From ‘AI everywhere in the SOC’ to ‘AI where it counts’

How selecting the right Microsoft Security Copilot use cases can enable more standardized, efficient, and effective security

Capgemini

AI presents a growing cybersecurity threat – and a crucial opportunity

Generative AI and LLMs are helping all kinds of enterprise functions operate faster and more efficiently. Unfortunately, they're also helping bad actors work faster, too.

The average time to exploit (TTE) shows the impact of LLMs and AI on cybersecurity. In 2022, before the full public launches of many common AI tools, the average TTE was 9.7 months. By 2026, ***that number had fallen to just nine hours***[1]. And estimates suggest this will continue, with AI agents reducing the time ***to exploit by 50% by 2027***[2].

Many organizations will attempt to fight fire with fire by deploying AI-powered cybersecurity tools. Given that AI features are bundled with other platforms they're already playing for, it certainly seems like a logical approach. But in practice, indiscriminate or unmanaged AI use can create more problems around cost, complexity, and risk. governance and operational consistency.

In this guide, we'll explore:

- The challenges of implementing AI in the modern security operations center (SOC),
- Key considerations for applying AI to a SOC use case
- How Capgemini's AI-Powered SOC, built on Microsoft technologies, delivers pre-built automations and AI prompts to create a more consistent and efficient cybersecurity operation.

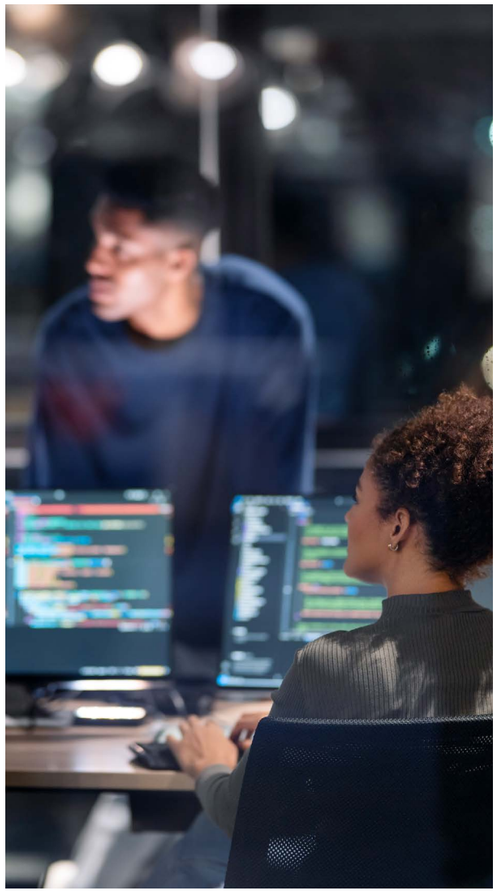

[1]Zero Day Clock, From Vulnerability to Exploitation Dashboard.
[2]Gartner Predicts 2025: Navigating Imminent AI Turbulence for Cybersecurity.

AI is a huge opportunity, but it's not a SOC cure-all

The simple truth is that AI can be a powerful tool when applied to the right task, but it's not a cure-all – particularly in the SOC. Indiscriminate application of AI can be particularly damaging for SOCs for four reasons:

AI resources are limited in security operations

Every time a SOC deploys AI to handle simple tasks, fewer resources are available for more complex analysis.

Role-based automation can be more efficient

Deterministic tasks such as enrichment, correlation, and routing are often poor fits for AI and are better suited to traditional automation.

Excessive AI use can complicate audits and governance

AI's probabilistic nature leads to inconsistent outcomes and processes that are difficult to monitor and observe.

Not every organization is ready for AI

AI doesn't transform bad processes into good ones; it only scales what you currently have. That means some teams may find they get worse results from AI if their SOC isn't built on strong foundations.

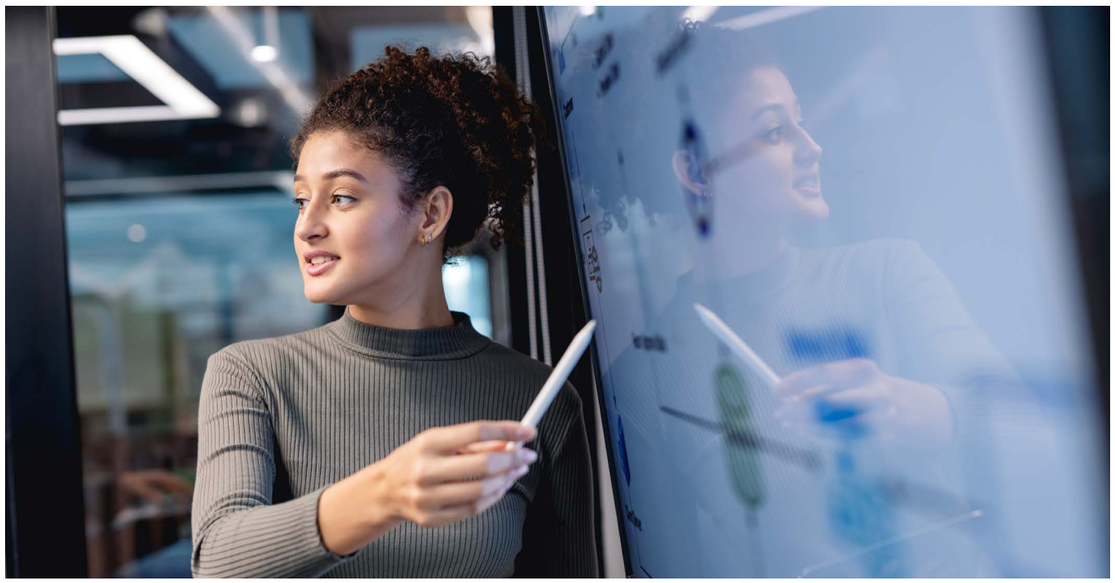

Getting ready for AI: Establishing Zero trust principles

AI can't exist in the SOC without governance. Any SOC AI journey must start by establishing the right controls to secure core infrastructure and ensure your SOC sends usable signals rather than noise. Zero Trust principles are a great starting point.

Microsoft sets out the steps for establishing *Zero Trust* as follows:

- **Verify explicitly** and always authenticate and authorize
- **Use least privilege access** and limit user access with Just-In-Time and Just-Enough-Access (JIT/JEA)
- **Assume breach,** minimize blast radius, and segment access. Verify end-to-end encryption and use analytics to get visibility, drive threat detection, and improve defenses.

A Zero Trust assessment can instill these principles throughout your SOC to harden and normalize upstream control points by clarifying trust boundaries, reducing standing access, governing non human identities, and improving telemetry quality and consistency across your environment.

This ensures AI solutions like *Microsoft Security Copilot* connect to systems that produce higher-fidelity signals, enabling AI to reliably triage, correlate, and respond.

With Zero Trust principles in place, the AI powered SOC can safely conduct workflow decisions at scale rather than amplifying existing weaknesses.

Four criteria for assessing Microsoft Security Copilot use cases

Once you've established your Zero Trust baseline, it's time to identify which tasks are appropriate for AI and which should be left to humans or conventional automation.

We'll be looking at this from the perspective of deploying Microsoft Security Copilot, as it's a tool many organizations will already have access to and can integrate with existing Microsoft tools like Microsoft Defender, Entra, Intune, and Purview.

There are four key criteria to consider when assessing potential use cases for tools like Microsoft Security Copilot:

Determinism

Does the task have a clear, repeatable logic that would be better served by traditional automation?

Judgement

Does the use case involve synthesizing, summarizing, or sorting large quantities of unstructured data?

Auditability

Can you ensure AI prompts and outputs are traceable and repeatable?

Frequency

How regularly will AI functions need to be called to complete the task, and how much will that cost?

By assessing your use cases across these four areas, it will quickly become clear where Microsoft Security Copilot will yield strong results and where other approaches might be more suitable.

For example, a high-frequency, highly deterministic task is better suited to automation, while a low-frequency task with high judgment requirements will be ideal for Microsoft Security Copilot or other AI cybersecurity tools. In cases where you need some judgment in high frequency, it may be better to use AI with strong gates in place and reuse outputs.

Vulnerability management use case assessment in practice

	Use Case	Determinism	Judgment	Auditability	Frequency
1	Exploitability Context Enrichment	3	4	4	4
2	Vulnerability – Identity Exposure Mapping	4	2	4	5
3	Vulnerability – Internet Exposure Assessment	4	2	5	5
4	Patch Priority Scoring with Business Context	2	5	4	4
5	KEV Alignment & Risk Justification	3	4	5	3
6	Lateral Movement Risk from Vulnerability	2	5	4	2
7	Vuln – Detection Coverage Gap Analysis	4	3	5	3
8	False Positive Vuln Validation	3	4	3	5
9	Vulnerability Recurrence Root Cause Analysis	2	5	4	1
10	Executive Vulnerability Exposure Brief	1	5		

SOC use case decision quadrant

Where Microsoft Security Copilot vs automation fits

High Judgment (Synthesis & Context)

Low Judgment (Simple & Repetitive)

AI on-demand / Advisory
(Low Frequency, High Insight)
- **09** Vulnerability Recurrence RCA
- **10** Executive Vulnerability Brief
- **06** Lateral Movement Risk from Vuln

Microsoft Security Copilot - Strong fit
(Governed Prompts, Human Gates)
- **04** Patch Priority + Business Context
- **01** Exploitability Context Enrichment
- **05** KEV Alignment & Justification
- **08** False Positive Vuln Validation*

Hybrid: Automation + AI narrative
(AI Explains Gaps)
- **07** Vuln - Detection Coverage Gaps

Automate - Not AI
(High Frequency, Rules-Based)
- **02** Vuln - Identity Exposure Mapping
- **03** Vuln - Internet Exposure Assessment

*After Automation Filters

Low Determinism (Ambiguous) → **High Determinism** (Rules-Based)

- Use AI where analysts debate, not where systems calculate.
- High Frequency + Deterministic = Automate.
- High Judgment + Low Frequency = Microsoft Security Copilot on Demand.

Generate the Priority Order

Move Now
- Patch priority scoring with business context
- Exploitability context enrichment
- KEV alignment and risk justification

Move with Controls
- False positive vulnerability validation
- Detection coverage gap analysis

Automate First
- Vulnerability - Internet exposure
- Vulnerability - Identity exposure

AI On-Demand Only
- Executive vulnerability exposure briefing
- Recurrence root cause analysis
- Lateral movement risk

Rethinking prioritization: moving beyond "critical high/critical medium"

Traditional SOC prioritization models rely heavily on static severity labels such as "critical", "high", and "medium". While useful, these models were designed for a pre-AI world and struggle to reflect the speed, scale, and complexity of today's threat landscape.

The problem with these severity scores is that they're static, while modern threats evolve rapidly. That means severity scores often fail to account for exploitability, exposure, or business context, and can leave SOC teams overwhelmed with large volumes of "high" alerts that still require manual triage.

An AI-powered SOC enables a shift from severity-based prioritization to context-driven decisions. Rather than trying to work out how critical an alert appears on the surface, AI allows SOC teams to dynamically prioritize risk by combining multiple dimensions of context:

- **Exploitability context:** is the vulnerability actively exploited or aligned to known exploited vulnerabilities?
- **Exposure surface:** is the asset internet-facing or accessible through identities or privileged access?
- **Business criticality:** what is the operational or financial impact if the asset is compromised?
- **Attack path potential:** can this issue enable lateral movement or chaining across the environment?
- **Detection and response readiness:** can existing controls detect and contain this threat, or are there gaps?

By synthesizing these factors, AI can help SOC teams continuously reprioritize work based on real risk, rather than static scoring models. This shift in thinking aligns directly with how Microsoft Security Copilot and automation should be applied within the SOC. Instead of relying solely on severity labels, organizations can adopt decision-led prioritization categories:

Act immediately (AI + automation):

High exploitability, high exposure, and significant business impact

Example
KEV-aligned vulnerabilities on internet-facing assets

AI-assisted investigation:

Scenarios with complex or uncertain context requiring synthesis across multiple signals

Example
vulnerabilities with potential lateral movement or unclear business impact

Automate and monitor:

High-frequency, low-judgment scenarios suited to deterministic automation

Example
enrichment, exposure checks, or identity mapping

On-demand analysis:

Lower priority or infrequent scenarios where AI is used selectively

Example
executive reporting or deep root cause analysis

How to find an efficient, consistent way forward with Microsoft Security Copilot use cases

Even when you identify a clear Microsoft Security Copilot use case for the SOC, there are still two major issues to overcome:

1. How do you ensure consistency with a technology that generates probabilistic outputs?
2. How do you scale your use case and ensure AI guards are replicable?

Both challenges stem from the same issue: when analysts create their own prompts for AI tools, costs and results are inconsistent.

That's where a prompt and automation library approach can help.

With a library of pre-built prompts, analysts can select from pre-defined inputs covering a range of use cases. That means every prompt is optimized to keep costs low while delivering consistent value.

And by combining prompt libraries with a collection of pre-designed automations, you can scale operations to complete deterministic, high-volume tasks with automation, and automate tasks that require high levels of judgment with AI.

For example, you can ensure AI provides consistent reporting to executives on exposure and threats while using pre-built automations for enrichment, suppression, correlation, and other routine data tasks.

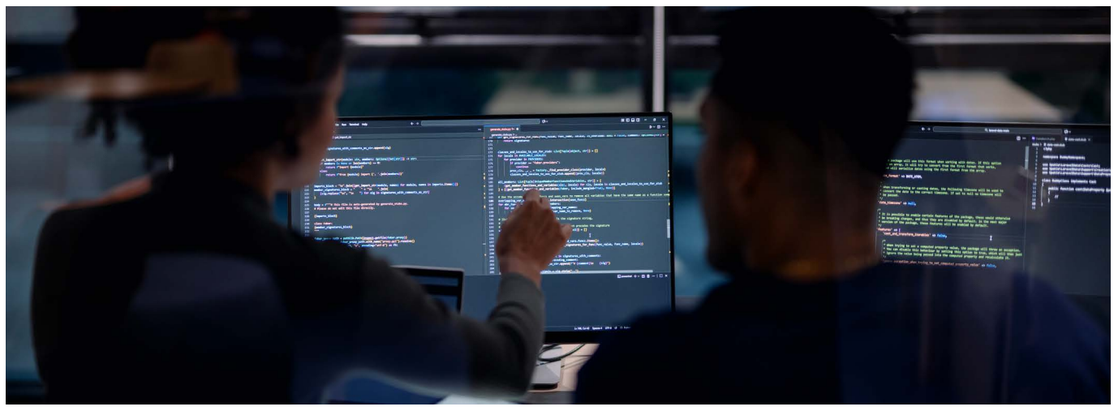

AI-Powered SOC transformation

Zero trust is relevant for all AI, Microsoft Security Copilot or Microsoft 365 copilot

Key Deliverables

- Repeatable delivery for Zero Trust gap assessment summary and findings report
- Domain-specific workshops and prioritized roadmap to align with Microsoft Workshops
- SOW's and delivery process for implementing golden policies & best practices
- MDR alignment and operationalization recommendations

Value to Clients

- Security posture baseline
- Clarity on gaps by domain pillars
- Actionable roadmap for remediation
- Securing AI implementations
- Tool hardening

Domain Pillars

- Identity
- Devices
- Data
- Network Infrastructure
- Security Operations

Capgemini's AI-Powered SOC offering includes pre-built workflows using Microsoft Security stack with Microsoft Security Copilot to offer a practical pathway to AI SOC use cases.

It starts with a Zero Trust workshop to assess gaps and prioritize roadmaps around your unique domain needs, and offers a process for implementing SOC best practices. This can help you:

- Enhance your security posture baseline
- Clarify gaps by domain pillars
- Create a clear remediation roadmap
- Secure your AI implementations

Once you've established your baseline, you can use Microsoft Security Copilot to power AI SOC use cases. By combining Microsoft Security Copilot with automations and prompt libraries from Capgemini, you get repeatable results while making the most of your AI resources.

Start your AI-powered SOC journey

To meet AI-powered threats, organizations need an AI-powered SOC. But that doesn't mean putting AI into every process. By establishing a strong operational foundation and using AI to scale where it makes sense, you can make the most of your AI credits while leaving deterministic tasks to proven automation.

The right assessment criteria can help you make smart decisions about AI deployment, but these decisions are best supported by extensive experience and proven security tools.

Together, Capgemini and Microsoft bring both.

Microsoft Security Copilot offers out-of-the-box agent experiences that can now integrate with common security workflows using Azure Logic Apps.

Capgemini's security experience and automation and prompt libraries help you make the most of these AI tools while supplementing them with clear gating and automations that deliver even more consistent, efficient security operations.

With a global network of Cyber Defense Centers, deep sector expertise, and thousands of cybersecurity experts worldwide, Capgemini empowers SOCs to scale up and meet the growing threat landscape head-on. Capgemini has been recognized as a leader in the ISG Provider Lens™ Cybersecurity – Services and Solutions report for five years consecutively and has also been recognized as a leader in Avasant's Cybersecurity Services 2025 RadarView™ report.

Author

Mona Ghadiri

Vice President, Global Offer Lead for Cybersecurity Defense

Cloud Infrastructure Services

mona.ghadiri@capgemini.com

For more practical advice on how to create your own AI-powered SOC

Read our eBook **Talk to our experts**

For more details, contact:

cybersecurity.in@capgemini.com

About Capgemini

Capgemini is the business transformation partner for enterprises in the age of AI. We help organizations imagine and build an intelligent, sustainable future, combining AI, technology and human ingenuity to transform how they operate, innovate and grow. With unique end-to-end capabilities spanning strategy, technology, engineering and intelligent operations, we bring together deep industry expertise and market-leading capabilities in AI, cloud and data to turn ambition into measurable business outcomes at scale. Supported by a robust ecosystem of partners and nearly 60 years of expertise, Capgemini is a responsible and diverse global organization of over 410,000 team members in more than 50 countries. The Group reported 2025 revenues of €22.5 billion.

Make it real.
www.capgemini.com

GIS_04.06.26_SC