LuceraAI, in-house SOC, MSSP or an AI SOC platform?

A plain comparison for security and IT leaders deciding how every alert gets investigated, who makes the calls, and what you can check afterwards.

The short answer

Every alert investigated, with evidence you can check.

For a lean security team that wants every alert investigated, with evidence it can check and limits it controls, LuceraAI gives that coverage without building a 24/7 SOC or handing judgement to an outside provider.

Each option suits a different team. An in-house SOC gives the most context and control at the highest cost. An MSSP gives round-the-clock people but less visibility into each decision. An AI SOC platform gives speed and automation, but often asks you to build and trust the automation yourself.

The four options in brief

AI SOC AGENTS, HUMAN-SUPERVISED

LuceraAI

A council of specialist AI agents investigates every alert from your EDR. Two independent models must agree on each verdict, a devil's advocate argues against it, and every verdict comes with its evidence. It acts only within limits you set, and anything uncertain goes to a person. With our managed service, LuceraAI analysts pick up escalations 24/7.

YOUR OWN ANALYSTS

In-house SOC

Your own analysts, tools and processes. Full context and control, but 24/7 cover means hiring, training and keeping a shift team, and alerts still queue when people are stretched.

A PROVIDER'S ANALYSTS

MSSP

A provider watches your alerts round the clock, usually across many customers at once. Quality and visibility vary widely by provider and contract.

AUTOMATION YOU BUILD

AI SOC platform

A security automation platform with case management and AI agents for investigation and response. Built for enterprise SOC teams, who design its automations and choose how much it does without a person.

✦ Side by side

The difference is the combination.

Every alert investigated, a checkable receipt for each verdict, and a person on every uncertain call.

	LuceraAI	In-house SOC	MSSP	AI SOC platform
Alert coverage	Every alert investigated, at 2pm or 2am	Limited by headcount; out of hours needs shifts	24/7; depth varies by tier and contract	High once automations are built and tuned
How decisions are made	A council of specialist agents; two models from different providers must agree	Your analysts' judgement	The provider's analysts and processes	Automations and AI agents you configure
How decisions are explained	A receipt per verdict: evidence, reasoning, confidence	Analyst notes; varies by person	Ticket summaries; full reasoning often not shared	AI-generated case summaries
People in the loop	Stops for a person when unsure, models disagree or evidence is missing	Always your people	The provider's analysts, not yours	Configurable, up to fully orchestrated
Control over actions	Per site: only actions you switch on, above thresholds you set	Full	Set by the contract; the provider acts	You build the runbooks and guardrails
Getting started	Read-only EDR connection; two-week evaluation in Observe mode	Hiring, tooling and training	Onboarding and contract	Build or adapt automations and integrations
Data	Hosted in AWS London; AI model processing in Europe; not used to train AI models	Stays with you	Varies; check location and access	Cloud platform; check region and AI data use
Cost model	Based on the number of devices you protect	Salaries for 24/7 shifts plus tools	Subscription, often per endpoint or tier	Platform licence plus engineering time
Best suited to	Lean security teams, and IT teams with no 24/7 SOC	Large organisations with budget and specific needs	Teams wanting monitoring fully outsourced	Enterprise SOCs with automation engineers

The AI SOC platform column is based on leading vendors' public product material. The in-house SOC and MSSP columns describe typical models; individual providers differ.

When each option is right

Pick the one that fits your team.

Choose LuceraAI

Your EDR raises more alerts than your team can investigate, you can't staff a 24/7 SOC, and you want every alert investigated with evidence you can check before anything acts on your estate.

Choose an in-house SOC

You have the budget and hiring pipeline for round-the-clock shifts, and need analysts who know your business intimately. LuceraAI can sit alongside them and gather the evidence first.

Choose an MSSP

You want monitoring and response fully outsourced and are comfortable with the provider making calls for you. Ask to see how each closed alert was decided.

Choose an AI SOC platform

You have an enterprise SOC with automation engineers, many tools to orchestrate, and time to build and tune the workflows yourself.

Questions to ask any option

1. Will every alert get a full investigation, including out of hours?
2. Can I see the evidence and reasoning behind each closed alert?
3. What happens when the system or the analyst is unsure?
4. Which actions can happen on my estate without my approval?
5. Who builds and maintains the playbooks or automations?
6. Where is my data held, and is it used to train AI models?
7. Can I test it on my own alerts, read-only, before I commit?
8. What can it not see or do? Ask for the limits in writing.

How to test LuceraAI

The fairest comparison is your own alerts.

1

Connect read-only

Connect SentinelOne or Microsoft Defender XDR with read-only access.

2

Two weeks in Observe

LuceraAI investigates your real alerts. Nothing acts on your estate.

3

Review the results

Against your baseline, with method notes and a plain list of anything we could not measure.

Your alerts. Our working.

Pricing is based on the number of devices you protect.

Run an evaluation: luceraai.net/evaluate

signals@luceraai.net
+44 (0)1483 948090
luceraai.net/demo