

Agentic AI engineering

Give AI the ability to act, with limits that hold.



BEST FIT Before an agent can act	ANCHOR LENGTH 2–3 weeks to assess	WAY IN Agent risk assessment	STARTS WITH Shadow Scanner
--	---	--	--------------------------------------

01 THE PROBLEM

Agent demos look alike. Operations do not. An agent that can spend, write or send inherits whatever credentials it was handed, and most teams have never tested how to stop one mid-run.

02 WHERE IT LIVES

AI STACK	DELIVERY STAGE
Applications	Design & assess
Delivery & lifecycle	Implement
Models	Enable
Orchestration & runtime	Monitor
Data platform	Iterate
Infrastructure	

03 METHODOLOGY

Scope → Limit → Gate → Log → Drill


- 1 Scope**
What each agent may touch, and must never touch, defined before it is built.
- 2 Limit & gate**
Spend and rate caps at the platform; a person approves anything irreversible.
- 3 Log**
Every action traced to an agent, a trigger and an authority, across hand-offs.
- 4 Drill**
Stop an agent mid-run in a rehearsal and time how long it takes.

04 HOW WE DO IT · THE 808 REPORT

808 Blast-Radius Chart

Each agent at the centre, with rings for the tools, systems and data it can reach. Red marks reach it has but does not need. That is the first thing to cut.

SAMPLE · FICTIONAL DEMO COMPANY · NOT
CLIENT DATA



05 THE ENGAGEMENT

5.1 ANCHOR ENGAGEMENT Agent Risk Assessment Fixed-fee engagement · 2–3 weeks PLANNED	5.2 ANCHOR ENGAGEMENT Guardrail Reference Build Productised build PLANNED	5.3 CONTINUING Agent Incident Drill Workshop · half day PLANNED
--	---	---

06 THE OUTCOME

- Agents that work inside boundaries you set
- A record of what every agent did, and for whom
- A kill path your team has actually rehearsed

MEASURED BY

Minutes to stop an agent mid-run, in a drill

THE TRADE-OFF

Approval gates slow the workflow at exactly the steps that matter. We place them only where a mistake is expensive or permanent.