



FROM PRINCIPLES TO A 90-DAY PLAN

Leading AI Transformation.

A practical playbook for executives who want AI adoption that is real, measured and enterprise-wide — not another round of pilots. The operating model, the maturity ladder, the governance and the first 90 days.

3

ADOPTION STREAMS,
RUN IN PARALLEL

5

MATURITY STAGES,
EVIDENCE-GATED

90

DAYS TO VISIBLE WINS
AND REAL FOUNDATIONS

THE UNCOMFORTABLE TRUTH

Pilots are easy. Adoption is hard.

Most organisations don't have an AI tooling problem. They have an adoption problem — and it looks the same everywhere.

Licences are bought, a pilot impresses in a demo, a few enthusiasts change how they work — and six months later the organisation is doing roughly what it did before, plus a subscription bill. The gap is rarely technology. It is the absence of an **operating model for adoption**: clear ownership, staged progression, honest measurement and the routines that turn individual enthusiasm into organisational capability.

This guide sets out the model we use to run enterprise-wide AI adoption programmes — the same model behind engagements that reached **100% adoption and a 30% productivity uplift within six months**. It is written for the executives who will sponsor the change: what to structure, what to measure, what to demand as evidence, and what to expect in the first 90 days.

**Momentum without results is theatre.
Results without momentum is luck. You
need to see both.**

Everything that follows is designed to produce both: leading indicators that show adoption is happening, and key results that prove it is working — with every claim backed by working examples, not sentiment.

IN THIS GUIDE

- **Four principles** that guide every adoption decision
- **Three streams** — why one programme must not pretend every team is the same
- **The maturity ladder** — five stages, each one earning the next
- **People & capability** — training, workshops, pair working, champions
- **Governance & measurement** — stage gates and metrics that can't be gamed
- **The first 90 days** — visible wins and real foundations

PRINCIPLES

Four principles that guide every decision.

When a hard call comes up — and it will — these are the tie-breakers. Agree them with your leadership team before the programme starts, not during the first crisis.

01

Diagnose before prescribing

Where adoption is slow, find out why before adding more tools or training. Fear of looking slow, unclear permission, a broken workflow and a missing skill all look identical from a dashboard — and each needs a different fix. Work surgically, down to the level of a single team or individual.

02

Climb the ladder — don't leap

Trust is earned at every rung: AI as assistant, enhanced individual practice, team workflows, then bounded autonomy. Teams that leap to automation without the habits underneath produce fragile wins that collapse under real load — and burn credibility you will need later.

03

Measure momentum and results

Leading indicators tell you adoption is happening; key results tell you it is working. Momentum without results is theatre. Results without momentum is luck. Insist on seeing both, every month, from a single source of truth.

04

Short wins now, strategic bets for later

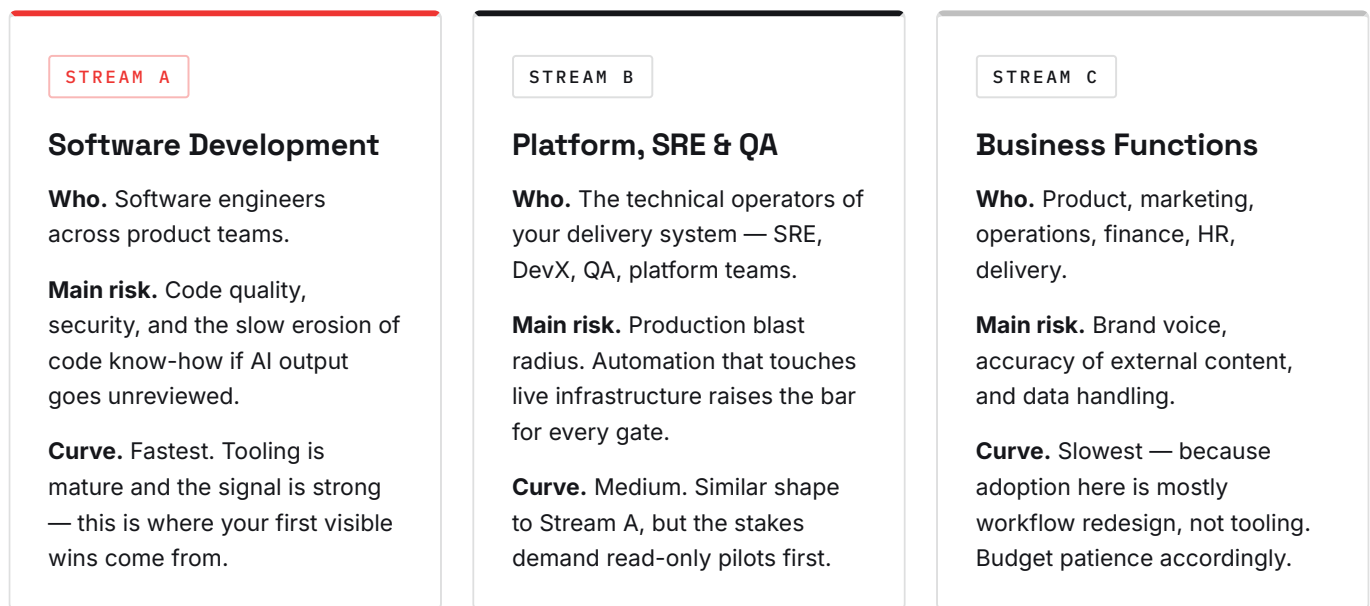
Every stream commits to both: visible quick wins in the first 30 days that build belief, and longer-horizon bets that compound over quarters. A programme with only quick wins plateaus; one with only long bets loses its mandate before it pays off.

The executive's role: you don't need to pick the tools. You need to hold the programme to these principles — especially when someone proposes skipping a stage "because the demo looked ready".

STRUCTURE

Three streams, deliberately split.

Different users, different risks, different maturity curves. One programme should not pretend they are the same — that pretence is why "AI rollouts" stall. Run three streams in parallel, each with a named owner.



Ownership: one name per stream

Each stream needs a single accountable owner with the authority to act — not a committee. Add an independent **metrics owner** who verifies evidence across all streams, and named subject-matter champions inside each function who connect programme intent to daily practice. If nobody owns a stream, that stream is where your programme will quietly die.

Why the split matters: a rollout that treats a marketing team like an engineering team will bore one and endanger the other. Same stages, same governance — but each stream climbs its own ladder at its own pace, with its own risks controlled.

THE AI AUGMENTATION MATURITY MODEL (AAMM)

Five stages. Each one earns the next.

Maturity is measured objectively and teams advance through evidence-based stage gates. You only progress when the data proves you're ready — no theatre, no vanity metrics.

STAGE	WHAT IT LOOKS LIKE	MOVE ON WHEN...
0 Aware	Ad-hoc, unmanaged AI use. Policy, approved tooling and baselines being established.	Everyone has safe, approved access; a usage policy is live; baseline metrics are captured.
1 Assisted	AI assists individuals on discrete, transactional tasks — drafting, querying, generating — with guardrails.	Most of the team has working access and uses AI daily, not just the enthusiasts.
2 Augmented	Best practice is embedded in individual workflows: shared configuration, prompt and context libraries, AI-authored output reviewed and shipped by humans.	Configurations are in place and AI-assisted work lands cleanly on a weekly basis, with measurable productivity gains.
3 Agentic	Standardised team workflows and playbooks. AI executes isolated, well-bounded flows end-to-end — triggered by humans.	At least one flow runs reliably end-to-end with AI doing the bulk of the work, and the practice is shared by the team — not one champion.
4 Autonomous	AI executes scoped workflows on automatic triggers — events, schedules, thresholds — within defined guardrails: bounded blast radius, override paths, humans in the loop. People review outcomes, not every step.	Stable when automated workflows deliver direct outcomes — e.g. a triggered audit finds an issue, drafts the fix, passes regression, and a human approves the change into production.

The ladder is climbed one rung at a time. A team that leaps is a team you will be rolling back.

PARALLEL LADDERS

Same stages, different shapes.

The ladder is one model, but each stream climbs it through its own practices.
Here is what the stages actually look like on the ground.

	STAGE 1 – ASSISTED	STAGE 2 – AUGMENTED	STAGE 3 – AGENTIC	STAGE 4 – AUTONOMOUS
Software Development	Ad-hoc prompting for code, tests and docs.	Project context files and IDE-integrated generation; AI-authored commits with human review.	Standard playbooks per repository. AI runs PR triage, test generation and refactor suggestions on demand.	AI opens pull requests from issues and runs scoped refactors on schedule, within a bounded blast radius.
Platform, SRE & QA	Engineers query AI for runbook summaries, log analysis and query authoring.	Shared prompt libraries for incident review; infrastructure-as-code reviews assisted by AI.	Automated change linting, alert summarisation and on-call handoff packs.	AI proposes and stages remediation for known failure patterns. A human approves before apply.
Business Functions	Individuals draft, summarise and brainstorm with an AI assistant.	Team-shared prompts, content templates and brand-checked outputs.	Standardised flows: launch briefs, recurring reports, ticket triage, research synthesis.	Scheduled briefings, automated first-pass reports, and monitoring of recurring inputs.

Read the table sideways...

to see how one stage translates across your organisation. Stage 3 in engineering (AI-run PR triage) and Stage 3 in marketing (standardised launch briefs) are the same organisational muscle: **a repeatable, team-owned workflow** — not a hero with a chat window.

...and downwards

to see why streams move at different speeds. Expect Stream A to reach Stage 3 while Stream C is still consolidating Stage 2. That is not failure — it is the model working. Force-marching every stream in lockstep produces exactly the brittle adoption you're trying to avoid.

CAPABILITY, NOT DEPENDENCE

How people actually learn to work this way.

Tools don't transform organisations; changed working habits do. Five mechanisms move people from "I tried it once" to "this is how we work" — and they are all social, not technical.

A

Role-based training

Not generic prompt courses. Engineers, operators and business teams get training built around their real work, in their real tools, on their real backlog. Training that doesn't touch Monday morning's tasks is entertainment.

B

Hands-on workshops

Each team brings its own live work and leaves with a working AI-assisted workflow — a configured repository, a functioning report pipeline, a reusable template. The output of a workshop is an artefact, never a slide deck.

C

Pair working

An experienced practitioner sits with your people on real tickets and real deliverables — modelling the habits, unblocking in real time. It is the fastest known cure for the "I tried it once and it didn't work" plateau, and it transfers judgement no course can.

D

Champions in every stream

Named champions carry best practice into their teams, deliver knowledge-share sessions, and surface what leadership can't see. Single-champion adoption doesn't pass a stage gate — but champion-led adoption is how every stage gets passed.

E Leader 1:1s and forums

Regular 1:1s with team leads surface the themes dashboards miss — blockers, wins, requests, quiet resistance. A recurring knowledge-share forum turns isolated wins into organisational memory. Together they are your programme's early-warning system: the diagnosis behind "diagnose before prescribing".

If external experts run your workflows forever, you bought a service, not a transformation. Every mechanism above exists to make your own people the experts.

GOVERNANCE

Who decides when a team levels up.

Stage progression is a decision, not a drift. Four criteria, applied by named people with clear decision rights — so promotion up the ladder means something.

The four stage-gate criteria

1 • Evidence

A team must show working examples, not intent. One repository, one workflow, one report — running, reviewable, real.

2 • Adoption depth

More than one person on the team must be using the practice. Single-champion adoption does not pass.

3 • Reliability

The practice survives a normal week without rework. Brittle wins are not wins.

4 • Reversibility

You can roll the practice back without disruption if it stops working. If you can't reverse it, you haven't governed it.

Decision rights

ROLE	AUTHORITY
Stream owners	Approve stage progression for any team in their stream. Final call on go / no-go. Resolve cross-stream dependencies and adjust criteria if reality demands it.
Metrics owner	Independently confirms whether the evidence meets the criteria. A deliberate check on stream owners' optimism.
Executive sponsor	Approves the methodology and holds the programme to outcomes. Reviews at 90 days: continue, adjust, or escalate.

In regulated or safety-critical environments, add a fifth criterion — an explicit safety / risk review — and gate Stream B behind read-only pilots before anything touches production. The ladder flexes; the discipline doesn't.

MEASUREMENT

Leading indicators and key results.

Two questions, two sets of numbers. Is adoption happening? Is it providing value? Calibrate targets stream by stream after baselining — never before.

LEADING INDICATORS — IS IT HAPPENING?

- Active AI users per team, week over week
- Knowledge-share session attendance
- Champion engagement: sessions delivered, teams visited
- Workflow maturity: shared configs, prompt & context libraries in place
- Themes from team-lead 1:1s — blockers, wins, requests

KEY RESULTS — IS IT WORKING?

- Cycle time on tracked workstreams
- Throughput: work shipped — PRs merged, tickets closed, content published
- Defect / rollback rate on AI-assisted output
- Hours saved on recurring workflows (self-reported, sampled)
- Quality of output holding or improving under higher pace

THE NON-NEGOTIABLE

Measure flow. Never people.

Use objective, team-level flow metrics to improve delivery — never surveillance-style tracking of individuals. This is a matter of ethics and effectiveness: the moment metrics become individual scorecards, your data corrupts and your adoption stalls. People game what watches them; teams improve what they own.

One dashboard, one owner, one version of the truth. Everyone reads the same numbers — even when they argue about what the numbers mean.

THE FIRST 90 DAYS

Visible wins, real foundations.

Ninety days is long enough to prove the model and short enough to hold attention. Here is what "on track" looks like at each checkpoint.

30

DAYS — SET THE FLOOR

- Usage policy live; safe, approved access for everyone
- Baseline metrics captured before anything changes
- Champions named in every stream
- First knowledge-share forum held
- Stage 1 confirmed across all streams

60

DAYS — FIRST QUICK WINS

- Two team-level workflows piloted in Stream A
- One Stream B pilot underway — scoped read-only first
- Stream C teams shipping AI-assisted work through shared templates
- Workshops and pair working running on live backlogs
- First 1:1 themes feeding back into the plan

90

DAYS — COMPOUND GAINS

- At least one team reaches Stage 3
- Stage-gate criteria validated against real cases
- Measurement framework adjusted from 60-day learnings
- At least one key result improved significantly for two or more teams — e.g. cycle time halved
- Leadership review: continue, adjust, or escalate

Keeping three streams one programme

- **A cross-stream forum, every two weeks.** Best practice and working use cases move between streams instead of being reinvented three times.
- **A shared playbook library.** Any practice that reaches Stage 3 in one stream is published with enough context for another stream to assess fit.
- **A single source of metrics.** One owner, one dashboard, one version of the truth.
- **One brief, every month.** A single page from the programme leads: what changed, what we learned, what we try next.

FAILURE MODES

Five ways this goes wrong — and the early signals.

Every stalled AI programme we've been called into failed in one of five ways. Each has a tell you can spot in the first quarter.

PITFALL	WHAT IT LOOKS LIKE	THE EARLY SIGNAL
Tool-first thinking	Licences bought before workflows are understood; adoption measured in seats, not changed work.	The programme's main artefact is a procurement sheet, not a playbook.
Pilot purgatory	Impressive demos that never touch production work. Twelve pilots, zero standard practices.	No pilot has a named owner or a date to become the default way of working.
Single-champion adoption	One enthusiast per team carries everything; progress evaporates when they're on holiday.	Usage data shows one heavy user per team and a long tail of near-zero.
Skipping stages	Automation deployed on top of habits that don't exist; brittle wins that collapse under load.	A team "at Stage 3" can't show a Stage 2 practice that survived a normal week.
Vanity metrics	Dashboards full of active users and prompts sent; nothing about cycle time, throughput or quality.	Nobody can answer "what business result improved?" without changing the subject.

The common thread

All five are governance failures, not technology failures. The model in this guide — streams, gates, evidence, named owners — exists precisely to make each of these visible early, while they're still cheap to fix.

The honest question

If your current AI initiative stopped tomorrow, would any team notice within a week? If the answer is no, you have a pilot programme wearing a transformation's clothes — and it's worth finding out why.

PROOF

This model, in production.

A three-person ICAN team ran this exact programme for a 120-person blockchain organisation — engineering and business functions alike. Strategy, roadmap, tooling, training, pair working: the full journey. Where internal teams had the capability, we set direction and stepped back; where they needed a working example, we shipped it ourselves.

100%

 AI ADOPTION ACROSS
 THE ORGANISATION

+30%

 PRODUCTIVITY INCREASE
 IN SIX MONTHS

70%

 OF TEAMS AT AAMM
 STAGE 3 IN SIX MONTHS

WHERE ICAN COMES IN

AI Adoption Programme

The enterprise-wide programme in this guide — streams, gates, training, workshops and pair working — delivered by senior practitioners, not slideware.

AAMM assessment

An objective read of where every team sits on the maturity ladder today, and an evidence-based plan mapped to the stage gates.

Agentic AI pipelines

We train your teams to build and operate agentic pipelines that automate delivery and operations end-to-end — to your security and governance standards.

Book an AI readiness conversation.

Tell us where you are today. We'll bring an objective view of your AI maturity and a pragmatic plan to move forward. No obligation, no jargon.

INFO@ICANGROUP.CO.UK

+44 7470 473113

WWW.ICANGROUP.CO.UK

ICAN Consultancy is part of ICAN Group — builders at heart. Offices in London (Group HQ, One Pancras Square) and Istanbul (Engineering). Find us on LinkedIn: [linkedin.com/company/ican-consultancy](https://www.linkedin.com/company/ican-consultancy). © 2026 ICAN Group. You are welcome to share this guide in full, with attribution.