

# The Agentic SOC

## Your Teammate for Tomorrow, Today

Today's security operations centers (SOCs) are caught between the explosion of AI-driven threats and the persistent friction of their own fragmented tools and manual processes. The result is a hidden tax on their most valuable resource—the mission-driven team keeping the lights on.

New research from Microsoft reveals the stark reality of this operational gap. It's near-universal, with **91% of security leaders**<sup>1</sup> confirming these gaps have directly caused a serious, business-disrupting incident. Their analysts spend a full day each week—more than **20% of their time**<sup>2</sup>—manually aggregating data. Drowning in this manual toil, teams leave **42% of security alerts**<sup>3</sup> uninvestigated. Each ignored alert is a potential foothold, allowing an attacker to operate undetected and turn a minor issue into a major breach.

This model grows more unsustainable by the day, demanding a fundamental operational shift. That shift is not just arriving—it's being led by a new class of "Frontier Firms," organizations that, according

to [Microsoft's 2025 Work Trend Index](#), are already rebuilding their businesses around AI. In security, these firms are defined by their move toward the agentic SOC, where semi-autonomous agents integrate directly into security workflows.

This isn't a future concept; it's an operational reality today. As proof, Microsoft's own internal security team now uses custom AI agents to autonomously investigate **75% of incoming incidents**.<sup>4</sup> This is the pivotal change: moving from manual analysis to automated containment, freeing human experts to focus on what they do best—architecting defenses and hunting for novel threats.

This white paper provides a foundational roadmap for this transformation. For practitioners, it promises a shift from reactive toil to high-impact work. For leaders, it offers a clear point of view on how to navigate this change—from building a unified data foundation to cultivating the human-AI teams of tomorrow.



### Contents

- The agentic imperative
- Stages of agentic SOC maturity
- Building the agentic operating model
- Choosing your agentic pathway

<sup>1</sup>"State of the SOC: Unify Now or Pay Later: What New Research Reveals," page 8, Microsoft, 2026

<sup>2</sup>Ibid, page 6

<sup>3</sup>Ibid, page 7

<sup>4</sup>Based on Microsoft research data. Accurate as of February 2026

# The agentic imperative

The timeline for SOC modernization has been dramatically compressed by three powerful forces:



## Industrialized AI threats

Adversaries now scale their attacks with AI, using generative tools to craft convincing phishing lures and automate reconnaissance, which drastically increases the volume, speed, and efficacy of sophisticated social engineering attacks.



## The shifting cloud and agent attack surface

Cloud infrastructure has become the primary target—and the attack surface is growing fast. As SaaS apps, AI agents, and multi-cloud deployments multiply, attackers gain new entry points that bypass the endpoint entirely. By compromising cloud accounts or agent access directly, they exploit blind spots that device-centric tools like EDR simply weren't built to see.



## Rising regulatory pressure

In response to these new risks, boards and regulators now demand stronger governance, documented human oversight, and rapid, accurate incident reporting.

These forces have rendered traditional SOC models—characterized by siloed tools and rigid, manual playbooks—insufficient. Security teams are overwhelmed, and **75% of security leaders<sup>5</sup>** are concerned their SOC cannot keep pace with new and emerging threats.

The agentic model directly addresses these pressures by scaling the SOC's ability to detect, decide, and act at the speed modern threats require. Agents absorb the operational load created by industrialized attacks, automate routine investigation across cloud and identity surfaces where traditional tools struggle, and strengthen governance with consistent, auditable actions. This operational shift elevates human analysts from reactive triage to high-impact strategic work.



### FROM THE FRONTLINES

## Focus on what matters

Agents automate the noise of triage—correlating data and filtering alerts—to free you for consequential work: investigating complex attacks and hunting novel threats.

<sup>5</sup>"State of the SOC: Unify Now or Pay Later: What New Research Reveals," page 9, Microsoft, 2026

# Stages of agentic SOC maturity

The evolution toward an agentic SOC is not a series of sequential steps on a linear roadmap, but a progression through three interconnected stages in a dynamic framework. Progress in one stage directly accelerates the value of the others, creating a virtuous cycle of modernization where better data enables smarter AI, and smarter AI reveals new opportunities for data enrichment. In today's threat landscape, waiting to perfect one stage before beginning the next is a strategic delay the modern SOC cannot afford.

## SOC I

### Unify your platform foundation

A unified security platform is the architectural bedrock of the agentic SOC. Without it, agents become just another digital silo, trained in isolation and unable to correlate insights across domains. The real value of unification is what it stops: the endless "swivel chair" investigations where analysts pivot between consoles just to piece together an attack. By integrating signals from identity, endpoints, and cloud infrastructure, you provide the holistic context that both humans and AI need for decisive, cross-domain action. This architectural principle is what separates experimental automation from production-ready security operations.

## SOC II

### Accelerate operations with generative AI

The focus of this stage is to elevate the SOC by embedding generative AI to deliver instant context, accelerate analyst decision-making, and streamline high-volume workloads. Synthesizing signals across the entire security fabric—from cloud and identity to endpoints—cuts through noise, improves accuracy, and empowers analysts to focus on higher-impact investigations. The result is a more efficient and confident SOC that establishes the operational patterns and human-AI teaming models required for the next stage.

## SOC III

### Deploy agentic automation

This final stage marks the shift from AI assistance to true agentic automation. Here, specialized AI agents advance from augmenting individual workflows to autonomously orchestrating specific tasks—like isolating compromised devices or triaging user-reported phishing—and, over time, entire complex scenarios such as simulating attacks and coordinating multi-step responses. This elevates human analysts to strategic supervisors, focused on directing agent performance, guiding security initiatives, and building organizational expertise in autonomous operations.

# Building the agentic operating model

With the technology stages defined, the focus shifts to the operating model. This means building new capabilities for the people who will supervise agents and the governance processes that build trust in automation.

## Cultivating your team for the agentic age

As AI begins to handle more frontline analysis, security teams are redefining what human expertise looks like—prioritizing judgment, validation, and oversight over repetitive execution. These emerging trends are reshaping how teams learn, grow, and build trust with their new digital teammates.

This evolution elevates human expertise, creating new, critical roles for experienced analysts: **orchestrators** who design and direct fleets of AI agents, **validators** who review and approve AI-driven actions, and **truth-setters** who handle novel incidents that require deep human intuition.



FROM THE FRONTLINES

### From alert-handler to agent-orchestrator

The agentic model elevates your role from reacting to the queue to supervising it. You get to focus on hunting novel threats and solving complex problems—the work only a human expert can do.

However, this presents a new challenge: the mundane, repetitive tasks that AI excels at are the core learning ground for junior analysts. Leaders must define a new path for AI-assisted analyst development, creating training programs that prevent skill degradation and build the next generation of “truth-setters.”

Forward-thinking organizations are responding by creating AI-augmented learning environments where junior analysts train alongside agents—learning to validate AI decisions, identify edge cases, and develop judgment through supervised AI interactions rather than purely manual analysis.

## Using guardrails to build trust in agentic operations


Effective use of agentic AI hinges on verifiable control. Building that control requires a governance framework aligned with emerging requirements—from [NIST’s AI Risk Management Framework](#) to the EU AI Act’s [high-risk classifications](#).

As detailed in [Microsoft’s Guide for Securing the AI-Powered Enterprise: Data Governance and Security](#), effective agent governance applies proven data lifecycle security principles to a new class of non-human identities through three essential principles:



### Treat agents as digital employees

Every agent requires a human “*agent boss*” responsible for directing its work and ensuring alignment with business goals. Security teams must build comprehensive inventories of these digital workers, verify each agent’s access and system interactions, and enforce least-privilege permissions narrowly scoped to specific tasks.



### Define agent authority by risk

Classify actions into three tiers—autonomous (low-risk tasks like alert enrichment), human-approved (medium-risk actions like account isolation), and human-executed (high-risk decisions like breach declarations). This framework aligns the required level of human oversight to the risk of each action.



### Implement verifiable decision trails

Every agent action must generate an immutable audit record capturing what decision was made, what data informed it, agent confidence levels, and whether human review occurred. This satisfies regulatory requirements while enabling continuous performance improvement.

This framework of trust, however, cannot exist in a vacuum. It must be supported by an integrated technology architecture that provides the [visibility and control necessary](#) for safe, effective agentic operations.

# Choosing your agentic pathway

The journey to an agentic SOC begins with a foundational strategic decision: whether to adopt turnkey agents for immediate value or invest in a studio environment to build custom solutions. We engineer our platform to support both paths, providing a flexible framework that scales with your organization's maturity.

## The turnkey path

For organizations seeking immediate efficiency gains, the roadmap focuses on deploying pre-built agents to automate common, high-volume tasks.

- **Identify high-volume toil**  
Start by pinpointing the most repetitive, time-consuming tasks in your SOC, such as manually triaging user-reported phishing emails or enriching alerts with basic context.
- **Deploy pre-built agents**  
Deploy Microsoft's dozens of [cross-ecosystem agents](#) to address that toil. For example, the Phishing Triage Agent can autonomously analyze user submissions, dismiss benign messages, and escalate credible threats, while the Threat Intelligence Briefing Agent can deliver instant summaries of new campaigns or vulnerabilities on demand.
- **Reinvest analyst capacity**  
With manual work reduced, redirect your human experts to higher-value initiatives like proactive threat hunting, tuning detection rules, and hardening defenses against emerging risks.

## The studio path

For mature organizations wanting to tailor AI to their specific workflows, the roadmap focuses on building bespoke solutions using governed frameworks.

- **Target a unique business workflow**  
Identify a security challenge unique to your organization. For a healthcare provider, this might be a custom agent that monitors for anomalous access to patient record systems. For a financial firm, it could be an agent designed to detect insider trading patterns by correlating data access with external communications.
- **Build and deploy with governance**  
Use frameworks like [Microsoft Agent 365](#) to construct your custom agent within secure guardrails. Assign it a unique identity, enforce least-privilege permissions scoped only to its task, and ensure every action is auditable.
- **Iterate and expand**  
Use the success of your first custom agent to build organizational expertise. Gradually expand its capabilities or identify the next unique workflow to automate, turning your SOC into a center for security innovation.



FROM THE FRONTLINES

### An instant win against alert fatigue

Reclaim hours from the phishing queue. Turnkey agents automate the repetitive triage of user-reported emails, freeing you to focus on hunting threats instead of chasing noise.

## Combining paths and setting expectations

Crucially, these paths are not mutually exclusive. Many organizations begin with turnkey agents and evolve toward custom builds as they mature. Regardless of the starting point, governance comes first. Leaders must define a minimum agent standard with a documented purpose, owner, and allowed actions before deploying any agent, pre-built or custom.

Most importantly, leaders must define what "end-to-end" means for their organization to set clear expectations. Even a turnkey solution requires the tuning and oversight that ensure trust. The goal is not to remove humans, but to free them for the work that matters most.

# Your partner for the new era of security

The agentic SOC represents a fundamental, and necessary, shift in security operations. Navigating this new era requires more than just adopting tools; it demands a new operating model, new skills, and a new partnership between human expertise and machine intelligence.

Microsoft not only understands this new landscape, but we have a clear, trustworthy plan for the future—a plan proven at scale within our own SOC. Our approach provides the unified platform to see across your entire digital estate, the production-grade agentic frameworks to act with speed and precision, and the strategic guidance to help you build an adaptive foundation of trust and governance. This is the future of resilient, agent-boosted defense, and Microsoft is the strategic partner to help you build it with confidence.



[Learn more about our AI-enabled SOC solutions](#)



[Get the latest agentic SOC content](#)