Everest Group®

December 2025

Fortify Cyber Defense with AI-led Security Operations

Contents

Introduction

Over the past two years, breakthroughs in generative and agentic AI have transformed both cyberattacks and defenses. Once productivity tools, these technologies now enable adversaries to launch realistic impersonations, generate polymorphic malware, and overwhelm traditional Security Operations Centers (SOCs) with machine-speed attacks – rendering rule-based, human-dependent defenses inadequate.

In this Viewpoint, we examine the implications of this new reality and outline how organizations can recalibrate their approach to security operations. We explore:

- The rise of AI-powered adversaries and how attackers are leveraging generative and agentic AI to mutate payloads, scale social engineering, and evade detection, exposing the limits of rule-led SOCs
- AI-led security models that transition routine detection, correlation, and containment to machine speed while keeping human oversight, intent, and governance at the center
- The future of the autonomous SOC, where continuous learning, agentic execution, and safe rollbacks redefine resilience and elevate security from a checklist function to a business confidence driver
- Enterprise implications across resilience, efficiency, workforce transformation, governance guardrails, and executive-ready performance metrics that translate machine-speed defense into measurable business outcomes
- The services layer needed to operationalize the AI-led SOC, highlighting how providers can help enterprises redesign processes, co-manage operations, accelerate responses, and continuously tune AI models using live threat intelligence

By understanding how AI reshapes the risk landscape, enterprises can fight AI with AI – strengthening resilience, boosting efficiency, and aligning machine-speed defense with business objectives. Generative and agentic AI represent the next frontier of cyber defense. Their progression – from rule-based automation to systems capable of reasoning and acting independently – shifts cybersecurity from static detection to adaptive, self-learning defense operating at machine speed. As this Viewpoint focuses specifically on generative and agentic AI, it is important to outline how AI capabilities evolve across intelligence, context grounding, and decision-making.

Exhibit 1: How capabilities progress across intelligence, context grounding, and decision-making

Source: Everest Group (2025)

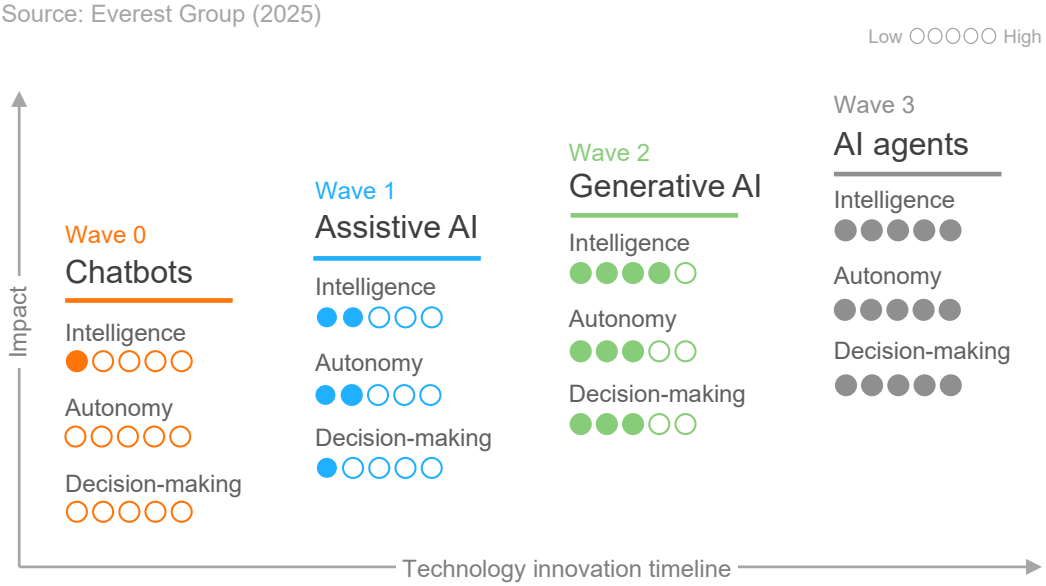

The rise of AI-powered adversaries

Realistic, real-time, and relentlessly iterative attacks

Machine speed deception at enterprise scale

Attackers now operate at machine speed as generative AI creates highly convincing phishing content, deepfake voices, and realistic personas. Agentic AI systems further amplify this threat by coordinating large-scale campaigns that register domains, launch fake sites, and dynamically adapt evasion tactics, turning what was once a manual process into a fully orchestrated machine-driven operation that continuously mutates to bypass static detection rules and human reviews. Cross-channel campaigns align deepfake voices and faces with CAPTCHA-fronted redirect chains and malvertising or SEO-poisoned downloads, so the experience feels routine, blends into trusted workflows, and compresses approval cycles. The result is a high-volume, highly personalized stream of low-signal alerts that overwhelms manual triage, allowing high-impact actions such as payment changes, credential harvest, and session hijacking to slip through normal business processes.

Exhibit 2: Example scenario: the high-volume and low-signal threat

Source: Everest Group (2025)

1 Recon

Attacker action (AI-enabled)
Uses AI to scan public information and build a map of the company and its vendors, learn how leaders write and sign emails, and prepare a list of likely targets and relationships.

AI capability in play
Agentic web agents with tool use, LLM entity extraction/ summarization, and style and persona modeling

How it evades triage
Largely off-network or low-rate benign queries; looks like normal browsing/search with no clear indicators to alert on

Business impact
Improves attacker accuracy and reduces warning signs, allowing them to bypass early human and system checks

2 Attack setup

Attacker action (AI-enabled)
Generates highly realistic provider and executive look-alike emails and domains.

AI capability in play
Style mimicry, persona cloning, rapid domain generation

How it evades triage
Each alert is a low-risk singleton and routed to low priority

Business impact
Exposure increases across mailboxes

3 Noise flooding

Attacker action (AI-enabled)
Sprays many plausible but minor anomalies over days

AI capability in play
Automated campaign management and low-level variation

How it evades triage
Volume dilutes attention and no cross-signal correlation

Business impact
Signals are buried in alert backlogs

4 High-impact action

Attacker action (AI-enabled)
Lures finance managers to fake invoice portals, captures credentials, accesses CFO mailboxes, injects payment instructions into real threads.

AI capability in play
Thread hijacks, tone and signature replications, session takeovers

How it evades triage
Individually, these look medium or routine and remain unlinked to initial phishing

Business impact
Business email compromise pathway is established

5 Outcome

Attacker action (AI-enabled)
Issues convincing wire requests from CFO accounts and passes casual checks

AI capability in play
Context-aware drafting that matches prior CFO communication

How it evades triage
Message appears normal in a legitimate thread and bypasses secondary checks

Business impact
Large transfers executed before detection

This example reflects just one instance of how AI-powered deception unfolds across multiple stages, exploiting both human behavior and system silos. Similar incidents have been observed across industries where generative and agentic AI amplify attackers' reach and precision. Exhibit 3 highlights real-world cases that demonstrate how quickly the impact escalates when such threats remain unchecked.

Exhibit 3: Case studies depicting the impact AI-powered threats can have if left uncontrolled

Source: Everest Group (2025)

Threat	Business Impact	How AI-led security could have prevented it
Multi-participant deepfake video meeting impersonates senior finance leader and colleagues, pressing for urgent transfers.	~US$25 million wired across multiple accounts; approval workflow subverted via convincingly forged presence and voice.	Behavioral risk scoring on approvals AI workflow anomaly detection Agentic hold and step-up verification for high-value payments
Mass-generated fake CAPTCHA verification pages on AI-native builders (for example, low-friction hosting) blind scanners and boost phish delivery.	Higher delivery and conversion rates for credential thefts; large-scale account takeovers and fraud from business-as-usual-looking flows.	AI crawl that renders CAPTCHA flows AI pattern clustering to detect and block phishing kit variants Real-time brand page anomaly detection
Search/Ad lures for AI tools (look-alike sites, social ads) deliver information stealers/ backdoors via SEO poisoning and malvertising.	Endpoint compromise at scale → credential exfiltration (SaaS, cloud, VPN) → business email compromise and lateral movement.	AI brand and impersonation detection across web and ads Adaptive sandboxing that simulates user behaviors Auto-correlation to isolate compromised sessions
Deepfake KYC/liveness bypass opens mule/vendor accounts that pass onboarding, later used for fraud and laundering.	Synthetic accounts clear controls, transact for months, and default or facilitate large fraud; downstream chargebacks and regulatory exposure.	AI liveness and synthetic media detection Graph analytics on identity, device, and IP links Continuous model retraining on new deepfake patterns

Attackers now use generative AI for realism and agentic AI for speed, which shrinks the timeline from recon to impact and overwhelms rule-based SOCs. Exhibit 4 shows where human-led triage breaks and why an AI-first approach is required. It summarizes key limitations of rule-led SOCs and how they create structural gaps in defending against AI-powered attacks.

Exhibit 4: Rule-based SOCs' limitations

Source: Everest Group (2025)

	Missing capability
Siloed signals Root cause: Fragmented tools and disconnected telemetry prevent end-to-end visibility	AI-driven cross-domain correlation unifies identities, endpoints, cloud, and network data into one contextual view
Static detection Root cause: Rule-based logic detects only known patterns	Adaptive learning models identify new behaviors and continuously refine detection thresholds
Analyst overload Root cause: Manual triage and repetitive investigation tasks consume analyst bandwidth	Autonomous triage and AI-driven prioritization surface only high-risk alerts for human review
Trust assumptions Root cause: Human validation biases and static access rules create blind spots	Continuous trust assurance validates users and sessions in real time using behavioral signals
Slow responses Root cause: Human gating delays containment actions and approvals	Policy-bound autonomous containment executes pre-approved actions instantly with full audits and rollbacks
Context fragmentation Root cause: Alerts lack linkage to business or identity contexts, limiting risk prioritization	Context-aware reasoning connects detections to assets, users, and business impact for precise containment
Reactive postures Root cause: Lack of predictive insights leads to delayed responses to emerging threats	Predictive and adversarial-aware analytics forecast attack paths and trigger proactive hardening

AI-powered threats have shifted the balance in favor of attackers by exploiting scale, speed, and trust signals that traditional SOC models cannot reliably defend against, leaving enterprises exposed even when controls appear to be in place. The clear implication is that security operations must evolve from rule-led detection and manual containment toward AI-driven defenses that can learn continuously, correlate signals across domains, and act at machine speed under human intent.

AI-led security: fighting AI with AI

From analyst-centric SOCs to AI-led operations with human oversight

AI-enabled threats demand AI-enabled defense, but outcomes hinge on where that defense is orchestrated. SOC is the execution layer that converts tools into governed decisions at scale. An AI-driven SOC redesigns workflows, roles, metrics, and controls such that models prioritize risks, agents automate investigations and responses under human oversight, unified data pipelines enrich and explain decisions, and governance manages privacy and model risks, turning potential into repeatable, auditable outcomes.

The shift to AI-led security operations under human oversight

Enterprises are moving from analyst-centric workflows to machine-led operations, where connected AI agents triage alerts, gather context, and execute bounded actions, and analysts set policies, supervise exceptions, and handle novel cases, with transparency through audit trails of agent reasoning and decisions. This forms the foundation of an AI-driven SOC that works alongside people rather than replacing them. Leading SOCs have implemented workload management, such as:

- **Autonomous triage:** AI agents collect evidence across telemetry, summarize context, and issue or propose verdicts for routine alerts under policies
- **Investigation copilots: A**ssist analysts by fetching contexts, building timelines, and drafting case notes with audit trails that explain why a conclusion was reached
- **Policy-bound containment:** Automatically runs common actions when confidence is high, including isolating endpoints, revoking sessions, resetting credentials, and disabling risky tokens
- **Cross-domain correlation:** Unified data platforms let AI link signals from email, web, identities, endpoints, and cloud, reducing alert volumes so fewer low-value alerts reach humans
- **Human oversight as SOC pilots:** Analysts review exceptions, set policies, and approve actions when uncertainties cross defined thresholds, with complete logs for audits
- **AI governance in the SOC:** Defined policies for data handling and access, model and prompt lifecycle management, evaluation and red teaming, drift monitoring, approval thresholds for actions, and end-to-end audit and change control
- **Open integrations:** Agent workflows execute across multiple security tools through standard interfaces while honoring enterprise guardrails and change controls

While these capabilities mark strong progress toward AI-led security operations, many SOCs are still early in their adoption curve. Most implementations remain partial, limited to triage, enrichment, or containment within predefined scopes. The real challenge lies in ensuring that these AI-driven elements operate seamlessly across the entire threat life cycle. Each stage of an attack, from reconnaissance to recovery, still exposes pressure points where manual controls, latencies, and data silos slow down responses. Exhibit 5 illustrates the threat lifecycle stages and highlights the pressure points where traditional, manual approaches struggle against machine-speed threats.

Exhibit 5: Threat lifecycle stages and the pressure points that overwhelm manual controls

Source: Everest Group (2025)

Threat lifecycle	Traditional approach			How it is being challenged
	Stage 1	Stage 2	Stage 3	
Signal ingestion and detection	Source onboarding and parsers	Rule and signature detections	Threshold-based alerts	AI-generated noise floods ingestion, polymorphic content evades signatures, short-lived infrastructure weakens reputation checks
Triage and prioritization	Threat enrichment	Analyst reviews	Severity-based queueing	Volume and speed exceed human throughput, cross channel coherence is missed, high-risk actions occur before the alert is touched
Context and correlation	Manual pivots to assets and library	Rule-based correlations	Indicator matching	Synthetic identities appear valid, weak signals across tools fail to combine, CAPTCHA and redirect chains hide payloads
Investigation and scoping	Artifact collection	Sandbox and log reviews	Scope confirmation	AI obfuscation defeats simple heuristics, encrypted collaboration limits visibility, attacker iteration outpaces analysis
Containment and response	Analyst recommendations	Approval and change controls	Endpoint isolation credential resets	Decision windows compress to minutes, cloud and SaaS sessions persist beyond device actions, and kit factories replenish domains rapidly
Recovery and assurance	Rebuild and restore	Password resets and communications	Manual verification and rollback as needed	Stolen tokens and refresh sessions survive rebuild, lateral movement into SaaS leaves residual access, and limited automated verification creates blind spots

What else do SOCs need now?

- **Identity-first detection with graphed IAM context:** Continuous discovery identifies both human and machine accounts, providing user visibility and automated risk scoring. These insights tie into authoritative identity stores and governance systems to link accounts with their owners, entitlements, and behavioral baselines. As a result, event detection and investigation become identity-aware rather than siloed. Posture is continuously managed, and response actions – such as credential revocation, step-up authentication, and targeted policy hardening – are executed automatically. Decoy identities can also be deployed to generate high-fidelity alerts
- **Risk-based exposure management with closed-loop remediation:** SOCs must unify exposure signals across networks, endpoints, and cloud; use AI to prioritize exploitable risks; and automate fixes and protections across first- and third-party controls so that vulnerability noise drops and risk actually closes
- **Predictive and adversarial-aware analytics:** SOCs must deploy predictive models and User and Entity Behavior Analytics (UEBA) to forecast likely attack paths, including zero-day exposure, while detecting anomalies early. Pair this with controls that recognize manipulation of data and model inputs, ensuring integrity and trust
- **LLM-powered protection of the communication layer:** SOCs should be able to detect intent-driven phishing and deep social engineering in mail and collaboration streams, auto-remove malicious messages, and take immediate identity and endpoint actions to contain compromise
- **Comprehensive use-case coverage on one platform:** SOCs must span Extended Detection and Response (XDR), data security posture, cloud detection and response, SaaS and identity signals, infrastructure-as-code checks, and exposure management within a single data and policy plane such that proactive and reactive controls reinforce each other and tool sprawl is reduced
- **Explainable AI that closes the talent gap with persona-aligned workspaces:** SOCs must implement explainable AI that provides decision traces, contributing signals, confidence levels, and evidence packs in plain language so that junior analysts ramp faster, seniors review more efficiently, and leaders can audit why the system acted. Role-based views then allow SOC analysts, responders, hunters, and identity administrators to work in one place with the right permissions
- **Unified telemetry and context fabric:** SOCs should bring together front-end user activity and collaboration with backend workload, data, and infrastructure events, plus rich identity context, so profiling, risk scoring, and detections use the full picture rather than tool silos
- **Global analytics and extensive data coverage with strong governance:** SOCs must leverage insights from a wide range of environments to identify emerging attack patterns early, train AI models on diverse and high-quality data to improve detection accuracy and reduce false positives, and maintain strict data controls for residency, access, and compliance so that the intelligence does not compromise trust

Beyond these operational capabilities, AI is redefining how detections are built and maintained. Generative AI refines rule signatures, reviews code for logic gaps, and generates new detection logic from analyst notes or threat reports. Agentic AI then tests these detections against real or simulated attacks, tunes thresholds, and proposes safe updates for approval. Together, they make detection engineering faster, more accurate, and continuously adaptive – evolving defenses at the same pace as emerging threats.

The future of the autonomous SOC

Potential future state of more autonomous SOCs

Turning adaptation into resilience

Continuous learning is now vital for modern SOCs as adversary tactics and behaviors evolve at machine speed, creating constant drift between trained models and live environments. Static updates no longer suffice; SOCs need embedded mechanisms to detect drift, generate and test new playbooks from fresh intelligence, and deploy them rapidly with rollback safeguards. This transforms every alert and analyst action into input for smarter defenses, compressing the cycle from threat emergence to containment.

Modern SOCs embed feedback loops across data parsing, triage, detection, and response, ensuring systems continuously adapt. Parsers learn from analyst corrections, playbooks evolve from past incidents, and detectors are validated before deployment. Resilience is now measured not by static control strength but by how quickly defenses adapt and recover – turning continuous learning into a core strength for cyber resilience.

Continuous learning operates through three interconnected loops addressing data drift, emerging threats, and changing business priorities – constantly retraining models, updating playbooks, and aligning responses with business intent for sustained defense agility.

> “A SOC that does not learn every day falls behind every day.”
>
> – CISO, Fortune 500 enterprise

Exhibit 6: Three loops of continuous learning in SOCs

Source: Everest Group (2025)

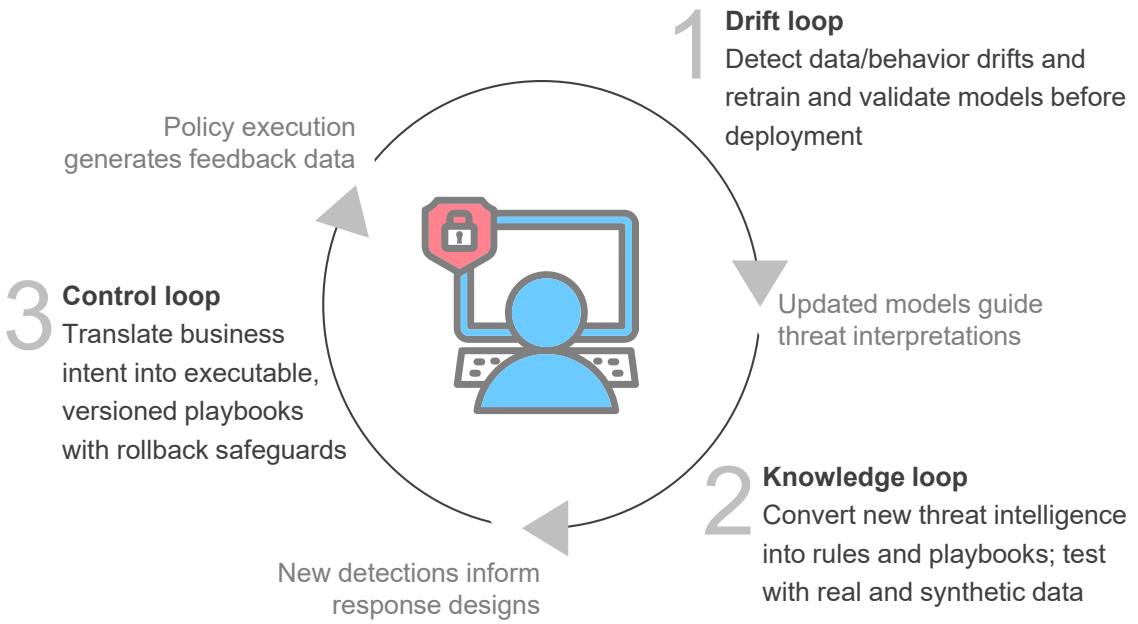

In the next generation of SOCs, human experts will focus on setting security intent and defining policies, while agentic AI systems carry out those directives with speed and precision. Unlike static automation, these systems are context-aware and goal-driven, capable of triaging threats, investigating anomalies, and initiating responses on their own. Once given a clear mandate, for example, containing a breach within defined risk limits, agentic AI can plan and execute a series of actions to achieve that outcome. Importantly, it does so within human-defined guardrails where every step is transparent, results are verified, and any unintended changes can be rolled back safely. In practice, this shifts the heavy lifting of real-time defense to AI, while ensuring every action remains aligned with the intent and risk tolerance set by human decision-makers.

This transition to agentic AI follows a clear progression that can be best understood as a ladder of autonomy:

- **Assisted:** AI helps analysts with summaries, recommendations, and data pulls. Humans make every security change
- **Supervised:** Agentic systems take bounded actions in real time, such as isolating a device or blocking access, while analysts monitor and review important actions
- **Delegated:** Within strict policy limits, AI operates on its own, and humans step in for exceptions, audits, and edge cases

Exhibit 7: The ladder of autonomy

Source: Everest Group (2025)

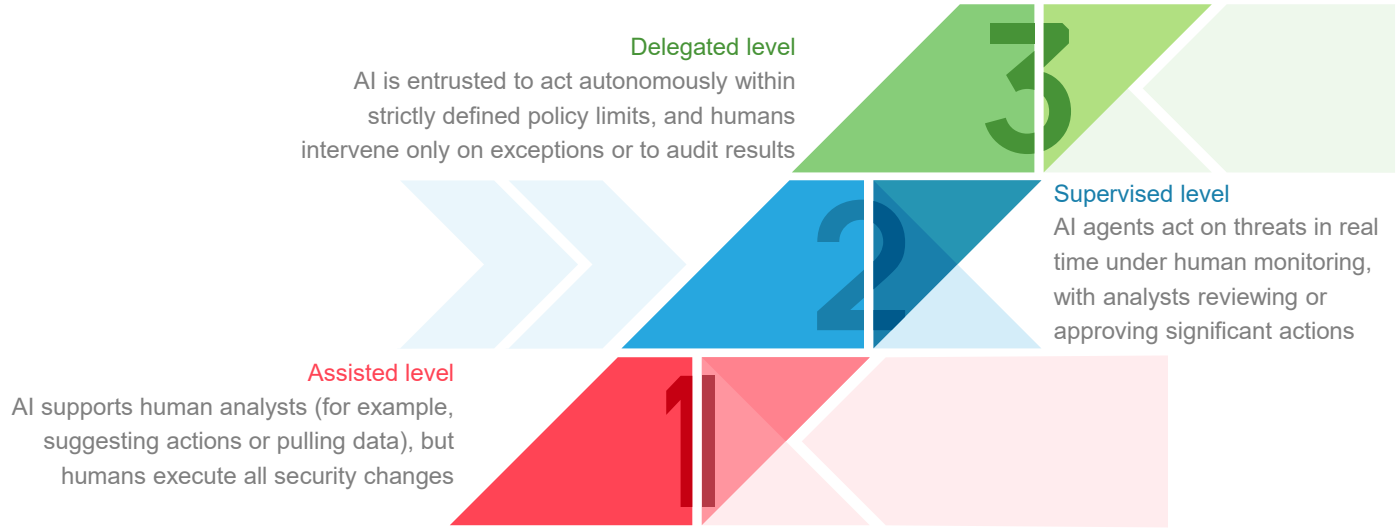

Outcome-driven business assurance with AI-enabled SOCs

Forward-looking organizations are shifting from process-based security to outcome-driven assurance. SOCs are now evaluated by how effectively they maintain business continuity and resilience, not by alert volume. Executives expect cybersecurity to quantify risk reduction and recovery performance – demonstrating protection of critical systems and rapid restoration. This shift emphasizes SLOs focused on minimizing impact, downtime, and disruption, with assurance reports highlighting faster containment, successful remediation, and business impact avoided.

Real-world examples are beginning to validate this shift from compliance-driven security to business outcome assurance:

- **Autonomous threat containment:** A global electronics manufacturer faced a novel ransomware attack that evaded traditional playbooks. An AI-driven SOC platform detected the abnormal encryption behavior and autonomously triggered containment – isolating infected machines and blocking the attack's spread in under a minute. This swift reaction prevented any significant downtime
- **Ransomware rollback in action:** When ransomware began encrypting files on a small set of endpoints, agentic AI immediately isolated the devices, stopped the process, and preserved evidence. It identified the last known good state and opened a ticket for the backup team to restore from approved backups with integrity checks and dual control. Humans approved each step; recovery met the firm's RTO and the blast radius was contained without paying ransom

The future SOC, therefore, is not just a technical command center but is evolving into a business assurance function, giving top executives peace of mind that cybersecurity is effectively safeguarding the enterprise's key outcomes in real time.

Enterprise implications

The enterprise perspective

- **Turning adaptation into resilience:** Agentic and generative AI are transforming security operations, making resilience, efficiency, workforce evolution, and governance measurable through metrics like MTTD, MTTR, containment rate, and recovery assurance. Cybersecurity now acts as a quantifiable enabler of business continuity and confidence
- **Resilience: from reactive recovery to assured continuity:** AI-driven SOCs prioritize business continuity over incident response. Agentic AI autonomously isolates assets and triggers pre-approved playbooks, while generative AI reconstructs incidents and proposes remediation, reducing impact and recovery time while assuring uninterrupted operations
- **Efficiency: turning machine speed into measurable value:** Efficiency arises from precision, not volume. Agentic AI removes redundancy across domains, and generative AI consolidates fragmented data into clear narratives, cutting investigation time, reducing false positives, and converting machine speed into measurable business value
- **People: evolving from operators to intent setters:** Analysts shift from manual execution to strategic oversight as AI handles triage and containment. Generative AI mentors teams by explaining patterns and summarizing alerts, driving faster upskilling, lower burnout, and broader defense coverage with fewer resources
- **Governance: assurance through transparency and testing:** Greater autonomy requires transparent control. Enterprises now deploy model assurance frameworks to test AI behavior, enforce guardrails, and log every automated action. Generative AI documents and explains outcomes, enabling continuous assurance and policy-aligned accountability

> "The real breakthrough was not faster detection; it was confidence. We can now quantify our security posture in the same language the business uses for resilience and continuity."
>
> – CISO, global consumer enterprise

Exhibit 8: SOC's evolution into a system of confidence

Source: Everest Group (2025)

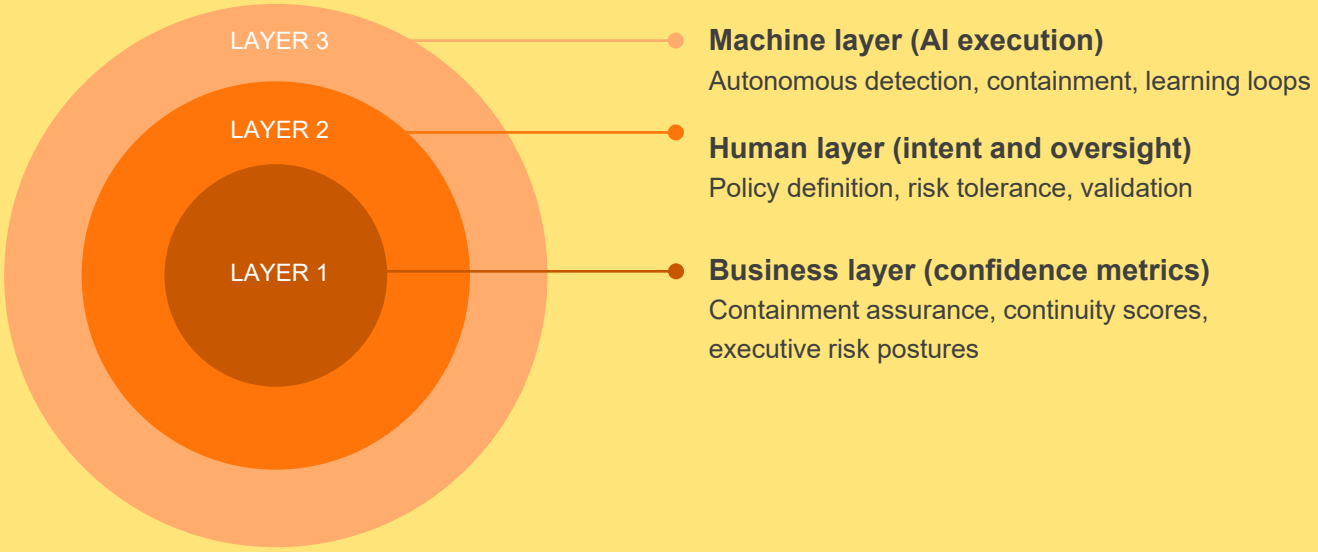

The services layer: operationalizing the AI-led SOC

The need for providers to operationalize the AI-led security narrative

The AI-led SOC is a governed services layer that sits over existing Security Information and Event Management (SIEM), Endpoint Detection and Response (EDR), Identity and Access Management (IAM), cloud, and ticketing stack with policy-bound generative and agent-driven automation. It does not replace tools; it binds them by normalizing telemetry, applying AI assistants to read and reason over enterprise context, and orchestrating responses within explicit guardrails and approvals. It should support deployment across major clouds and allow the use of enterprise-approved foundation models so organizations can align with residency, security, and procurement preferences. Success is evidenced by faster detection, rapid but reversible containment with human oversight, measurable reduction in analyst efforts, and auditable execution that aligns with privacy and access-control policies.

Must-haves in services layer

- **Interoperable AI fabric:** SOCs should merge outputs from all tools into a single, role-based view connected via APIs to enterprise systems. Identity-aware access and end-to-end orchestration ensure seamless, governed response
- **Deep context through retrieval-augmented grounding:** Dual-tier retrieval-augmented generation keeps AI decisions aligned with enterprise context. Analysts can query curated knowledge bases and temporary session data to reason over the latest playbooks and threats within governance limits
- **Quantified time compression across SOC workflows:** AI orchestration should collapse multi-tool processes, like threat triage, into unified, automated workflows completed in minutes, boosting containment speed and analyst efficiency
- **Persona-aware access and orchestration fabric:** A persona-based orchestration layer enforces identity and access governance through role-specific views and single sign-on, ensuring traceable, policy-bound actions
- **Extended SOAR and risk orchestration:** SOC platforms should trigger cross-tool playbooks from a single prompt, monitor MTTD and MTTR, run automated threat lookups and risk mapping, and self-update playbooks as patterns evolve
- **Dedicated tenancy and privacy-anchored governance:** Deploy SOCs as dedicated, privacy-controlled tenants to maintain data sovereignty. Session data stays internal, access is tightly governed, and shared learning occurs only through sanitized, reusable workflows

How providers can run it

To bring these capabilities to life, service providers should adopt a structured approach to rollout, governance, and continuous improvement, as detailed below:

- **Phased rollouts with guardrails:** Advance with discipline: diligence → design → deploy → tune → continuous improvement. Define prompt frameworks, engineer workflows, containerize into the customer cloud, validate outputs, enforce guardrails, and constrain responses to relevant context
- **Managed and co-managed SOC models:** Blend round-the-clock coverage with copiloted workflows. The provider maintains the AI fabric and playbooks; the enterprise sets intent, thresholds, and approvals. Prerequisites include named SMEs, tool APIs, hosting, and connectivity with a clear responsibility matrix from day one
- **Incident response and crisis management:** Use gen AI assistants to assemble timelines, executive briefs, and evidence automatically; drive → execute → validate → rollback steps within policy to keep containment safe and auditable
- **Continuous model tuning, threat intelligence, and simulated-attack loops:** Treat content as code and evaluate prompts and automations against fresh Tactics, Techniques, and Procedures (TTPs), sector intelligence, and continuous attack-simulation; refine retrieval scopes, thresholds, and actions weekly, to ensure that defensive learning stays ahead of offensive mutation

Exhibit 9: The services layer blueprint: from stand-up to measurement

Source: Everest Group (2025)

S	Signal integration	Normalize telemetry (SIEM, EDR, IAM, cloud, tickets) and expose clean APIs for AI use
E	Enrichment and knowledge	Gen AI assistants read policies / playbooks / threat intelligence (session-based uploads + approved library)
R	Response orchestration	Single-prompt, cross-tool playbooks that turn analyst intent into actions
V	Verification and oversight	Human approvals for material steps; time-bound rollbacks; full audit trails
I	Identity and risk context	Tie actions to user/device/app posture (step-up authorization, token revoke, least-privilege moves)
C	Containment and rollback	Safe, reversible actions across domains (endpoints, email/web, cloud/network, identities)
E	Evaluation and tuning	SLOs (MTTD/MTTR, autonomous-containment rate, analyst time saved), continuous tuning, attack-simulation loops

Conclusion

AI has permanently changed the tempo and texture of cyber risks. Offense now iterates at machine speed; defense must respond in kind without losing control. The path forward is clear in this Viewpoint: move routine detection, correlation, and bound containment to AI operating under policy, keep people focused on intent and oversight, and run the SOC as a continuously learning system measured by business-relevant outcomes.

Leaders have a straightforward mandate: translate policy into machine speed security operations without losing control by implementing a services layer that binds existing controls with AI. This layer enriches context, recommends actions, and orchestrates responses under clear approvals, audits, and rollbacks. Deployed within the enterprise cloud and capable of operating across major clouds and model choices, it scales core strengths by correlating signals, applying identity and business context, and containing issues quickly and safely. Progress should remain deliberate and governed, with an explicit autonomy envelope by risk tier, human approval for material steps, and a standing rhythm of continuous tuning across content, playbooks, and policies.

The payoff is tangible and near term: begin with a small set of high-yield use cases, instrument executive-grade service levels such as MTTD, MTTR to respond to containment, autonomous action rate, and analyst time saved, and operate in a co-managed construct while skills and confidence build. As assistive workflows demonstrate reliability, expand to preapproved and conditional actions where guardrails and rollback are well-defined. The result is faster containment, a smaller blast radius, and clearer assurance to the business; in short, an AI-first services layer converts intent into repeatable outcomes and delivers resilience at the speed modern threats demand.

Everest Group®
With you on the journey

Everest Group is a leading research firm helping business leaders make confident decisions. We guide clients through today's market challenges and strengthen their strategies by applying contextualized problem-solving to their unique situations. This drives maximized operational and financial performance and transformative experiences. Our deep expertise and tenacious research focused on technology, business processes, and engineering through the lenses of talent, sustainability, and sourcing delivers precise and action-oriented guidance. Find further details and in-depth content at www.everestgrp.com.

This study was funded, in part, by
Cognizant and Palo Alto Networks

For more information about
Everest Group, please contact:

+1-214-451-3000

info@everestgrp.com

For more information about
this topic please contact the author(s):

Kumar Avijit, Vice President
kumar.avijit@everestgrp.com

Arjun Chauhan, Practice Director
arjun.chauhan@everestgrp.com

Gautam Goel, Senior Analyst
gautam.goal@everestgrp.com

Notice and Disclaimers