SANCTUMSHIELD · AGENTIC GOVERNANCE FOR THE OTHER 99%

AI Attack Path Tabletop

One path, four steps. At each step: where would it be stopped here, and who owns that control?

A one-page exercise for the people who own your code, your cloud, your secrets and your AI services. Allow 45 minutes. Walk the path below in order. At each step, write down where your environment would stop it, the one person who owns that control, and the date it was last checked. Most teams find that some steps are already well covered. The exercise puts those controls on the record and shows where the next one belongs.

Educational template, not legal advice.

Last updated: November 9, 2026

Step	What the attacker has now	The question for your team	Where it would be stopped here	Control owner (one name)	Last checked
1. A hard-coded token in a code repository	A key or token written into code that someone can read	Could a key or token sit in code here? What would catch it before it is pushed, and who would replace it?			
2. The secrets manager	That token, used to sign in to the secrets manager	What can that token read there? Is it limited to one job, and would anyone notice an unusual read?			
3. Administrative credentials	Administrator keys or logins read from the secrets manager	Who would notice them used from a new place? Who may revoke them without waiting for approval?			
4. Your own AI services	Administrator access to your AI services and the data they can see	Would a spending alert or usage log show someone else using them? Who may cut off the account?			

Credit what already works. In one publicly reported 2026 case, branch protection on infrastructure code stopped part of an AI-assisted attack. A control you already run counts. Write it in the row where it stops the path.

The four actions, decided before you need them

Action	Who may do it without approval (person or role)	Backup	Last practiced (date and result)
Revoke keys and logins			
End signed-in sessions			
Pause deployments			
Cut off a cloud account			

After the walk-through. Circle the first step where the path stops today. Any row without an owner gets one named this week. Any action without a practice date gets a short exercise on the calendar. Date this page, and walk the path again when your AI services, keys or cloud accounts change.

The idea for this exercise comes from a November 2026 board newsletter on frontier AI.

Terms: AI credential register, LLMjacking, Containment authority (pre-decided), Agent borrowing human identity, Proven Stop. The Complete assessment records the credential register and the four actions with their owners and dates, in the Executive Risk Report and the Board Memo.

Free, and there is nothing to sign up for. The Board AI Register, the Responsibility Register and the free Shadow AI Check are at sanctumshield.com/templates and sanctumshield.com/check.

© PIGENAI LLC · SanctumShield · sanctumshield.com · Last updated November 9, 2026 · This is an educational template, not legal advice.