

HUMANSPARK AI ADOPTION SERIES · 2027 EDITION

AI Adoption Roadmap

Run People, Process, Position, and Policy in parallel to extract real value from AI

BY **Alastair McDermott**

Executive Summary

Two years ago, AI adoption looked like a sequence: assess, pilot, expand, transform. That model is now too slow.

Since 2024, three things have shifted. The tools are good enough to deploy broadly. Your staff are already using AI on their phones, with your company data, whether you've sanctioned it or not. And governance built in at the end is too late.

This edition of the AI Adoption Roadmap reflects that change. The four stages aren't gone. They've collapsed from a sequence into four parallel tracks that you run from day one:

People.

Capable AI tools in every hand; team trained to use them. Target within twelve months: daily use in the roles where AI fits the work, weekly use everywhere else.

Process.

Systematically discover, score, and pilot the repetitive work that shouldn't need humans. One successful rollout per quarter.

Position.

Use AI to do what your business couldn't do before. Slower to compound; biggest payoff over three to five years.

Policy.

The rules running alongside all three. Start as a one-pager. Mature as the other tracks deepen.

The first 90 days lay the runway. Day 30: one-page policy, primary vendor selected, tools deployed, literacy training delivered, task discovery underway. Day 60: first Process pilot running with a fixed kill-date.

CONTINUED

Executive Summary

Day 90: pilot shipped or archived - never "continue piloting."

Then you repeat, quarterly. The compound effect over a year outweighs any single transformation project. This Roadmap is how you get on the right side of it.

The Four Ps: Your Parallel-Track Strategy

Part of a Complete AI Toolkit

This AI Adoption Roadmap is one of three companion whitepapers:

1. **AI Adoption Roadmap** (this paper) - the practical playbook for running an AI programme: People, Process, Position, Policy.
2. **Risk and AI Adoption** - the safety, compliance, and governance companion. Read alongside Policy.
3. **Opportunity and AI Adoption** - the value-creation companion. Reads naturally with Process and Position.

All three share the same Four Ps spine; each takes a different angle.

WHY THE MODEL CHANGED

We dropped the four stages because the tools got good enough and staff were already using them. Governing at the end no longer works, so all four tracks start on day one.

Why Four Ps and not Four Stages

The first edition of this Roadmap (March 2025) used a four-stage model: Assess and Prepare, Gain Momentum, Expand and Integrate, Transform and Lead. It was the right model for 2024 - tools weren't capable enough to deploy broadly, governance was an emerging discipline, and going one gate at a time made sense.

Three things changed.

The Four Ps: Your Parallel-Track Strategy

The tools got good enough. Capable models now do work that needed teams in early 2024. Microsoft 365 Copilot now ships Anthropic's Claude models alongside OpenAI's, default-on for commercial tenants outside the EU/UK from January 2026. Autonomous agents are real and shipping.

Your staff are already using AI. Every business now has shadow AI - people running free tools on their phones with company data, building custom assistants, pasting client information into unapproved systems. The choice "should we use AI" is already answered, by your staff. Your choice now is whether to channel it or let it run as unmanaged risk.

Governance at the end is too late. In the staged model, formal policy came at Stage 3. With shadow AI and rapid tool deployment, that runway is gone.

So the four stages compressed into four parallel tracks.

[SHADOW AI]

Staff using AI tools on personal accounts with company data, outside the organisation's sanctioned platforms. Accumulates compliance and security exposure in parallel with productivity gains.

The Four Ps at a glance

TRACK	WHAT IT DOES	DAY-1 COMMITMENT
People	Make every person more effective	Licensed tools to the whole workforce
Process	Systematically automate the repetitive work	One quarterly pilot rhythm, kill-dates fixed
Position	Do what your business couldn't do before	One Position experiment in motion within 90 days
Policy	Keep the other three safe	One-page policy live by day 30

The Four Ps: Your Parallel-Track Strategy

You run them in parallel from day one. The mechanics underneath each track are still recognisable from the four-stage model - surveys, Champions, sandbox pilots, governance committees. The change is the shape of the programme, not the techniques.

The next four sections take each track in turn.

| How the weighting changes by sector

The four tracks are universal. The weighting isn't:

- **Professional services** (law, accountancy, consulting, agencies): People usually leads. Time and judgement is your raw material.
- **Industrial and innovation-led SMEs**: Process often leads. Concrete, defensible operational wins; OT/IT considerations sit inside Policy.
- **Retail and consumer businesses**: Position often leads, because customer-facing AI is where the visible value sits.

Your specific weighting comes out of a Phase 1 Diagnostic for your business. The tracks themselves apply to all three.

Implementation Guide

TRACK 1: PEOPLE

Make every person in the business more effective. Not a pilot group. Everyone.

That single sentence is the People track. The mistake most organisations make is putting AI in the hands of a small group, hoping it diffuses, and watching nothing happen because diffusion doesn't beat shadow AI. Your staff already use free AI tools. Your choice: give them better tools, training, and support, or leave them improvising on consumer accounts.

[CHAMPIONS NETWORK]

Three to five people across departments whom colleagues already go to for help. I'd pick problem-solvers over presenters: people who fix real problems with real tools.

What good looks like

- **Licensed AI tools for everyone.** Pick a primary vendor with a backup. The realistic shortlist for most SMEs is ChatGPT Business, Claude Team, Microsoft 365 Copilot and Google Workspace Gemini.
- **Baseline literacy training for everyone.** Half a day, hands-on. Cover what AI is, what it does well, where it fails, and your Policy's data rules.

STAGE 1
People

STAGE 2
Process

STAGE 3
Position

STAGE 4
Policy

Implementation Guide

- **A Champions network for local support.**
Three to five people with deeper training, who colleagues go to first for help.
- **A target.** Within a year, daily use where AI fits the work (admin, sales, operations, finance), weekly use elsewhere.

| Survey before you deploy (briefly)

The first edition of this Roadmap had a long staff survey section. Since then, the survey has become faster and lighter. You're not discovering whether staff use AI - they do. You're calibrating. Five questions:

1. Which AI tools do you currently use, and for what?
2. What's something you do every Friday that drains your energy? (The "Friday Frustration" question - your richest source of pilot candidates.)
3. What concerns you most about AI in your work?
4. Would you be interested in becoming an AI Champion?
5. What's one task you wish AI could do for you?

Run this in the first two weeks. The results inform Process discovery and tell you who your Champions are.

Implementation Guide

| Champions, lighter

The Champions framework still works.

Lighter than v1 in two ways:

- **Fewer, deeper.** Three to five Champions across departments beat fifteen with shallow training.
- **Practitioner-led.** Champions solve real problems, not deliver training sessions. Pick problem-solvers, not presenters.

Train them as a half-day or full-day workshop. Cover how AI works in plain English, practical use cases, the risk and policy rules, and hands-on time with the tools. Then give them a Slack channel and a fortnightly check-in.

| Sector callout

For **professional services**, People leads. Your raw material is people's time and judgement, and the productivity uplift shows up fastest here.

For **industrial SMEs**, deploy People to the knowledge-worker layer first - engineering, R&D, finance, quality - then expand into operations as Champions emerge. Operators on the shop floor benefit later, often through Process work.

Implementation Guide

| Common mistakes to avoid

- Treating People as "give everyone an account and walk away." Without training and Champions, daily use usually stalls early.
- Locking the tools to a tiny pilot group "to manage risk." Shadow AI is the actual risk; restrictions accelerate it.
- Training without a tool. People forget within a week.
- Tooling without training. People use AI for what they already know it does, and miss most of the value.

Implementation Guide

TRACK 2: PROCESS

Systematically automate the repetitive work that shouldn't need humans.

The most common mistake is treating Process as a transformation programme. It isn't. It's a steady-state rhythm that runs quarterly and indefinitely, delivering small wins rather than a single big bang.

[FRIDAY FRUSTRATION]

Ask people what task drains them every Friday afternoon, week after week. The answers surface repetitive work fast, and the best of them become your pilot candidates.

What good looks like

- A regular **discovery question** to find candidates. I use the "Friday Frustration" question - what task is draining you on a Friday afternoon, week after week? Repetitive, annoying tasks float to the surface fast.
- A scoring framework. **RATES** - Repetitive, Annoying, Time-consuming, Error-prone, Scalable. Tasks scoring high on all five are Process candidates.
- A pilot rhythm with **fixed kill-dates**. Run a pilot for 30 days, then make a binary decision: scale it, or archive it. There is no "continue piloting" - that path leads to pilot purgatory and AI programmes quietly die there.
- A small, focused team per pilot. Five to eight people. Include a sceptic deliberately.

STAGE 1
People

STAGE 2
Process

STAGE 3
Position

STAGE 4
Policy

Implementation Guide

One concept to name early - it shapes every pilot decision:



Verification tax is the time spent checking the AI's work to make sure it's right. Every AI use case carries one. It's the truth-telling number on whether AI is genuinely saving you time - if checking takes longer than doing the task manually, the use case isn't ready, no matter how impressive the output looks. Always measure verification tax explicitly during a pilot, not just the time AI saves on the first-pass output.

[RATES SCORING]

We score each candidate on five things: Repetitive, Annoying, Time-consuming, Error-prone and Scalable. A task that scores high on all five is a strong candidate.

The AI Working Group

The AI Working Group is the body that runs Process and co-authors Policy. It is not a strategic committee that meets quarterly to debate the future of AI. It is the operating team that scores pilots, picks the next one, runs them, and makes the kill-or-scale decisions.

Membership, kept small:

- IT - infrastructure, data security, integration.
- Operations / Business units - the people whose work is changing.
- Marketing or Sales - if customer-facing AI is in scope.
- HR - for workforce impact and training.

Implementation Guide

- Legal or Compliance - lightweight, especially early on.
- AI Champions - hands-on people from the People track.
- Senior leader as sponsor - to signal authority.

Cadence: every two to four weeks, ninety minutes, with a Slack channel between meetings for tool-sharing.

Implementation Guide

| The pilot loop

A Process pilot has five steps:

1**Discover.**

Find candidate tasks (Friday Frustration question, staff input, AI Working Group brainstorm).

2**Score.**

Apply RATES. Pick the strongest candidate that's also in the **Green Zone** - low risk, internal-only, text-heavy, no regulated data. Save the harder candidates for later quarters.

3**Pilot.**

Small group, fixed kill-date (30 days is right for most things), defined success metrics agreed up front. Measure the verification tax - how long checking the AI's work takes versus doing it manually.

4**Decide.**

On kill-date, the call is binary. Scale to the broader user base, or archive the pilot. Both are good outcomes; "continue piloting" is the bad outcome.

5**Roll forward.**

If scaling, plan the rollout; either way, start the next discovery cycle.

One rollout a quarter is strong. Zero for three quarters means the pilots are too ambitious or kill-dates are slipping.

Implementation Guide

| Pilot categories worth considering

Today, Process candidates fall into four broad categories:

- **Automating repetitive tasks.** Meeting notes, report drafts, document summarisation, invoice coding, basic data entry.
- **Improving efficiency in existing processes.** Customer support triage, content curation, internal helpdesk responses, order routing.
- **Enhancing decision-making with AI insights.** Trend identification in customer feedback, sentiment analysis, anomaly detection in financial data.
- **Multi-step task delegation via lightweight agents.** The newest category. Agents completing small, bounded tasks - research and summarise, draft and route, ingest and classify - then hand back to humans. Higher verification tax; pilot more carefully.

[VERIFICATION TAX]

The time spent checking the AI's work. It's the honest number on whether AI is saving you time, and if checking takes longer than doing it yourself, it isn't ready to scale.

| Your AI Sandbox

An "AI Sandbox" - a controlled testing environment for new tools - still applies. Now, also include on-device and small-model AI for sensitive workloads where data never leaves your infrastructure. Several major vendors ship on-device or local models, and open source is moving fast.

Implementation Guide

This is now realistic for sandbox work that can't tolerate data leaving the building.

Sandbox principles:

- Synthetic data for testing. Never real customer or financial data in unapproved tools.
- Testing isolated from production systems.
- A small group of trusted testers; not the whole workforce.
- Document everything tested, especially the failures.
- Coordinate with IT before you start.

[KILL-DATE]

A fixed end date for a pilot, usually 30 days out. On that date the call is binary, scale it or archive it, because a pilot that just keeps going is where projects quietly stall.

Sector callout

For **industrial SMEs**, Process often leads the entire programme. There is real repetitive operational work to automate - invoice coding, site log transcription, compliance reporting, document review, technical literature scanning. Wins here are concrete and quickly defensible to a board.

For **professional services**, Process follows People closely. Back-office work like proposal drafting, research synthesis, and document review is unusually responsive to current tools.

Implementation Guide

TRACK 3: POSITION

Use AI to do what your business couldn't do before.

Most organisations under-invest in Position because the first two tracks feel easier to measure. People shows up as productivity. Process shows up as cost or hours freed. Position shows up as competitive advantage over three to five years. By the time you can measure it, your competitors are already ahead.

This is why Position must start in parallel from day one, even if it starts small.

| What Position looks like

The pattern is the same in every sector: AI lets you do something that used to require specialists, used to take much longer, or used to be economically impossible.

STAGE 1
People

STAGE 2
Process

STAGE 3
Position

STAGE 4
Policy

Implementation Guide

A few practical examples:

- Customer segmentation that used to take three months of consulting work now takes three days using a frontier model and your customer data.
- A proposal that used to need a specialist on the team can be drafted by a generalist in an afternoon and tightened by the specialist after.
- A forecast that used to run quarterly because it was expensive to produce now runs weekly because the cost has fallen.
- A sales call that used to be analysed by sampling 5% can have 100% reviewed for coaching insights.

Position work starts with one person experimenting with a model on a strategic question, growing as confidence builds.

| How to start small

The most common Position-track first move:

1. Pick a strategic decision your leadership team faces every quarter. Pricing review. Customer cohort analysis. Competitive positioning for a product line. Candidate selection for a new market.
2. Frame it as a structured prompt to a frontier model with full business context - the data, the constraints, the history.

Implementation Guide

3. Compare what the model produces against what your team produces. Not to replace your team - to identify gaps. What did they miss? What did the model miss?
4. Iterate. Within a quarter you have a useful supplement. Within a year, a new decision-making practice.

This costs one person a few hours a week. The point is not immediate ROI; it's to build organisational comfort with using AI on questions that matter.

A practical example: a manufacturing company runs a monthly production forecast. Today it takes the Operations lead a day to gather data and run scenarios. The Position experiment: pose the same decision to a frontier model with six months of historical shipments, inventory turns, and customer bookings. In the first month, the model's forecast is wrong because it missed a business rule nobody documented. By month three, the model and the Operations lead review in parallel; disagreement exposes blind spots on both sides. By month six, the Operations lead uses the model's forecast as a starting point and adjusts for known external factors - a 60% time saving.

Implementation Guide

When to start a Position experiment vs. deferring

Start early if:

- Your decision is recurring (monthly, quarterly) so iteration happens naturally.
- You have historical data available - at least six months - that the model can learn from.
- The cost of a wrong decision is high enough to justify an hour or two per cycle of model review.

Defer until later if:

- You make the decision once a year. Iteration cycles are too slow.
- You have no data, or the data is small or chaotic. The model will hallucinate.
- The decision is low-stakes enough that gut feel is fine. Use the budget elsewhere.

If you defer, revisit quarterly. As data builds or decisions become more frequent, the Position experiment becomes sensible.

Implementation Guide

| The five priority areas for Position

Today, Position investment usually falls into one or more of these:

- **Strategic AI use cases aligned with business goals.** Serve goals the business already has, not let AI suggest the goals.
- **Scalable data infrastructure and AI platforms.** Position depends on data being accessible. Treat data infrastructure as a Position-enabling investment, not a separate IT project.
- **Ongoing AI skills development.** This overlaps with People, but Position needs deeper skills - prompt engineering, evaluation design, agent supervision.
- **AI governance structures.** As Position work scales, the AI Working Group from Process becomes the de facto AI governance body. Recognise that and resource it.
- **Continuous monitoring and adaptation.** Position assumes AI models will keep changing. Annual strategic reviews are too slow; quarterly reviews are about right.

Implementation Guide

| What to watch next

Three developments worth tracking for Position investment:

- **Model Context Protocol (MCP) and similar interop standards.** How models connect to enterprise data is becoming the strategic question. The model itself matters less; the plumbing matters more.
- **On-device and small-model AI.** As on-device options from the major vendors and open-source models mature, Position work on sensitive data becomes realistic without sending data to a third party.
- **Agentic workflows running for hours.** The question now isn't "can a model do this task" but "can an agent run a multi-step workflow with supervision." Pilot carefully; verification tax is the unsolved problem.

| Sector callout

For **retail and consumer businesses**, Position often leads the whole programme. Customer-facing AI - personalisation, search, recommendation, support agents - is where AI shows up most visibly to your customer and is therefore where competitive advantage shows up first.

Implementation Guide

For **industrial SMEs**, Position shows up in forecasting, quality control, predictive maintenance, and customer insight. Often the biggest medium-term wins, but they need People and Process working before the investment makes sense.



Implementation Guide

TRACK 4: POLICY

Policy keeps the other three tracks safe, running alongside them, not after.

The first edition put Policy at Stage 3 - after surveys, training, working groups, and pilots. That order is wrong now. Your staff are already using AI on personal accounts with company data. Every day without a policy, you accumulate liability. The right move: a one-page Policy v0.1 in the first thirty days, mature it from there.

| What good looks like

- **A one-page policy by day 30.** Better than a twenty-page policy nobody reads, written in month four.
- **A short list of approved tools, and the data each one may handle.** Keep it easy to update. Confidential customer data goes only into approved tools. Public or draft material can go into any tool your AI Working Group has cleared.
- **Clear rules for the common cases:** emails, meeting transcripts, source code and client information.

STAGE 1
People

STAGE 2
Process

STAGE 3
Position

STAGE 4
Policy

Implementation Guide

- **A vendor check.** Ask more than what the tool does. Where does our data go, how is it stored, who can access it, and how do we get it back if we leave?
- **A review cadence.** Quarterly for the first year, then annually, with extra reviews when new tools or regulations appear.
- **Policy ownership.** One person accountable - the AI Working Group's chair, the compliance lead or, in smaller firms, the founder or MD.

| The minimum viable Policy v0.1

A one-page policy that covers:

- Approved tools list (with the version / tier - free vs paid, business vs personal account).
- Data classes by sensitivity, with rules for each class.
- The handful of explicit prohibitions (no client data in unapproved tools, no source code in personal accounts, etc.).
- Who to ask when in doubt.
- How to report an incident.

Don't try to write the comprehensive document first. The one-pager gets you operating safely; the comprehensive document gets written once you know which sections actually matter for your business.

Implementation Guide

| What the comprehensive policy covers

By month three or four, the one-pager grows into something more detailed. Standard sections:

- **Approved vs. unapproved AI tools** with the request-and-evaluation process for new ones.
- **Data usage guidelines** - which data classes belong in which tools.
- **Responsible AI use** - human review requirements for AI-generated content, especially anything customer-facing or high-stakes.
- **Human oversight** - where AI augments and where humans must remain in the loop.
- **Documentation, accountability, and training** - tracking what's in use, who's responsible, and what training is required.
- **Incident response** - what to do when something goes wrong.

Implementation Guide

| EU AI Act and the regulatory context

The EU AI Act came into force on 1 August 2024 and applies in stages: prohibited practices and AI literacy obligations from 2 February 2025; general-purpose AI model obligations from 2 August 2025; most remaining provisions from 2 August 2026; high-risk obligations from 2 December 2027, deferred by a 2026 amendment; high-risk AI in products already covered by EU product-safety law from 2 August 2028.

If your business operates in the EU - wholly or through EU customers - map your AI use cases to the EU AI Act risk categories (Unacceptable, High, Limited, Minimal) and apply the appropriate obligations.

If your business is outside the EU, the Act still matters:

- Customers, partners, and procurement ask about EU AI Act alignment.
- Other jurisdictions (Colorado, California, UK) follow similar patterns.
- Building to the EU AI Act now avoids costly retrofitting later.

Read the *Risk and AI Adoption* companion paper for detailed treatment of the EU AI Act and its operational implications.

Implementation Guide

| Sector callout

For **industrial SMEs**, Policy needs to account for OT (operational technology) environments alongside IT, and for IP-sensitive R&D data that absolutely cannot leak. On-device and locally-run models are relevant here.

For **retail and consumer businesses**, Policy is heavily weighted toward consumer data, consent, and disclosure obligations.

For **professional services** with regulated client data (legal, financial, healthcare advisory), client confidentiality clauses and sector-specific regulator guidance set the floor for the Policy.

The First 90 Days

The Four Ps run in parallel. Here is what running them in parallel actually looks like for the first quarter.

→ → ALL FOUR TRACKS

All four tracks start together: People gets tools out, Process picks candidates, Position names its questions and Policy drafts a one-page policy. That sets the rhythm.

Days 1-30: Lay the runway

By the end of month one you should have:

- **Policy v0.1** signed off. One page. Approved tools, approved data classes, the obvious prohibitions, who to ask.
- **A primary vendor selected** with a backup. Don't agonise; the shortlist is small (Microsoft 365 Copilot, ChatGPT Business, Claude Team, Google Workspace Gemini) and any of them moves you forward.
- **Tools deployed to the whole workforce** - not a pilot group.
- **Baseline literacy training** - half a day, hands-on, for everyone.
- **Staff survey** completed - the lighter one from the People track. Identifies your Champions and your initial Process candidates.
- **AI Working Group** formed - cross-functional, small, with clear remit.

STAGE 1
People

STAGE 2
Process

STAGE 3
Position

STAGE 4
Policy

The First 90 Days

- **First Process candidates** scored with RATES - strongest two or three shortlisted for the first pilot.

This is a busy month. Most of it can be done in parallel rather than sequentially. The bottleneck is usually procurement and IT setup for the workforce-wide tool deployment - start that on day one.

| Days 31-60: Run your first pilot

By the end of month two:

- **First Process pilot is live** - a Green Zone candidate (low-risk, internal, text-heavy).
- **Pilot group of five to eight people**, with a sceptic included.
- **Fixed kill-date 30 days from start**. If pilot kicks off day 35, kill-date is day 65. Success metrics agreed upfront.
- **Champions** deepen training and start supporting colleagues.
- **First Position experiment** is in motion - one person trying a recurring strategic question with AI assistance.
- **Policy** gets its first revision. The one-pager raises questions in week three or four; capture them.

The First 90 Days

| Days 61-90: Decision point

By the end of month three:

- **Pilot reaches kill-date.** Binary decision: scale or archive. No "continue piloting" - that's where most AI programmes die.
- **Verification tax measured.** How long checking the AI's work took versus doing it manually - the truth-telling number. If it exceeded the time saving, archive and try another candidate.
- **Next pilot is lined up.** Discovery has been running in parallel; the next candidate is ready.
- **Position track** has produced its first useful artefact - decision-supporting analysis, customer insight, forecast - even if used by one team so far.
- **Policy** is now a 3 - 5 page document addressing first-pilot issues.
- **First quarterly review** with the AI Working Group: what worked, what didn't, what to adjust.

| What you should see by day 90

- Weekly use in every role where AI fits the work.
- One pilot completed - either scaled or archived, both fine.

The First 90 Days

- One Position experiment in motion.
- A Policy that has matured beyond the original one-pager.
- A working AI Working Group with a quarterly cadence.

| Then you repeat

Quarter two looks similar in shape:

- People track: move those roles from weekly to daily use, deeper Champion engagement.
- Process track: next pilot, same loop.
- Position track: second experiment, building on what worked in quarter one.
- Policy track: first formal quarterly review, address gaps.

The compound effect over a year outweighs any single big-bang transformation project. By the end of year one, you should have run three or four Process pilots, have daily use in the roles where AI fits and weekly use everywhere else, your first Position experiment in motion, and a Policy that has matured from a one-pager into something proper.

The First 90 Days

Common ways the first 90 days goes wrong

- **Tool deployment delayed.** Shadow AI accumulates while you wait. Push hard on procurement.
- **Pilot picked from "interesting" ideas rather than RATES scoring.** "Interesting" isn't a criterion. Score it.
- **Pilot group too big.** Five to eight is right. Twenty is a rollout, not a pilot.
- **Kill-date softened.** The kill-date is the discipline. Slip it once, and pilot purgatory starts.
- **No Position experiment.** Easy to defer; pays off over years. Don't skip it.
- **Policy as a future project.** A one-pager in week one beats a perfect document in month four.

05

Next Steps



Next Steps

| Take the next step

You've seen the four tracks and the first ninety days. Putting them into practice turns a framework into a result. A few concrete next moves:

| Within your organisation this week

- **Run the Friday Frustration question** at your next leadership meeting. Five minutes. You'll spot at least two Process candidates.
- **Audit your shadow AI.** Ask three people what AI tools they use on personal accounts for work. This tells you what your Policy needs to address first.
- **Pick a primary vendor.** Microsoft 365 Copilot, ChatGPT Business, Claude Team, Google Workspace Gemini - any moves you forward. Switching later is cheaper than waiting now.

| Within the next 30 days

- Draft your one-page Policy v0.1.
- Form your AI Working Group.
- Deploy tools to the whole workforce.
- Run baseline literacy training.
- Start your first Position experiment - one person, one strategic question, one quarterly cycle.

Next Steps

| Build knowledge

- Connect with local tech communities and AI practitioner groups. Peer learning beats vendor pitches.
- Follow practitioners on LinkedIn - people with hands-on implementation experience. Find me at [@AlastairMcDermott](https://www.linkedin.com/company/humansparkai).

| Ready to move forward with AI?

The Four Ps framework works best when tailored to your specific situation. If you want to talk through where to start, the clearest next step is a 25-minute Focus Call.

I'll listen to where you are now, name the one track that should lead for your business, and we'll explore whether a Phase 1 Diagnostic makes sense - a structured assessment that maps your current state across the Four Ps and finds the highest-impact next move.

Book a time at cal.humanspark.ai/alastair/focus-call.

If you lead an industrial or innovation-driven SME in Ireland, the **IRDG AI Transformation for Business Leaders Programme** runs the Four Ps end-to-end across six weeks with a cohort of peers. Details at humanspark.ai.

Next Steps

| Why work with HumanSpark?

I'm Alastair McDermott, author of this whitepaper and founder of HumanSpark. My focus is Human-First AI: using deep technology expertise and practical business strategy to make AI adoption work for the people doing the work, not just for the spreadsheet.

My background runs from mainframe assembler programming through to modern AI systems - twenty years of business advisory work as a consultant, trainer, and author of several books in the Human-First AI Series.


What you can expect from working with HumanSpark:

- Real-world implementation experience across multiple sectors.
- Focus on rapid, demonstrable returns rather than transformation theatre.
- Plain English communication that builds confident teams.
- A track record of AI projects that ship and stay shipped.

If you're ready, the Focus Call is the place to start. If not, the rest of the toolkit is here to help you get there.

06

Resources



Resources

| Your Complete AI Toolkit

This document is part of a set of three companion resources that share the Four Ps spine. You are reading the **AI Adoption Roadmap**, which answers:



"What does an AI programme actually look like, and how do I run it?"

The other two:

- **Risk and AI Adoption** - the safety, governance, and compliance companion. Covers the EU AI Act, the risk taxonomy, and the operational practices that keep your Policy track honest.
- **Opportunity and AI Adoption** - the value-creation companion. Covers the Four Pillars (Productivity, Capabilities, Decision-Making, Accelerated Learning), the data primer, and forward-looking trends.

All three share the same Four Ps spine:

- **People** - workforce-wide capability
- **Process** - systematic automation rhythm

Resources

- **Position** - new capabilities and competitive advantage
- **Policy** - the governance thread running alongside

Download all three at
humanspark.ai/toolkit.

Companion templates and guides

- **RATES Scoring Worksheet** - score your Process candidates against Repetitive, Annoying, Time-consuming, Error-prone, Scalable, with a pre-built scoring sheet and examples. humanspark.ai/rates
- **SparkOS Guide** - the seven-component custom-assistant system for building AI that actually knows your business.
humanspark.ai/sparkos
- **Voice Profile Template** - the Core Voice Profile framework for teaching AI to write in your authentic voice. humanspark.ai/voice-profile
- **Red Team Prompt** - the adversarial-review prompt for stress-testing your AI work.
humanspark.ai/red-team
- **AI Tools Comparison** - a platform comparison covering ChatGPT, Claude, Microsoft Copilot, and Gemini for SME use.
humanspark.ai/tools

Resources

| Talk to me directly

If you'd like to work through where to start with the Four Ps in your own business, book a 25-minute Focus Call.

cal.humanspark.ai/alastair/focus-call