GTK CYBER

EXECUTIVE RESOURCE

AI Security Vendor Evaluation Checklist

Every security vendor has an AI story now. Some are real, many are not. Use this before you buy, in a demo, or in a POC to separate real capability from marketing.

1. Pin down the claim

"AI-powered" is not a claim. Make the vendor be specific.

- [ ] What problem does the AI solve, specifically?
- [ ] What does the AI do that rules, signatures, or heuristics cannot?
- [ ] Where does the AI sit in the detection or response workflow?

2. Interrogate the training data

- [ ] What data was the model trained on, and how recent is it?
- [ ] Was it trained on your industry's data or general data?
- [ ] How often is the model retrained?
- [ ] What happens on inputs outside the training distribution?

3. Understand the false-positive rate

- [ ] False-positive rates in customer environments similar to yours
- [ ] How alert volume changed after deployment
- [ ] What tuning is required, and who does it
- [ ] A vendor claiming near-zero false positives has not deployed at scale or is cherry-picking

4. Prove it on your data

- [ ] A POC on your data (or realistic synthetic data matching your environment)
- [ ] Success criteria defined in advance
- [ ] Access to raw detection output, not just a dashboard
- [ ] If the vendor will not run a POC on your data, that is the answer

5. Demand explainability

- [ ] Can the model explain why it flagged a specific alert?
- [ ] What features or evidence drove the detection?
- [ ] Can analysts access the underlying evidence, not just the verdict?

6. Do not buy AI to buy AI

- [ ] The specific problem you are solving
- [ ] What you do now and why it is insufficient
- [ ] What success looks like in measurable terms
- [ ] What the non-AI alternative would cost

Adapted from GTK Cyber practitioner analysis. Full article: gtkcyber.com/blog/evaluating-ai-security-vendors