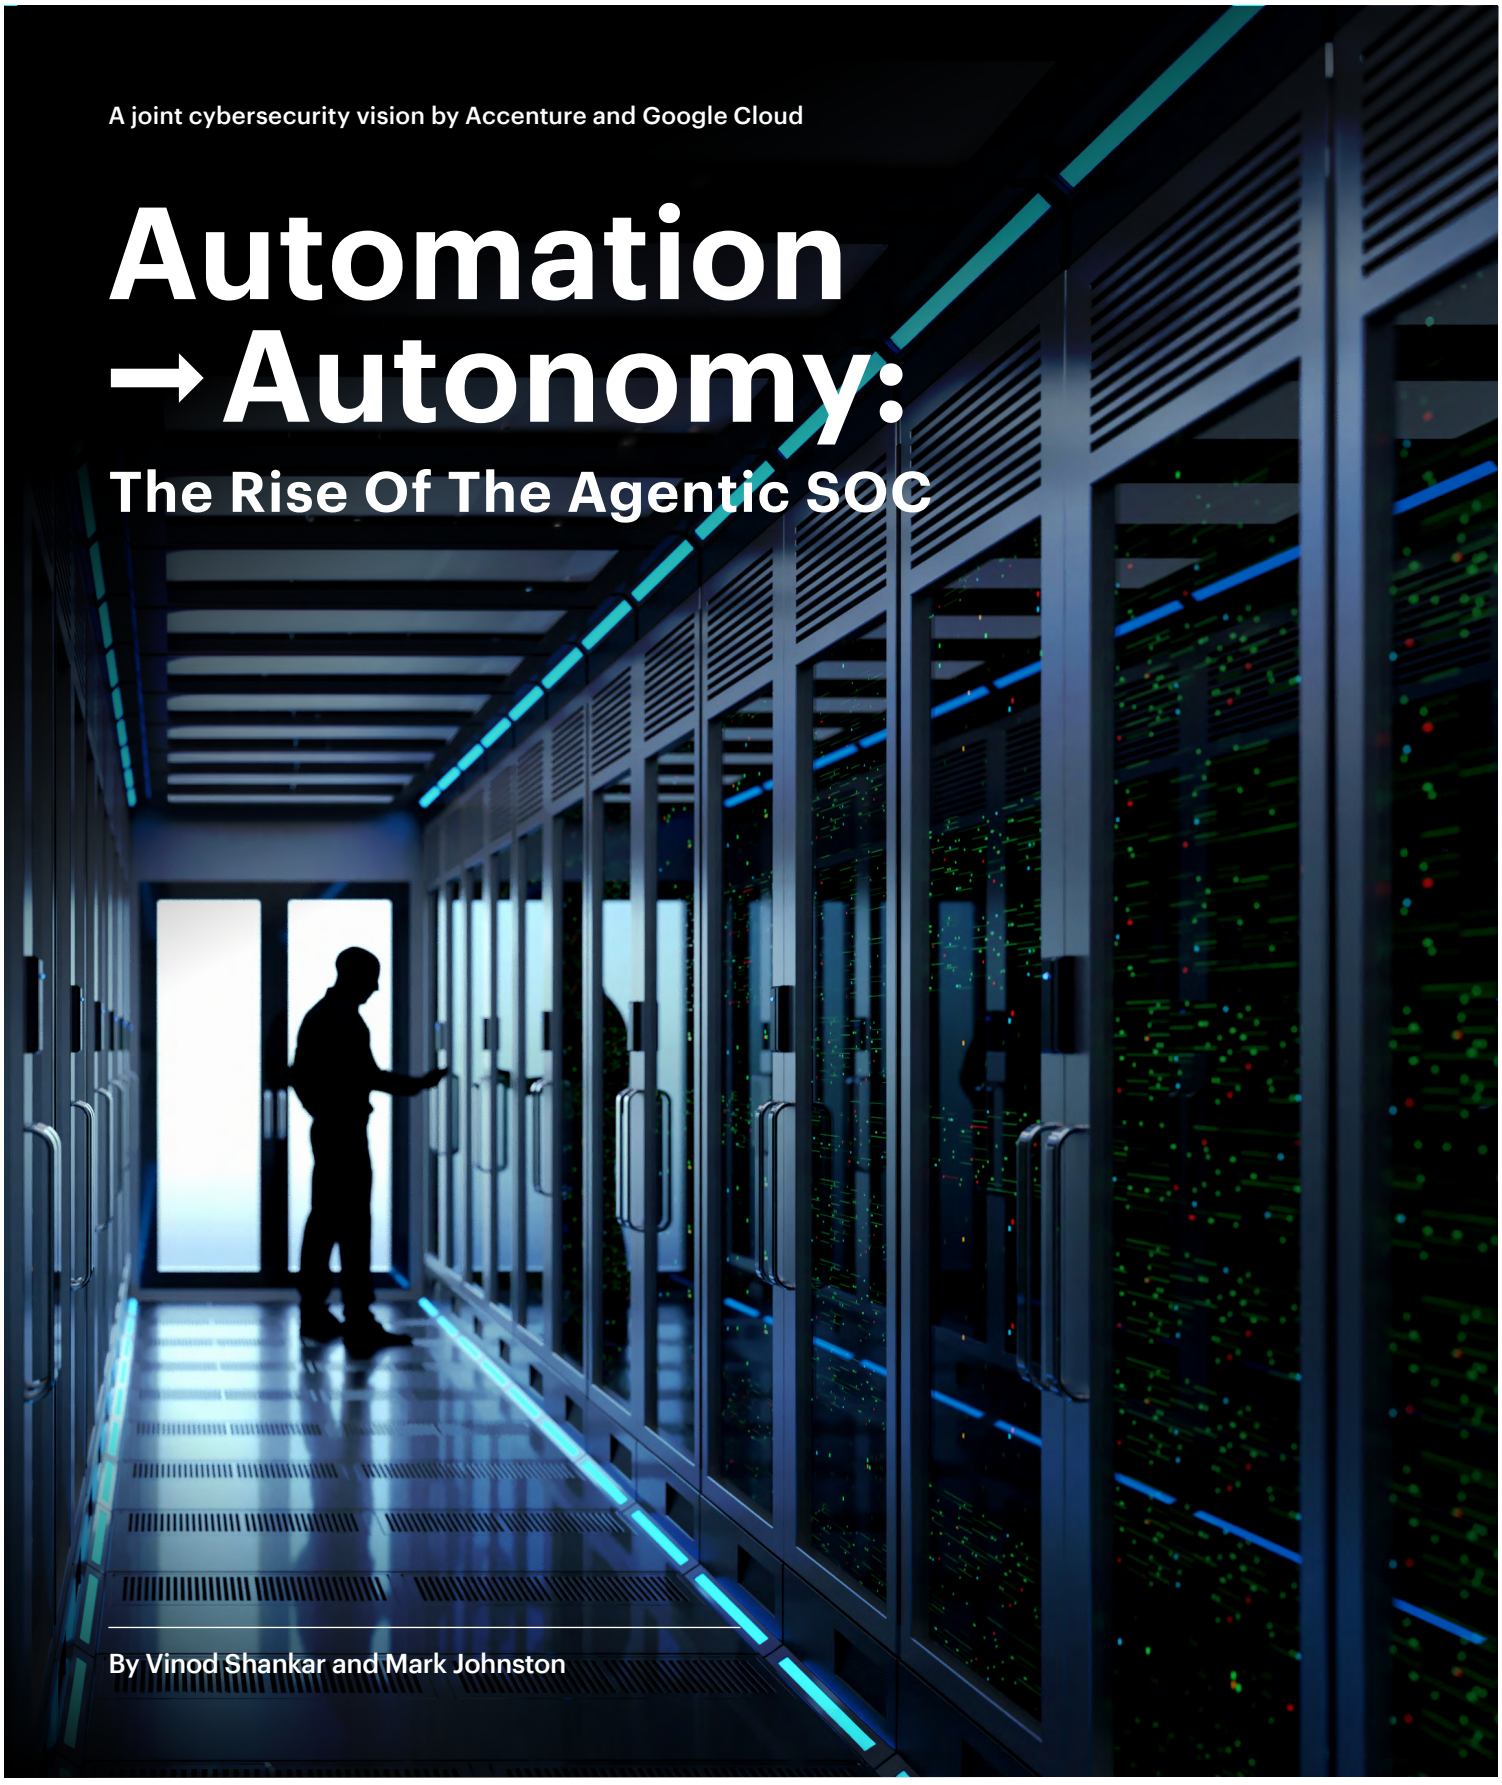
accenture | Google Cloud



A joint cybersecurity vision by Accenture and Google Cloud

Automation → Autonomy:

The Rise Of The Agentic SOC

By Vinod Shankar and Mark Johnston

1 **Executive summary** Page 3-4

2 **The SOC at an inflection point** Page 5-7

3 **The agentic leap** Page 8-11

4 **Reality vs. rhetoric** Page 12

5 **The convergence model: Building an autonomous security core** Page 13-15

6 **The road ahead** Page 16-18

7 **Conclusion** Page 19-21

Executive summary
The SOC at an inflection point
The agentic leap
Reality vs. rhetoric
The convergence model: Building an autonomous security core
The road ahead
Conclusion

1.

Executive summary

The context for reimagination

The modern Security Operation Center (SOC) faces unprecedented challenges across three critical dimensions:

Scale
Enterprises generate countless security events each day.

Latency
Human analysis cannot keep pace with machine-driven threats.

Entropy
Signal is buried beneath noise, eroding analyst confidence and focus.

Security teams find themselves overwhelmed by an unrelenting deluge of alerts, hampered by tool fragmentation, and exhausted by the repetitive toil of manual triage.

Over the past decade, the industry has invested heavily in automation — primarily through predefined workflows and Security Orchestration, Automation, and Response (SOAR) platforms — designed to accelerate ticket handling and reduce operational response times. While these solutions have delivered incremental improvements, they remain fundamentally reactive, rule-based systems that struggle to adapt to the dynamic threat landscape where adversaries operate with machine speed and human ingenuity.

The thesis: Moving beyond automation to autonomy

This paper argues that the core challenge facing the modern SOC is not a lack of technology, but a limitation of operating model. We believe the next phase of security evolution must therefore transcend incremental automation, moving toward artificial intelligence (AI)-led autonomy rather than merely AI-driven task execution. By combining Google Cloud's hyperscale data infrastructure and Gemini-powered agentic reasoning with Accenture's deep operational expertise and managed security services, we are elevating the SOC from event-handling operations to continuous analytical decision-making capabilities.

Such a transformation represents a fundamental shift from SOCs that simply react to threats, to intelligent systems that can surmount complex security challenges — combining the adaptive thinking traditionally associated with human analysts with the speed and scale advantages of AI.

Accenture's Perspective

"A modern SOC is not defined by headcount but by intelligence. Outcome-driven autonomy replaces playbook-driven automation and turns analysts into risk strategists, not alert handlers."

Google Cloud's Perspective

"With Gemini and Google Cloud SecOps, we aren't just summarizing threats; we are deploying AI agents capable of planning, investigating, and remediating complex attacks with a speed and context that no human — or simple script — can match."

Against this backdrop, SOCs must evolve to meet these new challenges head-on.

2.

The SOC at an inflection point

Why SOCs must evolve

Today's SOCs face a perfect storm of challenges. They are caught in a pincer movement between the exponentially growing volume of security data and the accelerating velocity of sophisticated threats. This challenge has been building for decades, but the scale and urgency have reached critical proportions.

The data deluge challenge

Enterprise infrastructure continues to expand in both size and complexity, dramatically increasing the attack surface organizations must defend. The modern enterprise now manages hybrid and multi-cloud environments, Internet of Things (IoT) devices, remote workforces, and increasingly complex supply chain relationships — each introducing new vectors for potential compromise.

On the one hand, organizations are generating more security data than ever. On the other hand, they are struggling to derive actionable intelligence from it. A typical large enterprise may process millions of security events per day, with security teams tasked with investigating thousands of alerts. The challenge is not simply volume — it is the signal-to-noise ratio that determines whether critical threats are identified before they cause damage.

Key challenges include:

- Security tools remain disconnected, creating fragmented data silos that impede holistic threat visibility. Organizations often deploy over 50 security tools that do not communicate effectively, forcing analysts to manually correlate information across multiple consoles.

- Data volumes expand continuously, generating thousands of daily security alerts that overwhelm human capacity. Security analysts face what industry experts call “alert fatigue” — the desensitization that occurs when overwhelmed by constant notifications, leading to genuine threats being missed.

- False positives consume analyst time and further contribute to alert fatigue, reducing the effectiveness of security teams. Studies suggest that security teams spend up to 25 percent of their time investigating false positives, time that could be better spent on strategic security initiatives.

- Mean time to detect (MTTD) and mean time to respond (MTTR) remain stubbornly high despite investments in technology. Even with advanced tools, many organizations take days or weeks to detect sophisticated threats — time that adversaries will exploit to achieve their nefarious goals.

- The skills shortage compounds these challenges. The cybersecurity workforce gap continues to widen, with organizations competing for limited security talent. Even well-funded security programs struggle to recruit and retain experienced analysts capable of managing this complexity.

Executive summary
The SOC at an inflection point
The agentic leap
Reality vs. rhetoric
The convergence model: Building an autonomous security core
The road ahead
Conclusion

The limitations of traditional automation

While SOAR platforms have existed since 2014 and have moved beyond simple scripting, they remain fundamentally constrained by their design. Current automation solutions function primarily as sophisticated decision trees — rigid, brittle, and lacking true cognitive capabilities. When an attack deviates even slightly from predefined playbooks, these systems fail to adapt.

Consider a typical enterprise SOC managing more than 10,000 security events daily. Traditional SOAR platforms can automate ticket creation and basic enrichment, but when a sophisticated threat actor uses living-off-the-land techniques — leveraging legitimate system tools to blend in with normal traffic — the playbook-based approach struggles to distinguish malicious activity from legitimate operations.

Conversely, with agentic SOCs, there will be faster responses, fewer uninvestigated alerts, and a shift of analysts' time from manual triage to complex, strategic investigations. In one real-life case study of a managed provider that processed millions of alerts, AI-powered alert triage led to the automated closure of roughly 60 per cent of benign alerts, with false-negative rates kept near 1 percent, dramatically reducing analyst workload and alert fatigue.

Executive summary
The SOC at an inflection point
The agentic leap
Reality vs. rhetoric
The convergence model: Building an autonomous security core
The road ahead
Conclusion

3.

The agentic leap

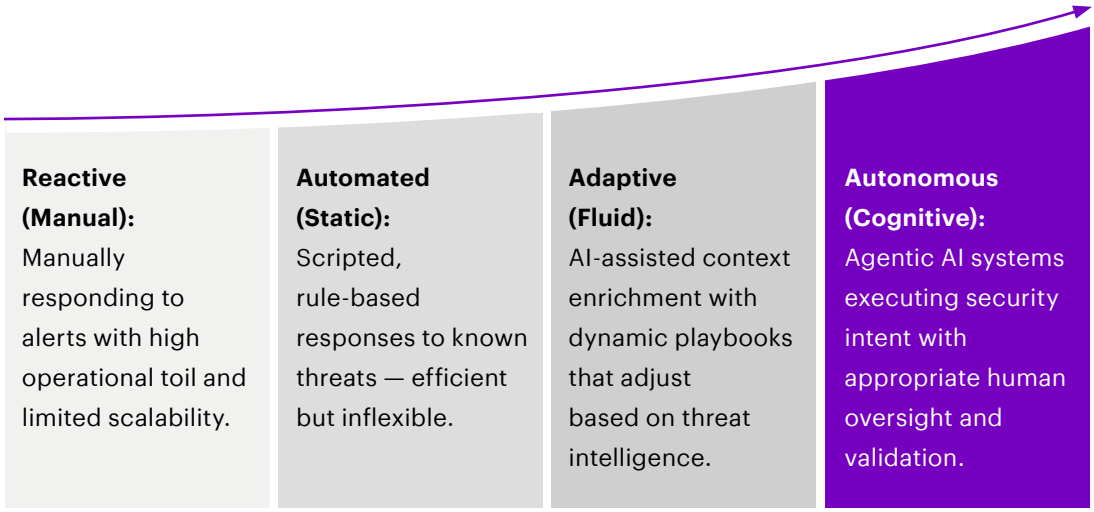
How SOCs should evolve

The future SOC must reason and adapt, not merely execute scripted responses. While human analysts have always brought reasoning capabilities to security operations, the challenge lies in scaling this cognitive approach to match the speed and volume of modern threats.

We must bridge the gap between human-defined static playbooks and adaptive, AI-guided operations. This transition to an Agentic Framework represents a paradigm where systems understand intent, context, and business impact — dynamically fine-tuning their approach based on the unique characteristics of each threat scenario.

The SOC evolution curve

Security operations have progressed through distinct evolutionary stages:



Defining the agentic SOC

An Agentic SOC represents a fundamental shift in architecture. Rather than a single tool, it functions as a mesh of specialized AI agents, each delegated with specific security responsibilities and operating with varying degrees of autonomy.

Unlike traditional SOAR playbooks that follow rigid “if-then” logic, an Agentic SOC receives high-level security objectives (for example, “detect lateral movement on Asset A”) and leverages reasoning capabilities to determine the optimal path to achieve these goals — while remaining aware of operational constraints.

These autonomous AI systems can independently identify, reason through, investigate, and resolve cyber threats with minimal human intervention — shifting the SOC from a reactive, manual alert-triage model to a proactive, intelligent defense system that operates at machine speed while preserving human judgment where it matters most.

The Google Cloud + Accenture difference

While traditional Security Information and Event Management (SIEM) and SOAR platforms rely primarily on correlation rules and signature-based detection, the Google Cloud SecOps architecture — powered by Gemini, Chronicle, and Mandiant threat intelligence — provides the foundational speed and intelligence capabilities required for real-time threat detection and extended detection and response (XDR).

Accenture leverages this technical foundation to architect and operationalize specific agent archetypes, each designed to address distinct aspects of the security operations workflow. Notable examples include:

- **The Triage Agent:** Google Cloud's SecOps Triage Agent automatically evaluates and prioritizes security alerts, reducing false positives and accelerating incident response. This capability is currently available in production within Google Cloud Chronicle SecOps.
- **The Detection Agent:** Continuously hunts for anomalies using Mandiant threat intelligence, dynamically adjusting detection logic based on emerging attack patterns — without requiring human intervention to update rules. This capability has been publicly demonstrated through Sec-Gemini, Google Cloud's AI-powered security assistant announced at Black Hat 2025.
- **The Context Enrichment Agent:** Autonomously queries disparate data lakes to construct comprehensive threat graphs, mapping relationships between users, assets, and attack timelines. This creates the contextual foundation necessary for informed decision-making.
- **The Cognitive Overseer Agent:** Synthesizes findings from multiple agent sources, evaluates risk and business impact, and recommends courses of action to human validators — explaining its reasoning and confidence levels.
- **The Response Policy Agent:** Executes containment and remediation actions within predefined guardrails, ensuring that automated responses align with organizational risk tolerance and compliance requirements.

Executive summary
The SOC at an inflection point
The agentic leap
Reality vs. rhetoric
The convergence model: Building an autonomous security core
The road ahead
Conclusion

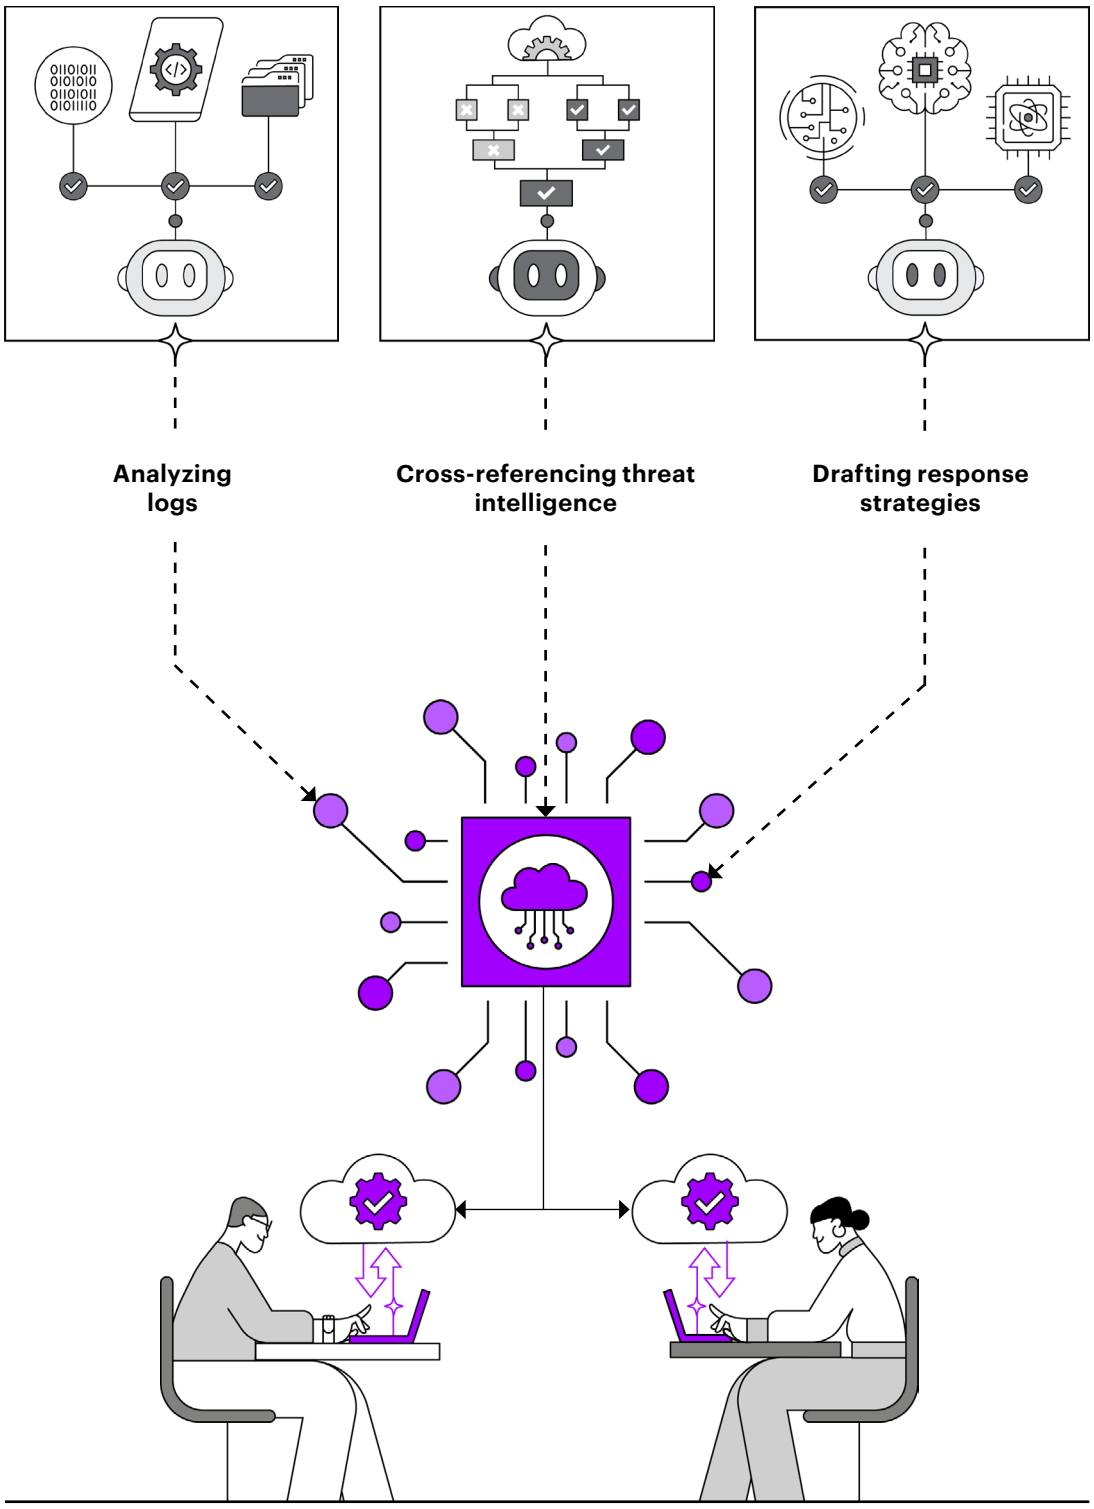


Figure 1: AI workflows converging into a central reasoning engine enabling action with human oversight.

4.

Reality vs. rhetoric

What types of SOCs are possible today?

> **"Every revolution starts as an exaggeration before it becomes reality. The trick is knowing which exaggerations are becoming true."**

To ground this vision in practical reality, we map current capabilities against the autonomous future, acknowledging both achievements and limitations:

Claim	Current capability	Limitation	Outlook
AI can triage alerts autonomously.	Achievable today via ML classification + LLM correlation.	Contextual errors and data silos can lead to inappropriate prioritisation or hallucinations.	**Will mature significantly with agentic context modeling and graph-based reasoning.**
SOCs will run with humans as strategic validators.	Unrealistic for most organizations today.	Significant governance, accountability, and trust gaps must be addressed.	**Achievable through tiered autonomy models.**
Threat intelligence enrichment can be automated.	Mature capability — widely deployed today.	Requires trust scoring and validation mechanisms to avoid false positives and low-confidence alerts.	**Continuous improvement expected through AI graph models.**

This realistic assessment helps organizations set appropriate expectations and plan incremental adoption pathways toward agentic capabilities.

5.

The convergence model: Building an autonomous security core

How to build a SOC of the future

To build a truly adaptable SOC, we must converge three foundational pillars: AI Reasoning, Zero Trust Architecture, and Unified Observability. These elements work together to create an intelligent and resilient security infrastructure.

The architecture of wisdom

This model transcends linear data processing, embracing graph-based knowledge representation that mirrors how security experts think about threats. The system leverages federated learning across the enterprise and integrates global threat intelligence from Mandiant to create a "collective defense" posture — where insights gained from one incident strengthen defenses across the entire organization.

Tiered Autonomy and human evolution

We propose a Tiered Autonomy Model that recognizes different contexts require different levels of human involvement. This framework consists of three operational modes: human-in-the-loop, human-on-the-loop, and human-above-the-loop.

The operational flow follows this pattern:



Most critically, the human role evolves from "Operator" to "Strategic Validator." AI agents propose sophisticated response strategies, while humans validate the business impact, ethical considerations, and alignment withorganizational risk appetite.

	Operational Mode	Example
Human-in-the-Loop	**Automation assists analysis and provides recommendations, but humans decide and execute all actions. Best for high-risk or ambiguous scenarios where business context and human judgment are critical.**	A malware alert is enriched with threat intelligence, attack path visualization, and potential business impact analysis. The AI agent provides a detailed briefing and recommended containment strategy, but the senior analyst makes the final decision on whether to isolate the affected system, considering operational dependencies and business requirements.
Human-on-the-Loop	**Automation executes predefined responses within established guardrails. Humans supervise operations and intervene only when necessary. Best for repeatable, well-defined risk scenarios with clear decision criteria.**	Phishing campaigns are autonomously detected, triaged, and analyzed. The AI agent automatically blocks malicious URLs, quarantines suspicious emails, and updates threat intelligence feeds. Analysts monitor aggregate metrics and exception reports, intervening only when the AI encounters edge cases or confidence thresholds fall below acceptable levels.
Human-above-the-Loop	**Automation reasons, decides, and acts autonomously. Humans focus on strategic oversight, risk outcomes, and continuous improvement of the AI systems. Best for mature security environments with established trust, robust controls, and comprehensive governance frameworks.**	An AI-driven risk engine continuously monitors the threat landscape and organizational security posture. It dynamically adjusts security controls based on real-time threat intelligence, business context (such as elevated risk during product launches or financial reporting periods), and compliance requirements. Security leaders review strategic metrics, risk trends, and AI decision-making patterns rather than individual incidents.

6.

The road ahead

What comes next?

The transformation toward agentic security operations is already underway in leading organizations. As we look to the immediate future, the Agentic SOC could continue evolving into increasingly sophisticated capabilities:

- ### Multi-agent collaboration ecosystems

  Rather than isolated agents operating independently, next-generation systems could feature agents that collaborate, debate competing hypotheses, and refine threat assessments through multi-perspective analysis before escalating findings to human validators. This deliberative approach mirrors how expert security teams discuss complex incidents, combining diverse viewpoints to reach more accurate conclusions. For instance, when detecting unusual network traffic, a “Detection Agent” flags potential data exfiltration. A “Context Agent” analyzes user behavior history and determines the activity aligns with an authorized data migration project. A “Risk Assessment Agent” evaluates the conflicting signals and recommends temporary enhanced monitoring rather than immediate containment. The “Cognitive Overseer Agent” then synthesizes these perspectives and presents a nuanced risk assessment to the human analyst.

Predictive security operations

Leveraging digital twin technology, advanced SOCs could simulate sophisticated attack scenarios against virtual replicas of production environments. These simulations enable organizations to identify and remediate vulnerabilities before adversaries can exploit them — shifting from reactive incident response to proactive vulnerability management. This approach transforms security testing from periodic assessments to continuous validation, ensuring that security controls remain effective as infrastructure evolves and new threats emerge.

Reinforcement learning loops

The most sophisticated future systems could implement continuous feedback loops that mirror adversarial machine learning — where “Red Agents” simulate attacker tactics, “Blue Agents” develop and refine defensive strategies, and “AI Auditors” evaluate the effectiveness and appropriateness of both offensive simulations and defensive responses. These reinforcement learning systems enable security capabilities to evolve at the same pace as threat actor techniques, ensuring defenses remain effective against emerging attack patterns. This is not merely an incremental upgrade to existing security infrastructure — it represents a fundamental evolution in how organizations approach cyber defense, blending human strategic insight with AI-powered speed and scale.

As security operations become increasingly autonomous, governance and trust become paramount. Cybersecurity frameworks like MITRE D3FEND provide standardized approaches to defensive techniques, ensuring that autonomous actions align with established security principles and industry best practices. Organizations must also establish comprehensive programs that address:

- **Explainability:** AI agents must provide clear reasoning for their decisions, enabling human validators to understand and verify autonomous actions.
- **Accountability:** Clear ownership and responsibility chains for AI-driven security decisions.
- **Auditability:** Comprehensive logging and forensic capabilities to review autonomous decisions and actions.
- **Safety:** Guardrails that prevent autonomous systems from taking actions that could disrupt business operations or violate compliance requirements.
- **Dependability:** Regular assessment of AI agent performance, accuracy, and alignment with overall objectives.

> **"Autonomy in security isn't the absence of humans — it's the presence of wisdom in the loop."**

Trust remains the foundation of all security operations — whether human-led or AI-assisted. The most successful agentic SOCs will likely be those that earn and maintain stakeholder confidence through transparency, reliability, and demonstrable value.

7.

Conclusion

The partnership between Accenture and Google Cloud is dedicated to delivering this future — moving clients from a state of constant alarm and reactive firefighting to one of confident resilience and strategic security management.

By combining Google Cloud's world-class security infrastructure, threat intelligence, and AI capabilities with Accenture's deep operational expertise and global managed security services, we aim to create security operations centers that can learn, adapt, and evolve at the speed of threat — while preserving the strategic judgment and ethical oversight that only humans can provide.

The transformation we envision is not a distant aspiration. It is starting now in leading organizations around the world. Early adopters are already seeing measurable improvements in detection accuracy, response times, and analyst productivity. More importantly, they are building security capabilities that scale with, rather than bottleneck, their business growth.

Organizations that embrace this evolution from automation to autonomy — balancing innovation with governance, speed with safety, and autonomy with accountability — could be better positioned to not only defend against today's threats but adapt to meet tomorrow's troubles. Those that hesitate risk falling further behind adversaries who are themselves leveraging AI and automation to increase the sophistication and deadliness of their attacks.

We look forward to a future of security operations that is intelligent, adaptive, and autonomous. The question should not be whether to embrace agentic security operations, but how quickly and effectively your organization can begin this transition.

> **"The SOC of the future isn't just automated — it's alive, learning, and accountable."**

About Accenture

Accenture is a leading solutions and services company that helps the world's leading enterprises reinvent by building their digital core and unleashing the power of AI to create value at speed across the enterprise, bringing together the talent of our approximately 786,000 people, our proprietary assets and platforms, and deep ecosystem relationships. Our strategy is to be the reinvention partner of choice for our clients and to be the most client-focused, AI-enabled, great place to work in the world. Through our Reinvention Services we bring together our capabilities across strategy, consulting, technology, operations, Song and Industry X with our deep industry expertise to create and deliver solutions and services for our clients. Our purpose is to deliver on the promise of technology and human ingenuity, and we measure our success by the 360° value we create for all our stakeholders.
Visit us at **accenture.com.**

About Google Cloud

Google Cloud is the new way to the cloud, providing AI, infrastructure, developer, data, security, and collaboration tools built for today and tomorrow. Google Cloud offers a powerful, fully integrated and optimized AI stack with its own planet-scale infrastructure, custom-built chips, generative AI models and development platform, as well as AI-powered applications, to help organizations transform. Customers in more than 200 countries and territories turn to Google Cloud as their trusted technology partner.
Visit us at **cloud.google.com/security**

Disclaimer: The material in this document reflects information available at the point in time at which this document was prepared as indicated by the date provided on the front page, however the global situation is rapidly evolving and the position may change. This content is provided for general information purposes only, does not take into account the reader's specific circumstances, and is not intended to be used in place of consultation with our professional advisors. Accenture disclaims, to the fullest extent permitted by applicable law, any and all liability for the accuracy and completeness of the information in this document and for any acts or omissions made based on such information.

Accenture does not provide legal, regulatory, audit, or tax advice. Readers are responsible for obtaining such advice from their own legal counsel or other licensed professionals. This document refers to marks owned by third parties. All such third-party marks are the property of their respective owners. No sponsorship, endorsement or approval of this content by the owners of such marks is intended, expressed or implied.

Copyright © 2026 Accenture. All rights reserved.
Accenture and its logo are registered trademarks of Accenture.