

The board meeting ends on time. Seventeen items, all green or yellow, and everyone walks out relieved to have "done the risk thing." Two weeks later a model update in the customer-success agent starts misclassifying tickets, the errors feed three other agents, and churn is climbing before anyone notices. None of the seventeen items mentioned it. The post-mortem ends the way every serious AI failure does: "This risk was not on the register."

The reflex is to add more items and schedule another review. That scales the document, not the protection. A longer list is a snapshot of what you already knew to name, not a map of what breaks you.

The AI Risk Register introduces the RISK framework: Register what matters, Inspect with evidence, Scenario-plan the unlikely, Kill-switch before deployment. Four moves turn a static compliance artifact into a living decision system that changes what your team ships. RISK is an operating heuristic, not a validated instrument: a lens you test against your register.

You will learn to:

- Register what matters: give every material risk one owner, a first evidence date, and an update trigger.
- Inspect with evidence, replacing "we talked about it" with a measured condition and a last-inspected date.
- Scenario-plan the unlikely with cards that trace the second-order path and name the tail condition before it cascades.
- Build tested kill-switches before deployment, so you stop a system on purpose.
- Wire the register into how the team sets risk appetite, escalates, and reports.

The result is a one-page living register with named owners, tested kill-switches, and scenario cards that trace what a static list never sees. It turns vague dread about AI into proportionate capability.


If you run material AI in production and refuse to meet your exposure in the post-mortem, start with the register you have and the effect it never named.

Len P. van der Hof, MSc

He builds operating systems for the strategic self. Treats entrepreneurship, AI and machine learning, marketing, philosophy, psychology, and health optimization as one engineering problem. MSc in Strategic Entrepreneurship, Rotterdam School of Management, Erasmus University.

The AI Risk Register

Build a **Living** Register That Names, **Owns**, and **Tests** Every AI Risk Before It Names You



Len P. van der Hof

Systems for the Strategic Self