

AI strategy & roadmapping

Know where AI actually pays before you buy anything.



BEST FIT Before you buy anything	ANCHOR LENGTH 6 weeks	WAY IN Free Readiness Scorecard	STARTS WITH Shadow Scanner
--	---------------------------------	---	--------------------------------------

01 THE PROBLEM

Most AI plans start from a vendor demo or a board question, and are built on the approved-tools list rather than what is actually running. Pilots then stall on data, risk, cost or unclear value.

02 WHERE IT LIVES

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

03 METHODOLOGY

Inventory → Score → Sequence


- 1 Inventory**
Shadow Scanner and interviews show what is really in use, by which teams.
- 2 Discover**
Two workshops turn scan findings and business pain into a long list.
- 3 Test**
Each shortlisted idea gets an owner, an access check, a data sample and a rights check.
- 4 Score & sequence**
Value, feasibility and data readiness scored; risk as a pass/fail gate; decisions dated.

04 HOW WE DO IT · THE 808 REPORT

808 Opportunity Grid

Every candidate plotted by value against readiness.
The quadrant decides the call: go now, prepare, not yet or drop. Dot size is the cost to run at real volume.

SAMPLE · FICTIONAL DEMO COMPANY · NOT
CLIENT DATA



05 THE ENGAGEMENT

1.1 ENTRY AI Readiness Scorecard Free assessment AVAILABLE	1.2 ANCHOR ENGAGEMENT AI Opportunity Map Fixed-fee engagement · 6 weeks · 5 workshops PLANNED	1.3 ANCHOR ENGAGEMENT Board Briefing Pack Fixed-fee add-on PLANNED	1.4 CONTINUING Quarterly Roadmap Review Retainer PLANNED
--	--	--	--

06 THE OUTCOME

- A sequenced roadmap of decision points across 6–18 months
- A scored register of every use case, with named owners
- Build, buy or leave positions, with vendor viability checked

MEASURED BY

Posture score at kickoff, re-measured each quarter

THE TRADE-OFF

Starting from an inventory takes longer than starting from a shortlist. Some of your ideas will come back as "not yet".

Vendor & tool evaluation

Test vendors against your workloads, not their demo.



BEST FIT Choosing between providers	ANCHOR LENGTH 4–6 weeks	WAY IN Published shortlist teardown	STARTS WITH Shadow Scanner
---	-----------------------------------	---	--------------------------------------

01 THE PROBLEM

A demo runs clean data on a happy path. Your environment does neither. And picking a vendor is partly a bet it still exists at renewal: one software category retired 1,367 products in a year.

02 WHERE IT LIVES

AI STACK

Applications
Delivery & lifecycle
Models
Orchestration & runtime
Data platform
Infrastructure

DELIVERY STAGE

Design & assess
Implement
Enable
Monitor
Iterate

03 METHODOLOGY

Criteria first → Bake-off → Cost → Viability


- 1 Criteria first**
Scoring written from your workloads and weighted before any vendor calls.
- 2 Bake-off**
The shortlist runs the same tasks, data and failure cases, side by side.
- 3 Three-year cost**
Licensing, inference, egress, integration and the staff to run it.
- 4 Viability & exit**
Funding, concentration and acquisition risk, plus what it takes to leave.

04 HOW WE DO IT · THE 808 REPORT

808 Fit & Viability Matrix

Each vendor scored on fit to your workloads and on viability at renewal. Bar length is three-year cost. The runner-up stays on the page, with where it would have won.

SAMPLE · FICTIONAL DEMO COMPANY · NOT CLIENT DATA



05 THE ENGAGEMENT

2.1 ENTRY Vendor Shortlist Teardown Free published comparison PLANNED	2.2 ANCHOR ENGAGEMENT Bake-Off Report Fixed-fee engagement · 4–6 weeks PLANNED	2.3 ANCHOR ENGAGEMENT Three-Year TCO Model Deliverable artifact · spreadsheet PLANNED	2.4 CONTINUING Renewal Watch Annual subscription PLANNED
--	---	--	---

06 THE OUTCOME

- A scored comparison you can defend in procurement
- A three-year cost model you can re-run yourself
- A recommendation that states its trade-offs

MEASURED BY

Cost and fit against the model at first renewal

THE TRADE-OFF

A real bake-off takes weeks and your team's time. A demo shortlist takes an afternoon. The difference shows at renewal.

Governance, risk & compliance

Controls that survive an audit, not just a checklist.



BEST FIT An audit or a board question	ANCHOR LENGTH 4–6 weeks	WAY IN Governance gap check	STARTS WITH Shadow Scanner
--	----------------------------	--------------------------------	-------------------------------

01 THE PROBLEM

Policy written to pass a questionnaire fails the moment someone asks for evidence. And a control built for a person clicking a button does not cover an agent running unattended at 3 a.m.

02 WHERE IT LIVES

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

03 METHODOLOGY

Inventory → Map → Control → Evidence

- 1 Inventory**
Every AI system listed and classified by risk, including the ones nobody approved.
- 2 Map obligations**
EU AI Act exposure, state-level US rules and SOC 2 mapped to what applies to you.
- 3 Build controls**
Approval and spend limits enforced in the system, each one testable.
- 4 Evidence & rehearse**
Logs that answer "which system, on whose authority", and a rehearsed response.

04 HOW WE DO IT · THE 808 REPORT

808 Control Heatmap

Obligations down the side, systems across the top. Each cell shows whether the control is tested, partial or missing, so the gaps an auditor will find are visible first.

SAMPLE · FICTIONAL DEMO COMPANY · NOT CLIENT DATA

	CRM	SUPPORT	DEV	HR	AGENTS
Inventory	TESTED	TESTED	PARTIAL	TESTED	MISSING
Testable controls	TESTED	PARTIAL	PARTIAL	PARTIAL	MISSING
Spend limits	PARTIAL	TESTED	MISSING	PARTIAL	MISSING
Evidence trail	TESTED	PARTIAL	PARTIAL	MISSING	MISSING
Use policy	TESTED	TESTED	TESTED	PARTIAL	PARTIAL

TESTED PARTIAL MISSING

05 THE ENGAGEMENT

3.1 ENTRY AI Governance Gap Check Free diagnostic PLANNED	3.2 ANCHOR ENGAGEMENT Control Set & Evidence Pack Fixed-fee engagement · 4–6 weeks PLANNED	3.3 ANCHOR ENGAGEMENT AI Use Policy Template Productised template + half-day tailoring PLANNED	3.4 CONTINUING Audit Standby Retainer PLANNED
---	--	--	---

06 THE OUTCOME

- A control set that holds when someone tests it
- Evidence you can produce without reconstruction
- An acceptable-use position staff have actually seen

MEASURED BY

Share of obligations with a tested control

THE TRADE-OFF

Enforced limits add friction. Some requests that went through on a credit card will now wait for approval. That is the control working.



Implementation & integration

From pilot to production, without a rebuild each time you switch.

BEST FIT Pilot to production	ANCHOR LENGTH 4–8 weeks	WAY IN Integration architecture review	STARTS WITH Shadow Scanner
--	-----------------------------------	--	--------------------------------------

01 THE PROBLEM

Pilots stall at the integration boundary: identity, data and monitoring turn out to be the real project. Build against one provider's conventions and switching later means a rewrite.

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

Boundary → Wire → Observe → Hand over


- 1 Boundary layer**
Provider code sits behind an interface you own, so switching is a change, not a rebuild.
- 2 Wire it in**
Identity, data paths and least-privilege access from the first commit.
- 3 Observe**
Cost, latency and failure on one view before go-live, not after the first invoice.
- 4 Hand over**
Your team runs it unaided in a working session before sign-off.

04 HOW WE DO IT · THE 808 REPORT

808 Boundary Map

Your systems on the left, providers on the right, and the boundary layer you own in between. Every crossing is labelled with its identity, data path and monitor.

SAMPLE · FICTIONAL DEMO COMPANY · NOT CLIENT DATA



05 THE ENGAGEMENT

4.1 ENTRY Integration Architecture Review Fixed-fee engagement · 1–2 weeks PLANNED	4.2 ANCHOR ENGAGEMENT Pilot-to-Production Sprint Fixed-fee engagement · 4–8 weeks PLANNED	4.3 ANCHOR ENGAGEMENT Runbook & Handover Pack Deliverable artifact + working session PLANNED
--	---	--

06 THE OUTCOME

- ☐ One workflow in production, done properly
- ☐ Runbooks written for the team who will operate it
- ☐ An exit path: switch providers without a rewrite

MEASURED BY


Time to switch provider on the boundary, tested

THE TRADE-OFF

One workflow done properly ships later than three pilots left at eighty percent. It is also the one that reaches production.

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.

Team enablement & training

Build the judgement in-house.



BEST FIT Less reliance on outside advice	ANCHOR LENGTH Program, scoped per team	WAY IN Lunch & learn	STARTS WITH Shadow Scanner
--	--	------------------------------------	--------------------------------------

01 THE PROBLEM

Training usually teaches a product's interface, not the judgement behind it, so it expires with the next release. When every decision needs an outside opinion, capability is rented, not built.

02 WHERE IT LIVES

AI STACK

Applications
Delivery & lifecycle
Models
Orchestration & runtime
Data platform
Infrastructure

DELIVERY STAGE

Design & assess
Implement
Enable
Monitor
Iterate

03 METHODOLOGY

Observe → Tailor → Practise → Codify


- 1 Observe**
Start from what teams already use, including the tools nobody approved.
- 2 Tailor**
Separate tracks for the people building, buying and accountable.
- 3 Practise**
Hands-on in your own systems and data, not a generic sandbox.
- 4 Codify & re-measure**
Playbooks that stay when people leave; a re-scan shows what changed.

04 HOW WE DO IT · THE 808 REPORT

808 Adoption Shift

Per team, the share of AI use that runs on sanctioned tools before and after the program. It turns training into a number rather than a feeling.

SAMPLE · FICTIONAL DEMO COMPANY · NOT CLIENT DATA



05 THE ENGAGEMENT

6.1 ENTRY Lunch & Learn Free or low-cost workshop · 60–90 min PLANNED	6.2 ANCHOR ENGAGEMENT Role-Based Training Program Fixed-fee program PLANNED	6.3 ANCHOR ENGAGEMENT Decision Playbook Library Productised template pack PLANNED
---	---	---

06 THE OUTCOME

- A team that runs the next evaluation without us
- Decision playbooks you keep
- A measured change in how tools are actually used

MEASURED BY

Sanctioned share of AI use per team, re-scanned

THE TRADE-OFF

Done well, this reduces how much you need us. That is the point of it.