AI decision models arrive in SOC tooling: what security leaders should weigh



A new class of AI "decision" models has appeared within weeks, with several vendors releasing products. The trend bears on how AI SOC platforms are built and assessed.

- TypeSafe AI released Jev on September 15 and described it as a System One model, referring to fast, intuitive thinking.
- OpenAI announced a Decisions API at DevDay and opened a public beta on October 6.
- Cribl published cribl-decision, an open LoRA adapter that can be run locally.
- The source says the category went from nonexistent to widespread within about a month.
- The excerpt is partial; it does not detail performance, pricing or security assessments.

What happened

According to The Hacker News, a new category of AI model focused on decisions went from nonexistent to widely discussed in roughly a month. TypeSafe AI released Jev on September 15 and positioned it as a System One model, a reference to Kahneman's idea of fast, intuitive thinking. OpenAI announced a Decisions API at DevDay and opened a public beta on October 6. Cribl published cribl-decision, an open LoRA adapter that organisations can run locally. The source notes that other vendors are entering the space as well.

Why it matters for SOC leaders

The headline links these decision models to the design of AI SOC platforms. The available excerpt does not describe benchmarks, accuracy, costs or specific security findings, so claims about capability should be treated as unverified until tested. What is clear is that options now range from hosted APIs to models that can run locally, which affects data handling, third-party risk and control over how automated triage decisions are made.

Practical stance

There is no need for urgency. Organisations considering AI-assisted triage can use this moment to set evaluation criteria before vendors approach them, rather than adopting a tool because of market momentum. Hosted and local deployment models raise different questions about where security telemetry is processed and who can review the decisions produced.

What to do now

1. Define evaluation criteria for AI-assisted triage, including accuracy testing on your own historical alerts, before any vendor engagement.
2. Compare hosted API and locally run model options for where security telemetry is processed and stored.
3. Check whether using external AI services with SOC data is consistent with PDPL obligations and your data classification rules.
4. Keep a human review step for automated decisions affecting containment or escalation until performance is proven.
5. Include AI model providers in third-party risk assessments.