SANCTUMSHIELD · AGENTIC GOVERNANCE FOR THE OTHER 99%

Board AI Register + Eight Questions

Fill this in once and you have a record your board can read out loud. Keep it current and you have a program.

Educational template, not legal advice.

Last updated: October 27, 2026

Page 1 — AI system register

One row per AI system your organization uses, whether it was bought, built, or arrived inside software you already owned.

System	Owner (name)	Allowed to	Not allowed to	Who can override	Stop / escalate trigger	Last proven stop (date, by whom)	Exposure triad (Y/N, approved Y/N)	S l a

Gateway row (added September 30, 2026)

One row type for an : a service that exposes many of your business systems to AI agents through one connection. Record the gateway as its own entry and list the systems it reaches as children beneath it.

Gateway	Type (A vendor-operated / B self-hosted)	Owner (name)	Scope record location	Identity source	Connected systems (children)	Review date

A gateway is one control point that reaches every system behind it, so it is the single most important entry in the register to keep current. For a vendor-operated gateway (Type A) the vendor holds the credentials and carries the data, so it is also a processor under GDPR Article 28 and belongs on your processor list. For a self-hosted gateway (Type B) the questions are who operates it, what its policy says, and whether its logs reach a person.

Owner is one human name, never a team. *Allowed to* and *Not allowed to* are written in the verbs that matter to a board: recommend, decide, act, spend, speak for us. *Who can override* names the person and the mechanism. *Stop / escalate trigger* is the condition that hands the work back to a person. *Last proven stop (date, by whom)* is the date a named person last actually halted or paused the system mid-task, in a test or a real event, and who did it. A stop nobody has exercised is a hope, not a control. *Exposure triad (Y/N, approved Y/N)* records whether the system can see confidential data, read content from outside the organization and send information outward (Y when all three are true), and, if so, whether a named person has approved that combination with a date (for example "Y, approved Y"). *Spend limit and approver* is the ceiling on what the system may spend, including any usage allowed past its quota, and the name of the person who approved it. *Last reviewed* is a date you enter, because a register that is never re-read is a historical document.

(in *Allowed to* and *Not allowed to*): whether it may start subagents and where agent-to-agent exchanges are logged, the systems it may write to, any schedule it runs on, whether it may keep running past its quota, and who approves the changes it proposes. If it answers from company documents for consequential decisions, *Evidence location* points to the document, version and page. Terms: Proven Stop, MCP Gateway, AI system register.

Page 2 — The eight questions

Eight questions a board can ask about any AI system the company runs. Each one has an answer that lives somewhere specific.

#	The board's question	Answered where	By whom	Date
1	What is this system allowed to do: recommend, decide, act, spend, speak for us?			
2	What is it not allowed to do, and who signed that list?			
3	Who can override it, and how fast?			
4	What evidence has to sit behind an output before anyone acts on it?			
5	If it goes wrong, can we reconstruct what happened?			
6	Who tests it, independently of the people who built or bought it?			
7	What makes it stop, or hand back to a person?			
8	Who is accountable for the outcome, by name?			

Reporting line (added October 8, 2026)

When an agentic system fails, damage moves faster than an escalation process. The reporting line is part of the control, not an administrative detail.

Question	Answer	By whom	Date
A standing AI and cyber risk committee exists (Y/N)			
The CISO has direct access to it (Y/N)			
The CISO reports to (name and title)			

Where these answers usually live

If your organization already produces governance artifacts, most of these questions already have a home.

- Questions 1, 2, 3, 4 and 7 are answered by an : authority and scope, the one-page not-allowed list with a named signer, override and escalation, evidence and human-review requirements, and stop triggers.
- Questions 5 and 6 are answered by an : record-keeping findings and the assessment method, refreshed on a cadence.
- Question 8 is answered by a , which names the accountable owner in the form a director reads.

Four artifacts. Eight questions your board can answer out loud. Refreshed on a monthly cadence, because the next question is always "and is that still true?"

Why a register, specifically

A regulator, an underwriter, or a director will ask for the register before asking for anything else. A register with named owners against each entry is the fastest way to answer that first question well.

EU AI Act Article 4 has required providers and deployers to take measures supporting the development of AI literacy among the people operating AI on their behalf since February 2, 2025. Article 50's transparency obligations have applied since August 2, 2026. A register is how an organization shows which systems those duties attach to.

Governance is an ongoing program, never a one-time audit. Date this page when you fill it in, and again when you review it.

The Shadow AI Check is twelve questions and no account; the method is published in full. The Academy teaches this register in module M20 (From Blueprint to Ledger), in every track.

© PIGENAI LLC · SanctumShield · sanctumshield.com · Last updated October 27, 2026 · This is an educational template, not legal advice.