



AI Opportunity Map

Six weeks from “what could AI do for us?” to a sequenced roadmap your team can act on and your board can follow — built from what is actually running, not from the licence list.

FORMAT	LENGTH	WORKSHOPS	STARTS WITH
Fixed fee	6 weeks	5 sessions	Shadow Scanner

This brief explains what the engagement includes, how each step works, what we need from you, and the published evidence behind the method. Every figure is sourced on the last page.

Who it is for

A CIO, CTO, head of architecture or business lead who has been asked for an AI plan — and wants one that survives contact with the data, the budget cycle and the auditors.

The problem it solves

Most AI plans start from a vendor demo or a board question. They skip the inventory, score ideas on how well they demo, and stall between pilot and production. Gartner expected at least 30% of generative AI projects to be abandoned after proof of concept by the end of 2025, citing poor data quality, weak risk controls, rising costs and unclear business value.¹

What it does instead

It starts from your live footprint, tests each idea against those four failure points before it enters the plan, and sequences the survivors into decisions with the conditions that must hold first. Some of your ideas will be marked “not yet”. That is the plan working.

What you walk away with

- 01 **Live footprint report.** What AI and SaaS is actually in use, by which teams, with a posture score you can re-measure.
- 02 **Scored use-case register.** Every candidate with value, feasibility, data readiness, a risk check and a named owner. A spreadsheet you keep.
- 03 **Sequenced roadmap.** Decision points across 6–18 months, each with the conditions that unlock it.
- 04 **Build, buy or leave.** A written position per use case, with vendor viability checked.
- 05 **Readout.** A short deck and a one-page summary for your executive team.

One sponsor, about 10 hours of their time. Around 8 hours each from 6–12 workshop participants.

It is not

An implementation, a vendor bake-off or a board pack. Those are separate engagements (4.2, 2.2 and 1.3), and the roadmap tells you whether you need them.

The trade-off

Starting from an inventory takes longer than starting from a shortlist. You get fewer, better-supported use cases rather than a long list of possibilities.

No referral fees

The 808 does not resell software or take referral fees from AI, SaaS or security vendors. “Buy nothing yet” is an allowed answer.

WHY THIS METHOD

EACH STEP ANSWERS A KNOWN FAILURE

We did not invent the failure modes. They are documented in public research. The engagement is built so that each one is checked before a use case reaches your roadmap.

78%

of AI users bring their own tools to work

Microsoft and LinkedIn's 2024 Work Trend Index (31,000 people, 31 countries) found 75% of knowledge workers use AI at work, and 78% of those users bring their own tools.² A plan built on the approved-tools list misses most of what is happening.

So we start with an inventory. Shadow Scanner shows what is actually running and who uses it, in week 1.

63%

lack or are unsure of AI-ready data practices

Gartner surveyed 1,203 data management leaders in July 2024. 63% did not have, or were unsure they had, the right data management practices for AI. Gartner predicts organisations will abandon 60% of AI projects unsupported by AI-ready data through 2026.³

So we test data before scoring. In week 3 every shortlisted use case gets an owner, an access check, a quality sample and a usage-rights check.

30%

of GenAI projects expected to stall after proof of concept

Gartner's July 2024 prediction named four causes: poor data quality, inadequate risk controls, escalating costs and unclear business value.¹

So our scoring model covers all four. Value, feasibility (including run cost) and data readiness are scored. Risk is a pass/fail gate, not a weight that can be averaged away.

1,367

products retired from one software category in a year

The 2026 Marketing Technology Landscape added 1,488 products and retired 1,367 — about 124 in and 114 out every month.⁴ Picking a vendor is partly a bet that it still exists at renewal.

So "buy" decisions carry a viability check. Funding, customer concentration and acquisition risk are recorded for every vendor named on the roadmap.

Gartner figures are predictions and survey results, not measurements of your organisation. That is the point of the engagement: it replaces industry averages with your own numbers.

SIX WEEKS, WEEK BY WEEK

1 SETUP + 5 DELIVERY

WEEK 0 Mobilise	WEEK 1 Inventory	WEEK 2 Discover	WEEK 3 Test	WEEK 4 Score & sequence	WEEK 5 Decide
---------------------------	----------------------------	---------------------------	-----------------------	-----------------------------------	-------------------------

WEEK	WHAT HAPPENS	YOUR TIME	OUTPUT
0 • Mobilise	Kickoff with the sponsor. Agree scope, success criteria and who we talk to. Your security team approves read-only access for Shadow Scanner. Workshop invitations go out. The free AI Readiness Scorecard goes to 8–15 leaders.	Sponsor 1.5 h · IT/security 2–4 h · Scorecard 8 min each	Signed scope, access approved, calendar locked
1 • Inventory	Shadow Scanner runs. Current-state interviews (45 min each) cover data, identity, existing pilots and budget.	6–10 interviewees, 45 min each	Live footprint report and posture baseline
2 • Discover	Two discovery workshops with the business, seeded with scan findings: where people already reach for AI is often where the need is. Typically 20–40 candidate use cases.	Participants 2 × 3 h	Candidate long list, each with a proposed owner
3 • Test	Data-readiness checks on the shortlist (usually 8–12): owner named, reachable by software, quality sampled, usage rights confirmed. Value estimates worked through with finance.	Data owners ~1 h each · Finance 2 h	Readiness evidence per use case
4 • Score & sequence	Scoring workshop. Risk gate applied. Build, buy or leave positions written. Vendor viability checked. The roadmap is drafted as decision points with unlock conditions.	Participants 3 h · Sponsor 1 h	Scored register, draft roadmap
5 • Decide	Sponsor review of the draft. Executive readout. Handover of the register and roadmap. The first quarterly re-scan date is set.	Sponsor 1.5 h · Executives 1.5 h	Final roadmap, readout deck, one-page summary

Why six weeks, not three or four. The calendar, not the analysis, sets the pace. Access approvals for the scan, two rounds of workshops and data owners' availability each take a week of elapsed time in most organisations. We would rather quote six and deliver six.

Can it go faster? Yes, if access is approved before kickoff and workshop dates are already held. Week 0 then overlaps week 1 and the engagement runs in five weeks.

THE FIVE SESSIONS

WHAT HAPPENS IN THE ROOM

SESSION	LENGTH · WHO	AGENDA	YOU BRING
1. Kickoff Week 0	90 min Sponsor, IT lead, us	What “good” looks like in six weeks. Interviewee list. Scan scope and data handling agreed in writing. Known constraints: budget cycle, regulation, live contracts.	Org chart, current AI policy if any, list of live pilots
2. Discovery A Week 2	3 h Business leads, 6-12	Scan findings, team by team. Where work is slow, repeated or error-prone. Ideas captured with who would own them. No scoring yet.	Two or three processes that frustrate you
3. Discovery B Week 2	3 h Same group + IT, data	Group and de-duplicate the long list. First feasibility read. Agree the shortlist that goes to data testing.	Data owners on call for questions
4. Scoring Week 4	3 h Sponsor, leads, finance, risk	Walk the evidence for each shortlisted case. Score together on the model overleaf. Apply the risk gate. Decide build, buy or leave.	Finance view of value ranges
5. Readout Week 5	90 min Executive team	The roadmap, the reasoning, what we recommend not doing, and the first decision you need to make.	Decision-makers in the room

Remote or in person

All five sessions run well online. In-person discovery workshops in Toronto are available, and tend to produce longer and better lists. Travel elsewhere is quoted separately.

How Shadow Scanner is used

Read-only, scoped in writing at kickoff. Findings are attributed to teams so the workshops can talk about real behaviour. We route findings to the people who can close the gap, not into performance reviews.

THE SCORING MODEL

THREE SCORES · ONE GATE · FOUR OUTCOMES

MEASURE	WEIGHT	WHAT A 5 LOOKS LIKE	WHAT A 1 LOOKS LIKE
Value	40%	A measurable gain in cost, revenue, time or risk, estimated with finance, on a process with real volume.	"It would be nice." No owner can name the number it moves.
Feasibility	30%	Fits existing systems and identity. Run cost modelled at real volume, not pilot volume. Skills in-house or easy to get.	Needs a new platform, a new team and a budget that does not exist yet.
Data readiness	30%	Owner named. Reachable by software. Quality sampled. Usage rights confirmed.	Spread across inboxes and documents, with nobody accountable for it.
Risk gate	Pass / fail	Regulatory exposure, usage restrictions, safety of any autonomous action, and whether it can be removed cleanly. A fail is never averaged away. It becomes a condition the use case must meet first.	

Every use case lands in one of four places

GO NOW	High score, gate passed. Goes into the first decision point with a named owner and a success measure.
PREPARE	Worth doing, but a condition is unmet: usually data access, an owner or a control. The roadmap says what to fix and by when.
NOT YET	Possible later, when something changes: a contract renewal, a platform upgrade, a price. The trigger is written down.
DROP	Low value or a hard risk fail. Recorded with the reason, so it is not reopened without new evidence.

ILLUSTRATIVE FORMAT · NOT CLIENT DATA

Decision point 2 · target Q3: extend the support-ticket summariser from the pilot team to all support queues.

Unlocks when: (1) the ticket system API replaces the weekly CSV export; (2) the support lead signs off on accuracy against 200 sampled tickets; (3) spend caps are enforced at the platform; (4) run cost at full volume stays inside the agreed budget line.

If it does not unlock by Q4: move to "not yet" and review at the next contract renewal.

WHAT WE NEED FROM YOU

THE ENGAGEMENT ONLY WORKS WITH THESE

- **One executive sponsor** with authority to approve the roadmap — about 10 hours over six weeks.
- **Read-only access** for Shadow Scanner, approved by your security team in week 0.
- **6–12 workshop participants** from the business lines the plan will affect.
- **Named data owners** for the shortlisted use cases, about an hour each in week 3.
- **A finance contact** for two hours to pressure-test value ranges.

If one is missing

Without a sponsor, the roadmap has nobody to approve it, and we will say so at kickoff rather than in week 5. Without scan access, we fall back to interviews and the Scorecard. The report will then say clearly which findings are self-reported.

FIT WITH THE FRAMEWORKS

SO THE WORK IS REUSABLE

FRAMEWORK	WHERE THIS ENGAGEMENT FEEDS IT
NIST AI Risk Management Framework 1.0 (NIST AI 100-1, Jan 2023) ⁵	The inventory and context work supports the framework’s <i>Map</i> function. The posture baseline and re-scan support <i>Measure</i> . Named owners and the risk gate give <i>Govern</i> and <i>Manage</i> something concrete to start from.
ISO/IEC 42001:2023 AI management systems ⁶	The register, owners and review cadence are structured so they can feed an AI management system if you pursue one. This engagement is not a certification or an audit.

AFTER THE ROADMAP

OPTIONAL, NEVER ASSUMED

1.3 Board Briefing Pack

The roadmap translated for the people who approve it.

1.4 Quarterly Review

A re-scan and a session each quarter to keep the plan current.

2.2 • 4.2 • 5.1

Bake-off, pilot-to-production, or agent risk, where the roadmap calls for it.

Service status changes over time. The current status of every engagement is published at the808.org/services/.

START HERE

TWO WAYS IN

Free: AI Readiness Scorecard

Eighteen questions on data, identity, governance and skills. Eight minutes, no email, scored in your browser. It is also the week 0 input to this engagement.

the808.org/services/scorecard/

A 30-minute scoping call

We confirm fit, scope and dates, then send a fixed-fee quote. If the Scorecard says you are not ready, we will tell you what to fix first.

contact@the808.org ·

the808.org/about/contact/

ABOUT THE 808

TORONTO · INDEPENDENT

The 808 is an independent advisory on cloud, AI and security, based in Toronto. It takes no referral fees from the vendors it evaluates and publishes The Weekly Brief, a ranked and source-checked digest of enterprise software, SaaS, AI and security news. The firm is led by Andrew John Baynes, who has 15 years of experience in enterprise technology: cloud infrastructure, platform strategy and vendor evaluation.

SOURCES

CHECK EVERY NUMBER

1. Gartner, "Gartner Predicts 30% of Generative AI Projects Will Be Abandoned After Proof of Concept By End of 2025," press release, 29 July 2024. gartner.com/en/newsroom/press-releases/2024-07-29-...
2. Microsoft and LinkedIn, "2024 Work Trend Index Annual Report," 8 May 2024. Survey of 31,000 people in 31 countries. blogs.microsoft.com/blog/2024/05/08/...
3. Gartner, "Lack of AI-Ready Data Puts AI Projects at Risk," press release, 26 February 2025. Survey of 1,203 data management leaders, July 2024. gartner.com/en/newsroom/press-releases/2025-02-26-...
4. chiefmartec / MartechTribe, "2026 Marketing Technology Landscape," May 2026. chiefmartec.com/2026/05/...
5. NIST, "Artificial Intelligence Risk Management Framework (AI RMF 1.0)," NIST AI 100-1, 26 January 2023. nist.gov/news-events/news/2023/01/...
6. ISO/IEC 42001:2023, "Information technology — Artificial intelligence — Management system." iso.org/standard/42001

This brief describes the standard engagement. The scope, dates and fee for your organisation are set out in a separate statement of work.