

Before You Buy Another AI Security Tool, Map What You Already Have

By Junior Williams

Draft for review · 14 October 2026

Before approving another AI security product, I would ask the team to walk through one AI workflow from beginning to end. What information goes in? Where does it travel? What can the system do with it? Who is responsible when something goes wrong?

Those questions give a product evaluation somewhere useful to start. A feature list tells you what a tool is designed to do. A workflow map helps you decide whether that capability addresses an actual gap in your environment.

Think about adding a security camera to a building. The decision depends on which entrance needs coverage, who watches the footage, and what happens when something requires attention. Buying a capable camera leaves those operating decisions unresolved.

The same reasoning applies to AI.

Follow one piece of work

Consider a hypothetical internal assistant that helps employees answer questions about company policies. It searches a document library, drafts an answer, and links to supporting material. An employee reviews the answer before using it.

The description sounds straightforward. A useful assessment would follow the information through each step:

- An employee signs in and asks a question.
- The assistant searches documents using a configured identity and permissions.
- Selected passages are sent to a model to produce an answer.
- The answer and its sources appear in the employee's session.
- Logs and other records may retain information about the interaction.

Each step raises a different control question. Does document retrieval respect the employee's access? Which information can leave the organization's environment? Where is the interaction retained? Can the employee tell when the answer lacks adequate support?

The map should also show any actions the assistant can take. An assistant that can update a record or send a message needs decisions about authority, approval and recovery in addition to decisions about information access.

Separate coverage from evidence

Existing identity, data protection, application and monitoring systems may address parts of this workflow. Their usefulness depends on how they are configured and connected to the actual process.

For each important control, record three things: the intended protection, the component responsible for providing it, and the evidence that it works in the agreed scope.

For example, “employees only receive documents they are authorized to access” is a requirement. Naming the identity platform describes part of the implementation. A permitted-access review and a controlled check using representative roles provide evidence about the result.

Keep the wording precise. “Configured” means someone has set it up. “Reviewed” means someone examined it. “Tested” should identify the behaviour checked, the conditions and the result. Those distinctions make the assessment easier to act on.

Use a small coverage worksheet

One worksheet per workflow is usually a workable starting point:

Record	Question to answer
Purpose and owner	What work does this support, and who accepts responsibility for its operation?
Information and access	Which sources, data categories and identities does it use?
Permitted actions	What may it change, and where is human approval required?
Existing controls	Which current capabilities address the identified risks?
Evidence and gaps	What has been checked, what remains uncertain, and what is missing?
Response and recovery	Who can stop the workflow, investigate a problem and restore service?

Keep the first version bounded. A named workflow with a clear owner and a short list of unresolved questions is enough to begin a useful discussion.

Let the gap determine the next step

The next step might be a configuration change, a permissions review, better operating instructions, or additional testing. It might also be a new product. The assessment should explain why the proposed action fits the gap and how the team will judge whether it helped.

If additional tooling is justified, evaluate the integration and ongoing work alongside the advertised protection. Someone will need to manage it, review its findings and respond when it raises a concern. Include that effort in the decision.

I would want the final assessment to give leadership a clear decision and give the delivery team something specific to implement. Both should come from the same evidence.

JWC's AI Exposure and Control Assessment is designed around that approach. To discuss a bounded review of one or more AI workflows, [book an introductory call](#).

Further reading

For a broader risk-management structure, the [NIST AI RMF Playbook](#) addresses AI inventories (Govern 1.6) and documented controls across system components (Map 4.2). The worksheet above is my suggested starting point, not a NIST certification checklist.