WHITEPAPER | EXECUTIVE INSIGHTS

Why the Traditional SOC Model isBreaking Down

A practical look at why legacy security operations struggle to keep up with modern threat environments, and what forward-thinking teams are doing instead to regain clarity and speed.

The failure mode is not a lack of tools. It is asking humans to reconcile all those tools at incident speed, while adversaries learn to move at machine speed.

The breakage is architectural

The traditional SOC is not failing because analysts have lost their edge. It is struggling because the operating model asks humans to behave like the integration layer for a system that now moves faster than human integration can. The work did not become less important. It became more fragmented.

Over the last decade, the defensive stack grew by acquisition, urgency and necessity. A SIEM to collect and correlate. SOAR to trigger workflows. CTI to enrich. EDR, NDR, IAM, cloud security, vulnerability management, ticketing and case management to handle the surrounding evidence and response. Each platform was rational on its own. Together, they often become a maze. The analyst's day turns into a tour of tabs: validate the alert, pivot to endpoint data, check identity, hunt through cloud logs, paste observables into threat intelligence, update the ticket, then repeat because one connector is out of date. Somewhere in there, a decision must happen, ideally before the attacker has lunch.

Recent research validates the pattern. Cisco's 2025 Global State of Security research reported that 78 percent of teams say their security tools are dispersed and disconnected, while 69 percent say that disconnection creates moderate to significant challenges. The same research points to the daily drag underneath those numbers: 57 percent lose investigation time to data management gaps, 59 percent have too many alerts and 55 percent deal with too many false positives [1]. Splunk's State of Security 2025 also found that 46 percent of respondents spend more time maintaining tools than defending the organisation [2]. Those numbers do not describe a lazy SOC. They describe a brittle architecture.

Why the old stack feels busy, but not fast

SIEM, SOAR and CTI still matter. They matter a lot. The problem is that their value collapses when they become destinations rather than services in a joined-up operating model. A SIEM is not a source of truth if identity context lives elsewhere. A SOAR playbook is not real automation if every exception requires a human to repair the connector, interpret an edge case or wait for approval in another queue. CTI is not intelligence if it arrives as another interesting feed that never changes a decision.

Legacy SOCs often confuse automation with choreography. The lights flash, the steps happen and the dashboard looks reassuring, but a human still has to keep the music playing. This is where tool sprawl becomes operational debt. Every console has a data model, query language, licensing boundary, permission structure and training curve. People end up learning platform behaviour instead of adversary behaviour. Changes in one tool break playbooks in another. Architects compensate with more pipelines, more dashboards and more integration work. Complexity quietly becomes the job.

This is expensive in the least glamorous way possible. It shows up as queue time, rekeyed facts, half-complete evidence packs, duplicated triage and missed context. The analyst is not just detecting threats. The analyst is also doing logistics for the investigation. That is a poor use of scarce human judgement.

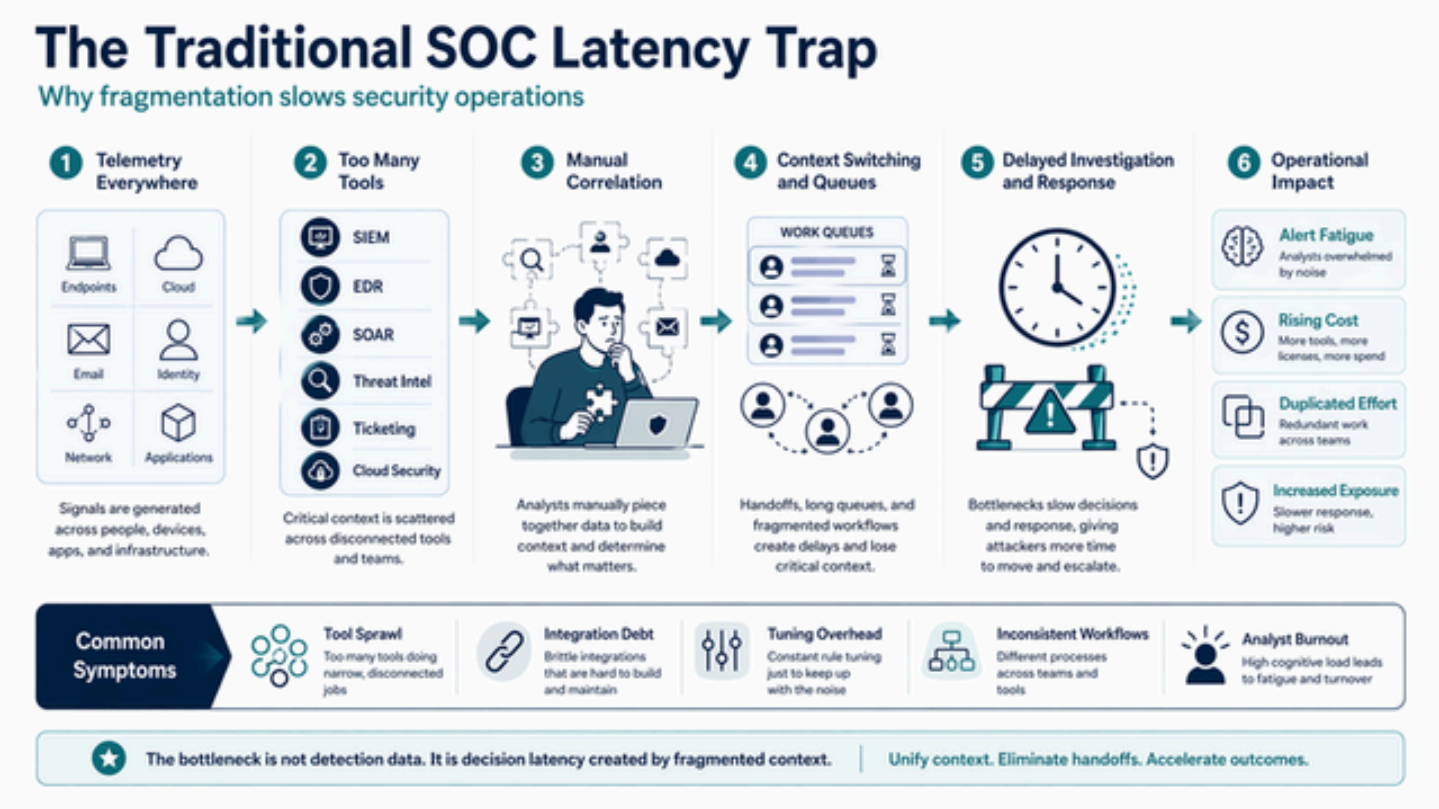

Figure 1. The operating shift from manual tool stitching to an agentic cybersecurity operating loop.

The economics of speed have changed

Security operations is now a time business. Speed to know, speed to decide and speed to contain are not operational luxuries. They are financial signals. IBM's 2025 Cost of a Data Breach report puts the global average breach cost at USD 4.44 million and attributes the year-on-year decrease partly to faster identification and containment. It also reports USD 1.9 million in average cost savings for organisations using AI extensively in security compared with those that do not [3].

That does not mean sprinkling a chatbot on top of a SOC and hoping for wizardry. It means the economic case for faster, better coordinated operations is now visible to the board. A SOC that can gather context, form a hypothesis, test it, present evidence and recommend containment faster is not merely more elegant. It is more resilient. The trick is to make speed governed and explainable, not frantic.

The Glasswing moment

Project Glasswing and Claude Mythos Preview make the same point from the other side of the battlefield. Anthropic introduced Project Glasswing in April 2026 as an initiative to give selected defenders access to Claude Mythos Preview for critical software security work. Anthropic describes Mythos Preview as a gated research preview of its most capable model for coding and agentic tasks, and says it has already identified thousands of zero-day vulnerabilities across critical infrastructure [4][5]. Its Frontier Red Team write-up describes a substantial leap in the model's ability to identify and exploit zero-day vulnerabilities in major operating systems and browsers when directed to do so [5].

The important lesson is not that everyone should panic. Panic is a terrible operating model and a worse procurement strategy. The lesson is that offensive and defensive cybersecurity are becoming AI-mediated. Australia's Cyber Security Centre put it plainly: frontier models are not creating new vulnerabilities, but they are reducing the cost, effort and expertise required to discover and exploit vulnerabilities that already exist [6]. If attackers get faster at finding the weak seam, defenders need faster operating loops to verify, contextualise, contain and learn. The manual clipboard SOC is not going to enjoy this race.

What forward-thinking teams are doing instead

The answer is not to abandon people, nor to throw away SIEM, SOAR and CTI. The answer is to turn those systems into capabilities that an operating layer can use. Telemetry, identity, cloud state, threat intelligence, vulnerability context and case history should flow into a governed agentic layer that can reason across them. Humans should set intent, risk appetite and command. Agents should gather context, compare patterns with recent cases, draft hypotheses, propose containment, document evidence and learn from analyst feedback.

The design shift is subtle but profound. First, reduce pivots. A case should carry its context with it instead of asking the analyst to go shopping for facts. Second, make reasoning visible. Every recommendation needs provenance, confidence and a reversal path. Third, separate autonomy from authority. Agents can do investigative work at machine speed, but the human operator decides where authority lands. Fourth, treat models as an evolving capability, not a fixed product. On-premise LLMs, private cloud LLMs and frontier cloud models each have a role depending on sensitivity, latency, capability and governance.

That is also where trust becomes practical. Analysts do not need a mysterious black box. They need a colleague that is fast, tireless, auditable and comfortable doing the boring parts without pretending to be the boss. Think less rogue intern with root access, more disciplined teammate who always brings receipts.

Figure 2. From alert queues to a governed operating loop: agents maintain momentum while humans command.

Where CiBRAI fits

CiBRAI's Cybersecurity Operating System is being designed around this shift. Rather than adding one more console to admire at 2 a.m., it blends human and agentic AI operators into a single operational fabric. It can orchestrate the latest on-premise and cloud LLMs so sensitive investigations stay controlled while teams still benefit from the latest

defensive reasoning. It connects SIEM, SOAR, CTI and the surrounding toolsets without making the analyst remember where every fact lives.

The 10x efficiency ambition is not magic. It comes from removing hundreds of small frictions: fewer pivots, fewer dead-end queues, fewer manually copied observables, faster evidence packs, reusable investigation patterns and a better way to learn from closed cases. The goal is not to replace analyst judgement. It is to stop wasting it on tool logistics.

The future SOC is a command system

Traditional SOCs break down when complexity becomes the work. The modern SOC has to make complexity the substrate. That means tools become services, agents do the stitching and humans command the mission. The analyst should not be the API gateway for a disconnected empire of consoles. The analyst should be the commander of a responsive operating layer.

That is where security operations is heading: not a bigger pile of connected tools, but a cybersecurity OS that can sense, reason, act and learn at the speed the threat environment now demands. The future is not less human. It is more deliberately human, because the machines finally take on the part of the job that machines should have been doing all along.

References

[1] Cisco Newsroom. (2025). Global State of Security Report reveals critical need for connected security operations. Source

[2] Splunk. (2025). State of Security 2025: The stronger, smarter SOC of the future. Source

[3] IBM. (2025). Cost of a Data Breach Report 2025. Source

[4] Anthropic. (2026). Project Glasswing. Source

[5] Anthropic Frontier Red Team. (2026). Claude Mythos Preview. Source

[6] Australian Signals Directorate. (2026). Frontier models and their impact on cyber security. Source

Prepared for CiBRAI article series: SIEM, SOAR, CTI and the future of agentic cybersecurity operations.

CiBRAI
CYBER INTELLIGENCE - BEHAVIOURAL RESPONSE