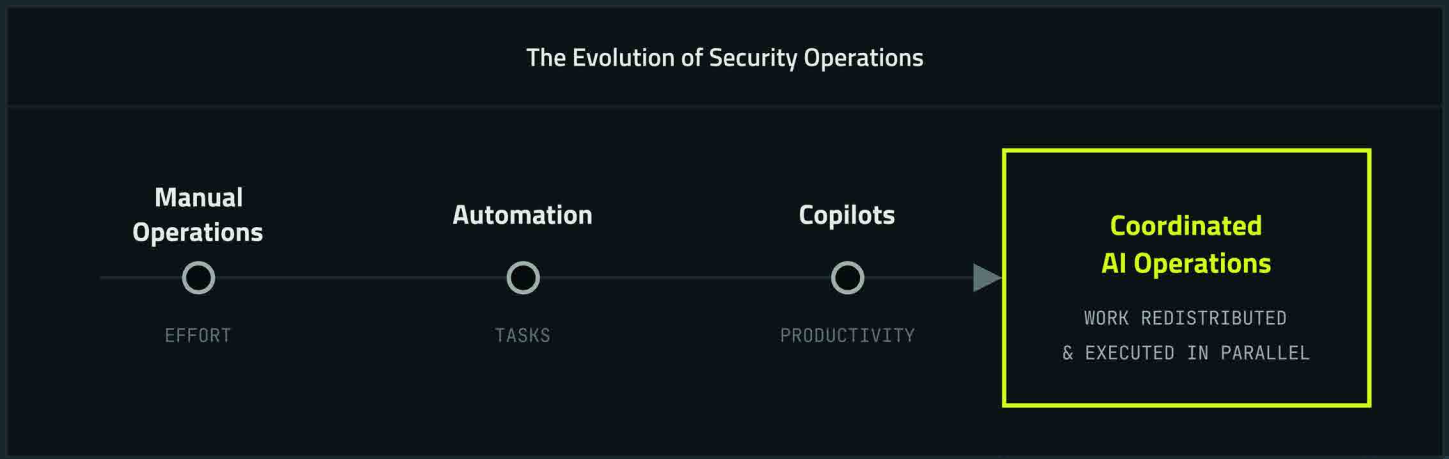

Chapter 5

The Fragmentation Problem Nobody Talks About

Before engaging with what a new operating model could look like, there is a problem that needs to be named directly. It sits underneath the scaling problem, makes it significantly worse, and rarely gets addressed.Today's security teams aren't solving isolated problems. Alerts, investigations, threat intelligence, vulnerability management, and threat hunting are all happening simultaneously and all feed into each other. A threat intelligence finding should inform how an alert gets investigated. A vulnerability discovery should shape what threat hunting looks for. An incident investigation should feed intelligence back into detection engineering. A pattern of low-severity alerts should surface into risk conversations that would otherwise never connect them.

In practice, almost none of those connections happen reliably.

The tools are different. The teams are often different. The workflows are definitely different. An alert lands in the SIEM and gets triaged by a SOC analyst. Threat intelligence often lives in a separate platform managed by a separate team. Vulnerability findings can sit in a system that the incident responder rarely has access to during an active investigation. Threat hunting is conducted episodically, often disconnected from what's currently active in the alert queue.

Context gets created and then lost. Decisions happen in silos. Work that should inform security decisions doesn't, because there is no structured mechanism for it to flow.

This fragmentation is not an accident. It is the natural result of procuring best-of-breed tools for specific functions over many years. Each tool solved its specific problem well. What nobody optimized for was the connections between them.

The risk of relying on all of these point solutions is real and largely invisible. Every time an analyst investigates an alert without access to relevant threat intelligence, they are making a less informed decision than they should be. Every time a vulnerability finding doesn't reach the analyst actively investigating a related endpoint, the investigation takes longer and potentially misses something important. Every time a threat hunt concludes without its findings being incorporated into detection logic, the value of that hunt is only partially realized.

Agentic SOCs have to address this segmentation directly. Not just by doing existing work faster, but by building a security operation where information flows across functions, context carries from one workflow to the next, and the work actually connects. The SOC needs to become something different than it is today.

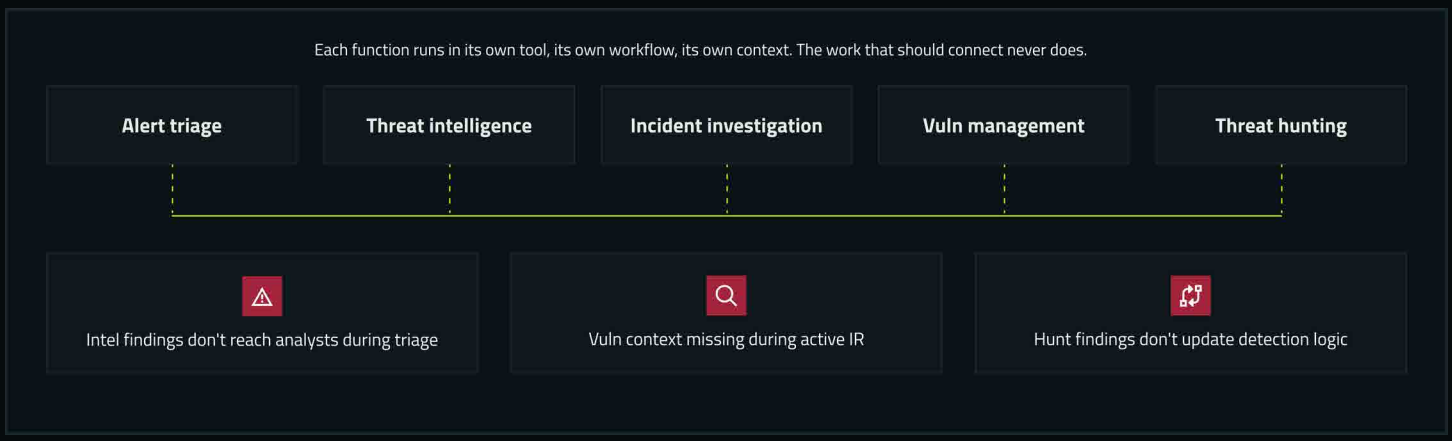

Today's security operation functions often happen in silos, even though they aren't solving isolated problems. The connections between them exist in theory. In practice, they break down.

Chapter 6

What Changes When Operational Capacity Increases

For most of the history of security operations, the primary constraint has been execution capacity. There are only so many investigations an analyst can complete in a shift. There are only so many incidents a team can actively manage simultaneously. Only so many workflows that can be reviewed, adjusted, and improved in parallel with active operations.

This reality shaped how SOCs were designed. Escalation processes were built to ensure expertise was concentrated where it was needed most. Alert prioritization frameworks were developed so analysts could focus on what mattered. Triage processes were designed to move low-risk alerts through quickly and redirect analyst time toward investigations that warranted deeper attention.

All of that was rational given existing constraints. When human execution is the only available engine, you optimize how humans spend their time.

Agentic AI introduces a different possibility: AI systems that can contribute meaningful investigative work. Not by answering a question an analyst asks, and not by executing a predefined playbook, but by actively progressing an investigation. Gathering evidence across multiple systems simultaneously. Assembling context from disparate sources. Adapting their approach based on what they discover. Moving the investigation forward without requiring a human to drive every step.

When that capability is governed appropriately, the volume of investigative work the organization can perform expands significantly.

Alerts that previously sat uninvestigated because there wasn't bandwidth to get to them can be actively worked. Investigations can progress in parallel rather than sequentially. Evidence gathering doesn't wait for an analyst to be available.

This isn't simply faster execution of the same work. There are exponentially more amounts of work getting done. Threats that would have gone unexamined get examined. Context that would have been skipped in the interest of time gets assembled. Connections that would have been missed because no analyst had the bandwidth to look for, are surfaced.

With these changes comes something equally important: the ability to connect work across functions. Threat intelligence findings can be automatically incorporated into active investigations. Vulnerability context can be pulled into incident response without requiring a separate analyst to bridge the gap. Threat hunting activity can feed back into detection engineering continuously rather than waiting for a formal review cycle.

The fragmentation that defines today's SOC begins to resolve. Not because the tools changed, but because there is capacity to make the connections that were always necessary but never achievable at scale.

The security organization doesn't just operate more efficiently. It operates as a more coherent system.

Chapter 7

Humans Roles Evolve — and Become More Strategic

One of the most persistent concerns about agentic security operations is that as AI systems do more operational work, human security professionals become less relevant. This concern is understandable, but also groundless.

As operational execution becomes more scalable, the strategic value of human judgment increases. What changes is not the value of a human's skillset, it is how their judgment is applied and how much more they are able to accomplish.

In the traditional SOC, human attention is distributed across every step of every investigation. Analysts gather evidence, evaluate findings, determine next steps, coordinate with other teams, make decisions, and document outcomes. At scale, this means senior analysts spend significant time on mechanics rather than strategy.

The most experienced security professionals in the organization are frequently occupied with the operational overhead of keeping investigations moving, rather than the higher-order thinking that their experience makes them uniquely capable of.

In an agentic SOC, operational execution is increasingly handled by AI agents operating under defined policies and human oversight. That frees human judgment to operate at a different level, and to do so more consistently, because the mechanical burden of investigation execution is no longer competing for the same attention.

Four areas become significantly more important for security leaders and their teams in this model.

1. Defining objectives and priorities

Security leaders are responsible for determining what the organization is trying to achieve and translating that into clear direction for the systems performing the work. Which risks matter most given the current threat environment? How should investigative resources be allocated when multiple significant investigations are active simultaneously? What outcomes should the operation be optimizing for? This is genuine strategic work that benefits from experience and organizational context. In the traditional SOC, it often gets crowded out by operational urgency.

2. Establishing authority and guardrails

In a human-operated SOC, analysts understand what they are permitted to do based on training, experience, and institutional norms. That implicit understanding doesn't transfer to AI systems. In an agentic model, those expectations have to be explicit.
What actions can the system take without human approval?
Where is a human decision required before proceeding?
What limits must be maintained regardless of circumstances?
Building and maintaining that policy framework is substantive work that requires security expertise and organizational judgment.

3. Governing outcomes and risk

More operational capacity doesn't eliminate the need for human evaluation of what an operation produces. Security teams still need to assess business impact, evaluate risk in context, and verify that the system's outputs are aligned with what actually matters. The role shifts from executing every step to ensuring the system is working as intended and its outputs are defensible. Humans remain ultimately accountable for every decision the operation produces.

4. Continuously improving the system

Every investigation produces information that should make future investigations better. Every outcome reveals something about how well the current approach is working. In a human-operated SOC, institutional learning is largely informal. It lives in the heads of senior analysts and gets shared through mentorship and experience. In an agentic SOC, that learning can be systematically captured and incorporated, building organizational capability that compounds over time rather than departing when experienced analysts move on.

Human Expertise Moves in the Agentic SOC.

Humans remain fully accountable for security outcomes. They direct, govern, evaluate, and continuously improve the system that performs the operational work. Their expertise is applied to the problems that actually require it, rather than being consumed by the mechanics of keeping investigations moving.

Human Expertise Moves

Traditional SOC Where analyst time goes	Agentic SOC Where analyst time goes
Pulling evidence from multiple systems	Setting investigation priorities
Context switching between tools	Reviewing and approving agent findings
Writing up investigation notes	Improving detection and response policies
Manually connecting intel to active alerts	Evaluating risk and business impact
Repeating triage steps for similar alerts	Building and refining governance policies

Human expertise doesn't disappear in the agentic SOC. Instead, the strategic value of human judgment increases.

Chapter 8

Governance is What Makes Scale Trustworthy

When operational work becomes more scalable, governance becomes more important.This is worth stating plainly, because the instinct in many technology conversations runs the opposite direction. Governance gets framed as the obstacle to moving fast, and better technology gets positioned as the way to reduce the friction of oversight.

That instinct is wrong for security operations specifically.

Security decisions have consequences. An investigation that concludes incorrectly can leave a real threat unaddressed. A response action taken without full context can disrupt legitimate business operations. A containment decision made prematurely can create more problems than it solves. These risks exist both in a human-operated SOC and an agentic SOC. The difference is the rate at which decisions and actions can occur.

When AI agents are performing investigative work across hundreds of concurrent investigations, gathering evidence, assembling context, making intermediate decisions about what to pursue next, and determining when findings warrant escalation, the volume of consequential activity happening inside the SOC increases significantly. Without appropriate governance, that volume creates uncertainty. Security leaders cannot be confident that the system is operating within acceptable boundaries, that decisions are being made with the right information, or that outcomes are aligned with organizational risk tolerance.

With governance, that same volume becomes an advantage. Security teams have visibility into what work is being performed and why. They can trace decisions, review conclusions, identify where the system's judgment diverged from what a senior analyst would have done, and improve the process over time. They have confidence that the system as a whole, and each agent within, is operating within defined authority and that the boundaries they've established are actually being respected.

This is the insight that experienced security leaders tend to arrive at quickly when they engage seriously with agentic security operations: governance is not something you impose on AI systems to make their use politically acceptable. It is the mechanism that allows you to trust what they produce. And trust is what allows you to expand what they do.

Organizations that approach agentic adoption by asking how to remove controls in order to move faster will create risk without confidence. Organizations that approach it by asking how to build enough control to operate confidently at a much larger scale will build something that actually delivers on the promise.

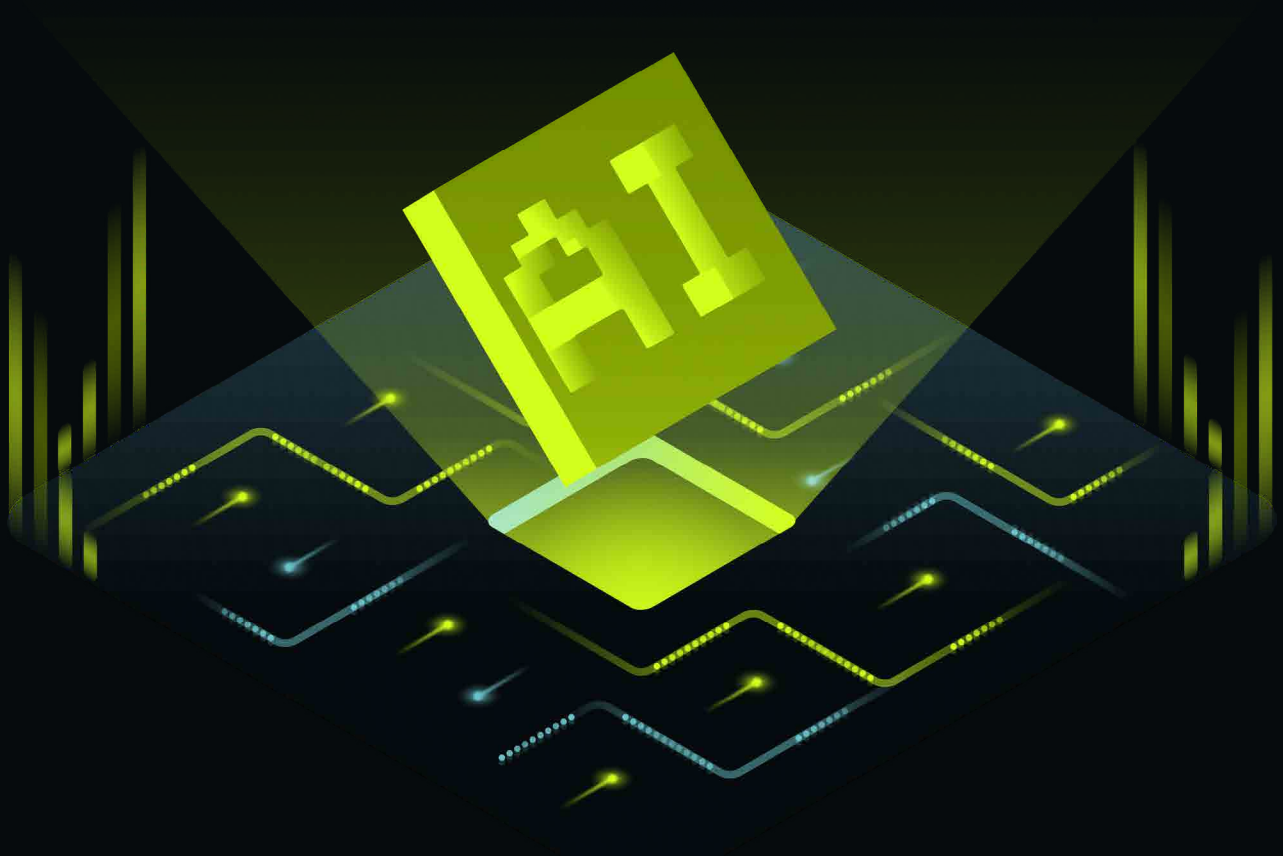

Chapter 9

Building the Agentic SOC

The transition toward an agentic SOC does not require replacing existing security investments or discarding what has been built. The SIEM, EDR, threat intelligence platform, case management system, and existing detection logic all remain important. The security ecosystem that organizations have assembled over the years continues to matter. What changes is how work flows across that ecosystem, who initiates and drives investigative activity, and how findings from one part of the operation inform what happens in another.

The most effective starting point is not attempting to transform every security workflow simultaneously. That only creates organizational disruption without demonstrating value quickly enough to sustain confidence in the program.

Instead, organizations should identify where additional operational capacity creates the greatest immediate impact and where the fragmentation problem is most costly. For most organizations, that points toward investigation-heavy workflows: alert triage, incident investigation, threat intelligence analysis, threat hunting, and vulnerability prioritization. These are areas that require significant context gathering, benefit from cross-functional information flow, and are currently constrained by the sequential, siloed nature of how the work gets done.

Agentic systems also differentiate most clearly from traditional automation in areas that require significant context.

The value of agentic execution is highest precisely where the workflow is adaptive and contextual, where what needs to happen next depends on what was just discovered, and where findings from one system need to be immediately connected to information from another.

This is also what separates agentic systems from traditional automation in practice. Automation requires someone to define the work in advance: the sequence of steps, the systems to query, the conditions to evaluate. An agentic system starts from a different point. The analyst defines the objective. The system determines how to achieve it.

Successful adoption also requires clarity about what success looks like before the program begins. Faster average time to triage. Broader coverage of alerts that previously went uninvestigated. More consistent investigation quality regardless of which analyst is on shift. Faster time to respond to confirmed incidents. Better flow of context from threat intelligence into active investigations. Improved connection between vulnerability findings and incident response. These are measurable outcomes that reflect real operational improvement, and defining KPIs in advance creates the accountability structure the program needs in order to improve over time.

The goal is not adding AI agents into existing workflows and leaving everything else unchanged. The goal is redesigning how work moves through the security operation, and doing so in a way that allows the different functions of the SOC to actually connect.

That requires thinking differently about roles, responsibilities, workflow design, and what security leadership is accountable for managing. It is organizational change supported by technology, not technology that substitutes for organizational thinking.

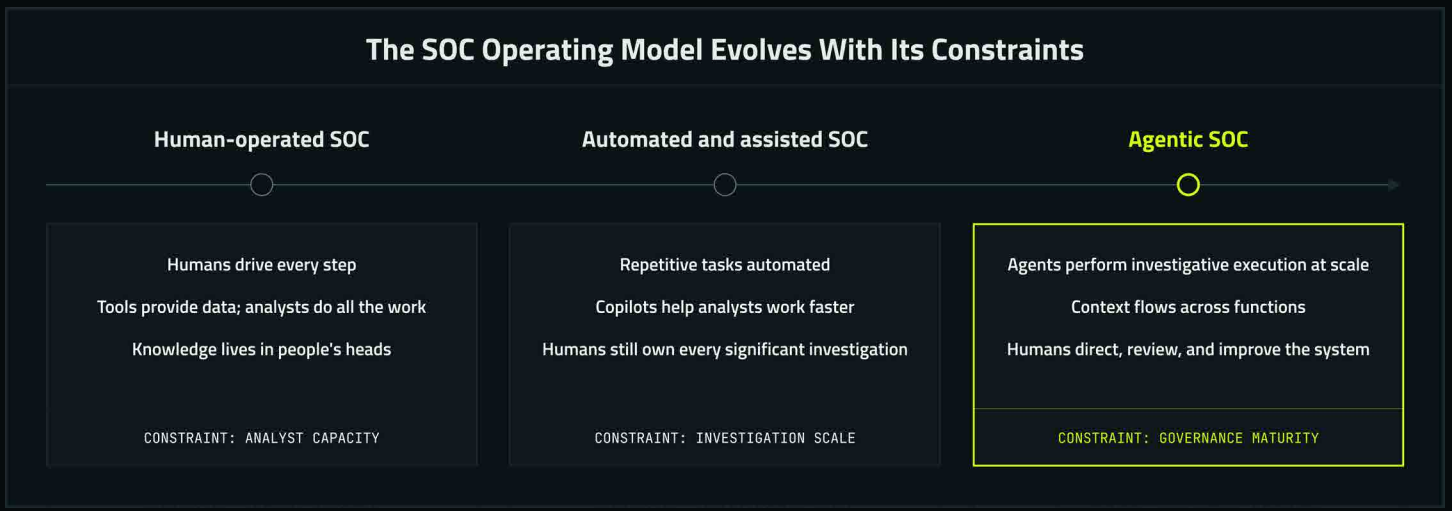

The SOC operating model evolves in response to changing constraints. At each stage, the bottleneck shifts — and so does the required investment.

Conclusion

A Security Operation Built for What's Ahead

Security operations have always evolved in response to changing constraints. Organizations invested in detection when visibility was the limiting factor. They invested in automation when repetitive work was consuming too much of an analyst's time. They adopted copilots when information access and individual productivity were the problems to solve.

Each investment addressed a real constraint. Each moved the practice forward. And each eventually exposed the next constraint, which is where the industry finds itself now.

Today's challenge is operating effectively when the volume and complexity of security work has outgrown what human execution alone can sustain, and doing so across an environment fragmented into disconnected tools, disconnected teams, and disconnected workflows that were never designed to function as a system.

Agentic security operations create the possibility of a different kind of SOC. One where investigative work proceeds at a scale that doesn't depend on analyst availability at every step. One where context flows across functions rather than stopping at the boundaries of each tool. One where threat intelligence informs investigations in real time, vulnerability findings reach incident responders when they're relevant, and threat hunting activity feeds back into detection without a formal review cycle to bridge the gap.

Realizing that possibility requires more than deploying new technology. It requires a new operating model and the governance capability to make that model trustworthy.

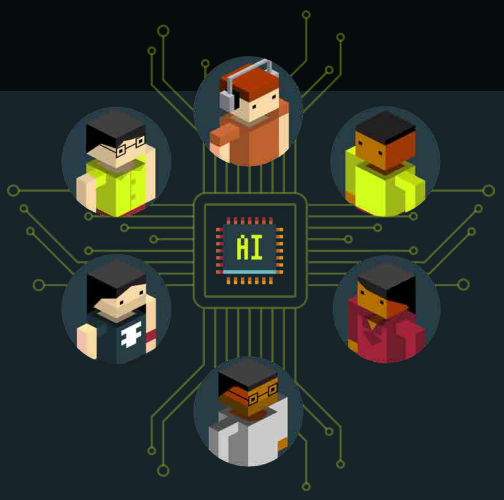

The Future SOC is Not One Without Humans.

It is one where human expertise is elevated, where security professionals direct, govern, and continuously improve a system capable of doing more than any team could accomplish alone. Where the walls between functions are replaced by connected workflows. Where institutional knowledge compounds over time.

That requires a platform built for coordination, not just automation. One where security professionals define the objective and the system determines how to achieve it, carrying context across functions, connecting work that operates in isolation, and adapting as new evidence emerges. One that gives security teams the governance controls and visibility to expand what AI agents do with confidence. That is what Bricklayer is built to deliver. **Learn more at bricklayer.ai.**

BOOK A DEMO TODAY
bricklayer.ai/book-a-demo