

How AI is Evolving and What Your SOC Needs to Do About It

6 priorities for productivity,
customization, and control

splunk>
a CISCO company



AI won't run your SOC — but it can make it a whole lot smarter

AI is changing everything — including how security teams operate. From accelerating triage to surfacing hidden threats, AI is beginning to **transform the SOC**. But let's be clear: the hype is loud, the technology is fast-moving, and the pressure to “do something with AI” is real.

The good news is that you don't need to build custom models or deploy next generation tools to see results. Many security teams already have AI-powered capabilities in their SIEM, SOAR, and data platforms. They just haven't fully tapped into them.

This guide lays out **six practical ways** analysts and SOC directors can evolve the SOC for the age of AI. These aren't abstract ideas or future-state ambitions. They're concrete shifts you can start making today to boost productivity, improve outcomes, and maintain control — all without replacing the human judgment at the heart of security operations.

AI won't run your SOC. But it can make your SOC faster, smarter, and more resilient — if you know how to put it to work.

Contents

- 1. Put embedded AI to work.3
- 2. Redraw the boundaries between human and AI4
- 3. Keep your SOC's AI accountable.5
- 4. Protect your organization from AI misuse6
- 5. Plan for AI's appetite.7
- 6. Prepare for agentic AI8
- What your SOC can do right now9

1. Put embedded AI to work

AI is already in your security stack — the question is whether it's doing anything useful.

Today's leading SIEM and SOAR platforms come equipped with embedded AI features designed to help reduce noise, correlate alerts, prioritize risk, and even assist with triage. But too often, those features are left underused or untouched. Why? In some cases, it's a trust issue. In others, it's a matter of unclear value, lack of time, or concerns about control.

A smart first step is to audit the AI already available in your environment. What's currently enabled? Where is it helping? Where is it getting in the way? This helps establish a baseline — not just for capability, but for comfort level.

Focus on low-risk but high-volume tasks to start. Let AI reduce alert fatigue by clustering duplicate events, helping prioritize incidents by severity, or enriching tickets with relevant context. When tuned properly, these embedded tools can reduce triage time, improve MTTD, and free analysts from repetitive workflows that drag down productivity.

Also look at where embedded AI fits into your broader threat detection, investigation, and response (TDIR) process. AI that supports **TDIR workflows** — not just siloed tasks — will deliver the clearest path to impact.

Check-in: Is your embedded AI pulling its weight?

- ☐ **Have you reviewed which AI-powered features are currently enabled in your SIEM, SOAR, or security data platform?**
You may already have alert groupings, noise reduction, or auto-triage capabilities running (or sitting idle!).
- ☐ **Are those features actually helping reduce noise or investigation time?**
Track metrics like MTTD or alert volume to evaluate impact.
- ☐ **Do you know where embedded AI is underdelivering — or even getting in the way?**
Look for frustration points: over-suppressed alerts, irrelevant correlations, or poorly prioritized incidents.
- ☐ **Are you using embedded AI on low-risk, high-volume workflows first?**
Start with alert duplication or auto-enrichment before moving to higher-stakes decision-making tasks.
- ☐ **Is embedded AI connected to your broader TDIR workflow — or just running in a silo?**
AI works best when integrated into your full detection, investigation, and response process.

PRO TIP

Start small. Pilot embedded AI on routine, low-risk tasks like ticket enrichment or alert clustering. This builds team trust and paves the way for broader adoption.

2. Redraw the boundaries between human and AI

AI isn't here to replace analysts. It's here to redefine their role.

As AI takes on more tasks in the SOC, the key challenge becomes knowing which responsibilities to delegate — and which to keep firmly in human hands. That starts with identifying where AI can assist: tasks like incident enrichment, semi-automated triage, or providing interactive guidance. When done right, these handoffs can shift the analyst's role from reactive responder to high-value decision-maker.

This isn't about "trusting the machine." Rather, it's about building workflows where analysts supervise and refine AI outputs, rather than doing everything themselves. It's the difference between manually sorting every alert and reviewing a prioritized, enriched queue. Over time, this "hand-in-hand" approach could reduce burnout, sharpen focus, and elevate your team's strategic impact.

Clarifying boundaries matters. What should AI never do on its own? Where is human oversight mandatory? Who owns the final decision, and how is that documented? These are the questions that will shape not just your playbooks, but your culture.

And remember: redefining roles is a skill in itself. As your SOC evolves, everyone's role on the team will, too. Supporting individual growth into reimagined roles is as critical as tuning the tech.

Check-in: Are you making the most of your human-AI handoffs?

- ☐ **Have you identified repeatable tasks that AI can assist with — like alert enrichment, deduplication, or low-risk triage?**
Start with low-stakes, high-volume actions.
- ☐ **Are security or AI engineers reviewing and refining AI outputs, or are they still handling everything manually?**
Look for areas where human oversight can guide, not redo, AI work.
- ☐ **Do your workflows clearly show who owns each step — AI or analyst?**
Define handoff points so accountability stays clear.
- ☐ **Are your analysts trained and supported in their evolving roles?**
As AI reshapes workflows, give your team the time and tools to grow with it.
- ☐ **Have you reviewed where AI might be either over reaching or under performing?**
Check for blind spots where the balance of trust and control needs recalibration.

PRO TIP

Start by mapping your current workflows and labeling what's human-led, AI-assisted, or ready for automation. You may be surprised where the opportunities (and friction points) really are.

3. Keep your SOC's AI accountable

AI can be powerful. It can also be very — and confidently — wrong.

Generative and embedded AI systems can produce outputs that look convincing but miss the mark. A hallucinated recommendation, a mismatched correlation, or a misprioritized threat could lead your team down the wrong path — or worse, create new risks. That's why accountability is non-negotiable.

Start with visibility. Make sure AI-driven decisions are logged, traceable, and reviewable. If a model flags an event or proposes a remediation action, your team should be able to ask: Why? Based on what? What changed?

Next, monitor the health of your AI systems. Over time, models can drift, especially if the environment changes or feedback loops are missing. Periodic reviews of false positives, false negatives, or mis-escalations help reveal when something's gone off track.

Human-in-the-loop safeguards are essential, especially for sensitive actions like auto-remediation, user lockouts, or policy changes. AI might recommend, but humans should always verify.

Establishing this discipline early doesn't just catch mistakes — it builds trust. Analysts are more likely to embrace AI when they know it's monitored and adjustable, not a black box dictating actions.

Check-in: Are you holding your AI to the right standards?

- ❑ **Are AI-driven decisions — like alert prioritization or remediation suggestions — logged and reviewable?**
Transparency is your first line of defense.
- ❑ **Do you audit AI outputs regularly for accuracy and consistency?**
Track overrides, false positives, and missed detections to identify drift or degradation.
- ❑ **Is there a feedback mechanism for analysts to flag incorrect or confusing AI behavior?**
Tighten your loop between output and improvement.
- ❑ **Do your most sensitive actions still require human approval?**
Use AI to assist — not autonomously act — on high-risk workflows.
- ❑ **Are AI models adapting to your environment, or are they stuck in a static configuration?**
Resist assumptions and update training data or rules as your SOC evolves.

PRO TIP

Build feedback loops directly into your SOC workflows. Give analysts an easy way to flag and correct AI missteps, and use that input to refine behavior over time.

4. Protect your organization from AI misuse

AI isn't just a tool — it's also a growing attack surface.

As more teams adopt generative AI tools to write code, summarize logs, or generate remediation steps, the risk of misuse grows. That risk isn't just external. Employees using AI tools without clear guardrails can accidentally leak sensitive data, introduce unvetted code, or rely on outputs that haven't been validated.

Threat actors are already experimenting with prompt injection attacks — feeding malicious instructions into seemingly benign outputs to manipulate AI behavior. Meanwhile, shadow AI tools used without security oversight can expose confidential data to third-party models with no accountability.

To protect your organization, start with usage guidelines. What tools are approved? What are the boundaries? Who needs to approve new AI integrations? Treat these tools like you would any privileged system: Log activity, restrict access, and review outputs for risk.

Also consider content scanning before anything generated by AI is stored or released. That includes scripts, configurations, documentation, or even user-facing content.

Establishing controls isn't synonymous with blocking innovation. It means giving your teams the confidence to use AI safely — without putting your environment, users, or reputation at risk.

Check-in: Are you securing how AI is used in your organization?

- ☐ **Do you have clear guidelines on which AI tools are approved for internal use?**
Avoid shadow AI by making the safe path the easy one.
- ☐ **Are employees trained on what types of data should never be shared with AI tools?**
Prevent accidental leaks by setting firm boundaries.
- ☐ **Do you monitor for prompt injection or other AI-specific attacks?**
Stay ahead of threats that manipulate how AI behaves.
- ☐ **Is AI usage logged and auditable like other privileged systems?**
Visibility isn't optional when tools touch sensitive workflows.
- ☐ **Are AI-generated outputs scanned before being stored or shared externally?**
It's probably best at least at the outset to be skeptical and verify, especially with content that reaches users or codebases.

PRO TIP

Monitor AI usage just like you would for sensitive systems. If an AI tool has access to data, users, or infrastructure, it needs the same level of visibility and control.

5. Plan for AI's appetite

AI thrives on data — a LOT of it.

Generative and predictive models often require massive amounts of input data to function well. That might include log files, asset inventories, user behavior records, or even sensitive business and customer information. But without careful planning, this appetite can quickly outpace your **governance strategy**.

The challenge isn't just storage — it's control. Where is your data going? Is AI being trained on all data or just inference? Will it be retained, shared, or reused? Too often these questions are overlooked until they create a compliance issue.

Start by mapping which AI systems in your SOC (and across your org) are accessing what types of data — and for what purpose. Then: Layer on privacy, **retention**, and compliance controls that match the data's sensitivity. If a model is accessing regulated data, make sure it's subject to the same governance as the data source itself.

Data minimization goes a long way here. Don't give AI access to everything "just in case." Give it only what it needs — and remove access when it's no longer necessary.

Finally, make sure AI workflows respect your existing controls. AI doesn't get a pass on compliance. It should inherit the same access policies, audit trails, and retention standards already in place.

Check-in: Are you managing what AI consumes — and where it flows?

- ☐ **Do you know which AI tools or features are accessing your data — and why?**
Clarify what's being used for training, inference, or storage.
- ☐ **Are your AI systems following your organization's existing data governance policies?**
They should inherit and not bypass your controls.
- ☐ **Have you applied retention limits to AI-accessed or AI-generated data?**
Limit risk by minimizing what's stored and for how long.
- ☐ **Do your AI workflows include audit trails and access logs?**
Track usage like you would for any system with access to sensitive information.
- ☐ **Are you limiting AI access to only the data it needs?**
Less is more. Keep inputs lean and risk lower.

PRO TIP

Don't let AI create shadow pipelines. Any time AI accesses, transforms, or stores data, treat it like a new data flow and secure it accordingly.

6. Prepare for agentic AI

AI that recommends is one thing. AI that acts is another.

And now for something completely different. Agentic AI is emerging fast. As a result, it's poised to shift security operations again. These systems don't just offer suggestions; they take autonomous action based on goals, policies, or environmental cues. From SOAR tools that auto-triage and remediate to AI assistants with write access to cloud environments, we're entering the age of AI that moves on its own.

The benefits are real: faster containment, scalable response, and reduced manual toil. But the consequences are equally real — especially if actions are taken without proper oversight, logging, or guardrails.

Now is the time to prepare. Start by creating an inventory of where autonomous action is already happening. Are playbook actions being triggered without review? Do any tools have privileged API access to make changes in production? Which permissions are delegated to AI agents — and are they scoped appropriately?

Just as important: define your rules of engagement. Which actions can AI take unilaterally? Which require human approval? How are decisions logged and reviewed? Create clear boundaries, then build mechanisms to enforce and audit them.

Agentic AI is no longer a future-state idea. It's already being embedded into SOC workflows. The question isn't whether to use it — it's how to use it responsibly.

Check in: Are you ready for AI that makes decisions, not just advises?

- ☐ **Do any of your tools already take autonomous actions — even small ones?**
Start by auditing existing SOAR or remediation playbooks.
- ☐ **Have you scoped what actions AI is allowed to take — and where it needs human sign-off?**
Draw a clear line between assistance and authority.
- ☐ **Are all AI-initiated actions logged and reviewable?**
If it acts, it needs to leave a trace.
- ☐ **Do your existing approval workflows account for agentic AI?**
Integrate oversight into existing escalation paths.
- ☐ **Have you defined and shared your “rules of engagement” for AI?**
Treat it like a new team member — one that needs guardrails.

PRO TIP

Write (and socialize) a clear “rules of engagement” policy for AI. Define what it can act on, what it must escalate, and how those actions are recorded.

What your SOC can do right now

AI is already reshaping security operations — but you don't need to start from scratch.

1. Activate what you already have. Audit and tune the embedded AI features in your existing stack.
2. Redraw the lines between humans and machines. Let AI handle the heavy lifting so your analysts can focus on decisions and strategy.
3. Keep the machines accountable. Monitor AI outputs, verify sensitive actions, and build human-in-the-loop controls.
4. Secure your AI use. Treat AI like any other privileged tool — with clear policies, logging and review.
5. Feed it wisely. Understand what data AI is using and make sure it aligns with your privacy and compliance standards.
6. Plan for what's next. Agentic AI is here. Define your rules of engagement before actions go unchecked.

AI doesn't replace your SOC — it upgrades it. With thoughtful implementation, your team can work faster, respond smarter, and stay in control.

Discover Splunk AI



splunk>
a **CISCO** company

Splunk, Splunk>, Data-to-Everything, and Turn Data Into Doing are trademarks or registered trademarks of Splunk Inc. in the United States and other countries. All other brand names, product names, or trademarks belong to their respective owners. © 2025 Splunk LLC. All rights reserved.

25_CMP_listicle_how-AI-is-evolving_v6

