

Seven questions to ask before you put AI in your SOC

A gut check from the framework Expel uses internally, built on 10 years of running AI and automation in production. Questions 1–3 place a task on the grid. Questions 4–7 decide what you actually deploy there.

1 Am I starting from the security outcome or the mandate?

Work backward from the security outcome you need, then pick the technology that serves it. Reactive deployment produces AI sprawl: lots of capability, no strategy, and analysts more overwhelmed than before.

2 What breaks if it gets this wrong?

This is your impact axis. At the low end, failure means noise and a wasted hour. At the high end it means an isolated domain controller or a wiped C-suite laptop. Rank tasks by consequence before you rank them by effort.

3 Would I hand this task to a new analyst on day one?

You already know how to scope trust. A junior hire doesn't get a complex forensics engagement, and your principal responder doesn't run tier 1 triage. Use the same judgment with AI, and expand its autonomy the same way, once it's earned.

4 Does this need AI, or would plain automation do?

Deterministic workflows handle a lot of repeatable work without the model costs. Some of the automation behind Ruxie, Expel's AI SOC manager, has run for 10 years with no AI model in the loop. Save the models for problems that need reasoning or pattern recognition.

5 Is the data underneath this clean enough?

Feed a model flawed telemetry and it will clear the queue with the wrong calls in it. When a low-consequence task keeps eating analyst hours, the problem is the process or the data source. Fix that first.

6 Who's accountable when the AI is wrong?

AI can't testify, take responsibility, or explain a bad call to your board. A person can. For high-impact work, keep a human as the approver: reviewing the output, verifying it, authorizing the action.

7 What would this have to prove to earn more autonomy?

Set benchmarks the way you'd set them for a promotion. Write down the accuracy and tunability you need to see before you expand what AI handles. It gives your team clarity, and it gives vendors a specific answer when they ask why you haven't turned their AI features on.

Where each task lands



Know what you're deploying



Automation

Deterministic rules. If X, do Y. No AI required.



Machine learning

Pattern recognition that improves with data volume.



Generative AI

Reasoning and synthesis, with probabilistic output rather than exact results.



Agentic AI

Autonomous planning and execution of multi-step work. Different infrastructure, different oversight, different cost of being wrong.