

THE EXECUTIVE GUIDE TO AI IMPLEMENTATION

From Overwhelmed to Operational

How leaders cut through the AI noise, find where value is actually trapped in their business, and get agents into production that earn against a number agreed before the work starts.

Rising AI Group
risingaigroup.com

SECTION ONE

You Are Not Behind

But you will be, if the next ninety days look like the last ninety.

If you are reading this, you are almost certainly experiencing one of the following.

- You know AI is changing your industry, and you feel the pressure to act.
- You have tried to get started and found the landscape confusing, contradictory and exhausting.
- You have watched competitors move faster and quietly wondered whether you are falling behind.
- You have had internal conversations about AI that went nowhere. No clear owner, no clear starting point, no clear return.
- You have been approached by vendors selling AI tools you cannot evaluate or connect to a business outcome.
- You have approved pilots that stalled, consultants who produced reports nobody acted on, or technology your teams simply did not adopt.

> *"Every week I receive three more AI pitches. Every week I feel less certain about which ones matter, and more certain that doing nothing is not an option."*
>
> A sentiment shared by almost every chief executive we speak to.

This is not a sign of slow thinking. It is the predictable result of a market that has produced an extraordinary volume of noise — tools, models, frameworks, vendors and promises — with very little practical guidance for operational leaders who need results rather than whitepapers.

The executives succeeding with AI right now are not the ones who understand it best. They are the ones who found a structured way to identify what matters in their specific business, and someone who could implement it without disrupting the operation.

That is what this guide is about.

8	250%	114–455%
production deployments	**average return**	**project-level range**
across four sectors, all running today	*measured against pre-deployment baselines*	*the spread, not the best case*

Every figure in this guide is drawn from a delivered engagement and measured against a baseline captured before the agent went live. Nothing here is a projection, and no market statistic is quoted that we cannot source.

SECTION TWO

Why AI Is Genuinely Hard

And why the standard advice makes it harder.

The dominant narrative around enterprise AI is written for two audiences: very large corporations with dedicated AI teams and multi-year transformation budgets, and technology startups building AI-native products from scratch. Most operating businesses sit in neither category.

You have a functioning, profitable business. You have operational complexity. You have teams with real workloads and limited appetite for disruption. And you have a board that needs to see returns rather than experiments.

Apply the large-enterprise playbook to that context and you get one of three outcomes. All of them are expensive.

The three failure paths

FAILURE PATH ONE

The Platform Trap

You invest in an AI platform — a tool, a licence, a system — before anyone can name the workflow it is for. Six months later the platform is running. Nobody is using it. The vendor blames change management. Your team blames the vendor. The investment is written off.

The mistake: *buying technology before identifying the business problem it solves.*

FAILURE PATH TWO

The Consultant Report

You engage a firm to produce an AI strategy. They interview your team, benchmark your competitors and deliver a thorough document with a roadmap. It is genuinely good. It sits on a shelf, because nobody inside the business knows how to turn slide ninety-four into something running in production on a Tuesday.

The mistake: *confusing a strategy document with an implementation plan.*

FAILURE PATH THREE

The Premature Hire

You recruit a Head of AI or a data science lead — a significant investment in both time and cost. They arrive to no data foundation, no executive alignment and no operational clarity, inheriting accountability without the authority to act on it. Around fourteen months later they leave, and the knowledge leaves with them.

The mistake: *assuming a hire solves what is actually an organisational readiness problem.*

None of these reflects poorly on the executives involved. They are the predictable result of applying the wrong method to a genuinely complex problem.

> *"The question is never whether AI can deliver value in your business. It almost always can. The question is where, and in what order."*

What actually works

The organisations achieving measurable results share three characteristics. They started with a structured diagnosis of where value is trapped, before committing to any technology. They deployed inside the workflows their teams already used, rather than building new ones. And they measured against a baseline captured before anything was built.

This is not a new insight. It is what separates successful technology deployments from failed ones in every domain. What is new is the pace: four weeks to a costed, prioritised plan, and eight to twelve weeks from there to an agent running in production and earning against the number you set.

The five forces behind the three paths

Five things predict whether an AI programme reaches production. The three failure paths above are what happens when one of them is missing — which is why the fix for each is completely different, and why a generic diagnosis is worse than none.

Vision alignment
Whether the organisation agrees on which problem AI is actually for. Weak here, and buying a platform feels like progress — because buying is easier than deciding.

Implementation discipline
Whether decisions reliably become things that run in production. Weak here, and the instinct is to commission a better plan.

Senior ownership
Whether one named person carries the outcome the way a CFO carries cash. Weak here, and the instinct is to hire someone who will. It is the single most reliable predictor we have found.

Tactical flexibility
Whether you can change direction six weeks in, once budget and contract are already committed.

Adaptive governance
Whether an agent can touch a regulated or customer-facing process safely, and whether anyone has been through that approval path before.

Which of the three are you closest to?

Most leaders know within about four seconds. If you would rather have it tested than guessed, the AI Readiness Diagnostic scores your organisation against these five forces and names the failure path your weakest one puts you nearest — along with the first move worth making.

Twelve questions, about four minutes | risingaigroup.com/assessment

SECTION THREE

Where the Value Lives

The three categories where AI delivers the fastest, most measurable return.

Across deployments in banking, retail, regulated manufacturing and shared services, the highest-value opportunities almost always fall into one of three categories. Not because other categories do not matter, but because these three combine volume, measurability and implementation speed.

01 High-Volume Repetitive Intelligence

The category with the fastest payback.

Every operating business has processes where skilled people spend significant time on work that is, at its core, highly repetitive — even when each instance feels unique. Answering the same categories of question. Producing the same kinds of report. Reviewing the same documents against the same criteria.

This is where agents have the most immediate impact. Not because the work is simple — it often requires genuine domain expertise — but because that expertise can be codified, deployed at scale, and made available around the clock without degradation.

Where this applies in practice

- Commercial managers spending hours each week manually pulling performance data
- Compliance teams reviewing specifications against regulatory requirements
- Policy teams answering the same questions from suppliers and internal teams
- Verification teams processing high volumes of documents to a fixed standard

> **Measured result:** in one deployment, 240+ FTE hours reclaimed every week across 80+ commercial managers, worth A$1.4M in annual net benefit at a 367% return. In another, over 30,000 documents a month verified, 90% of them in under ten seconds.

02 Compliance and Risk Intelligence

The category with the highest asymmetric value.

In regulated industries the cost of a compliance failure is not proportional to the cost of prevention. Agents that operate proactively — monitoring, flagging and alerting before issues escalate — turn an existential risk into a managed workflow.

They do not replace the compliance expert. They extend that expert's reach, so nothing falls through a process designed for human capacity.

Where this applies in practice

- Proactive specification review against regulatory requirements
- Contract clause monitoring against agreed obligations
- Supplier onboarding checks against internal standards
- Real-time policy adherence monitoring across operational teams

> **Measured result:** A$1.75M in annual net benefit at a 455% return, a 100% compliance rate, and zero incidents in the fourteen months since go-live. Adoption reached 98%.

03 Commercial and Decision Intelligence

The category with the greatest strategic leverage.

The quality of commercial decisions is limited by the speed and quality of information available to the people making them. In most businesses that information is trapped: in systems that need analyst support to access, in knowledge that lives in senior people's heads, or in external intelligence nobody has time to monitor.

Agents that open up that access — instantly, in plain language, without technical skill — change the speed of the entire commercial operation.

Where this applies in practice

- On-demand commercial analysis for every manager and buyer
- Personalised weekly briefings rather than requests to a data team
- Instant access to historical deal structures, margin data and supplier terms
- Decision support inside lending, pricing and approval workflows

> **Measured result:** analysis cycles three times faster. A 14-person policy team moved from 2.3 days to under an hour per document, at 100% consistency. A business lending decision moved from 25 days to two.

The critical question is not which category matters most in the abstract. It is which specific processes in your organisation hold the greatest concentration of volume, risk or bottleneck. Identifying that rigorously is the foundation of every successful deployment we have delivered. It is also precisely what the Value Creation Sprint does.

SECTION FOUR

The Value Creation Sprint

Four weeks to clarity, a prioritised plan, and a board-ready business case.

The Sprint was designed to answer one question: how do you make a confident, well-informed decision about AI without committing to an expensive multi-year programme before you know what the opportunity actually looks like in your business?

The answer is a rigorous four-week diagnostic. Not a broad consulting engagement, not a vendor pitch cycle, and not another strategy report — a focused assessment ending in five specific deliverables your organisation can act on immediately.

WHAT WE COMMIT TO

"In four weeks, a ranked and value-sized AI pipeline, an execution plan and a governance framework. From there, eight to twelve weeks to the first agent running in production and measured against the baseline we agreed before we started."

The five deliverables

Every Sprint concludes with five concrete outputs. These are not slide decks with aspirational recommendations. They are working documents your leadership team and board can act on.

01

AI Opportunity Pipeline, ranked by value

A map of your highest-impact opportunities across compliance, operations, commercial and customer experience, each ranked by return potential and implementation complexity, giving you a defensible prioritisation your board can approve.

Typical output: 8–15 ranked opportunities, with the top three to five identified for immediate deployment.

02

Value Sizing Model and Measurement Plan

Every opportunity carries a projected value built on your real transaction volumes and salary loadings — time saved, risk reduced, revenue enabled — expressed in the financial language your CFO expects, and paired with the KPI that will prove it.

Typical output: a sized financial model covering every opportunity in the pipeline, each with its own projected annual value.

03

Operating Model Recommendation

A practical recommendation for how AI should be governed, resourced and managed in your organisation — whether an internal centre of excellence makes sense, what capability you need to build, and what ongoing support is required. Grounded in your actual structure rather than a generic framework.

Typical output: a governance model that addresses the ownership gaps, scope creep and unmanaged risk that stall most deployments.

04

90-Day Execution Roadmap

A sequenced implementation plan for your first tranche — which agents to build, in what order, with what dependencies and at what resource requirement. Designed to put something in production within ninety days of the Sprint concluding, not twelve to eighteen months later.

Typical output: a deployment backlog with week-by-week sequencing, named owners and risk mitigation for the first three agents.

05

Executive Business Case Pack

A boardroom-ready package including the financial model, implementation timeline, risk register and strategic rationale — designed to be presented to your board or executive committee without further preparation.

Typical output: a complete board paper positioning AI investment with a clear, costed, time-bound delivery plan.

How the four weeks are structured

Each week has a specific focus, a clear set of activities, and a tangible output that builds on the week before.

Week 1
Diagnosis

- Structured interviews with four to six stakeholders across operations, compliance, commercial and technology
- Process mapping of your ten highest-volume or highest-risk workflows
- Data landscape audit — what exists, where it lives, how accessible it is
- Competitive benchmarking of the AI capability your direct competitors are deploying

Output: a validated map of your operational landscape and a longlist of 15–20 potential opportunities.

Week 2
Prioritisation

- Value sizing of the top ten opportunities using a structured financial model
- Implementation complexity assessed against your current technology and data environment
- Dependency mapping — which opportunities unlock others, and in what sequence
- Stakeholder validation against business strategy and board appetite

Output: a ranked shortlist of five to eight priority opportunities with validated value projections.

Week 3
Architecture

- Agent architecture design for the top three to five opportunities
- Operating model design and governance framework development
- 90-day deployment roadmap with sequencing, owners and resource requirements
- Risk register and mitigation plan for each priority deployment

Output: a complete deployment plan for the first tranche of agents, ready for board approval.

Week 4
Business case

- Financial model finalisation and sensitivity analysis
- Executive Business Case Pack prepared for board presentation
- Readout session with your executive team — findings, recommendations, next steps
- Transition briefing if proceeding to implementation

Output: five complete deliverables, a board presentation, and a clear decision point.

SECTION FIVE

What the Sprint Is Not

Clarity on what you are, and are not, committing to.

Before committing to any engagement, an executive deserves a direct answer to one question: what am I actually signing up for? Here is ours.

The Sprint is

- ✓ A structured, outcome-oriented diagnosis of your specific AI opportunity
- ✓ A practical, actionable deliverable set rather than a theoretical framework
- ✓ A commitment with a defined start date, end date and clear outputs
- ✓ A low-risk entry point — the fee is fixed and the scope is agreed before you commit
- ✓ An honest assessment — we will tell you if AI is not your right priority right now

The Sprint is not

- × A commitment to an ongoing engagement or retainer
- × A vendor pitch for any specific AI technology or platform
- × A strategy report with no implementation path
- × A broad digital transformation programme
- × Contingent on you having clean data or existing AI infrastructure
- × A prerequisite for working with us — though it is the fastest route to clarity

> *"We designed the Sprint to answer one question as efficiently as possible: where should this specific organisation deploy AI first, and what is that worth? Everything else follows from that."*
>
> Benedetta Angrilli, Founder and Managing Director

The investment

The Value Creation Sprint is a fixed-fee engagement. Four weeks, five deliverables, a board-ready business case, and no surprises — the fee is agreed in writing before any work begins, and the scope is calibrated to your organisation during the mapping session.

What the Sprint returns

A Sprint typically identifies several million in pipeline value; delivered engagements have landed between A$400K and A$1.8M in realised annual benefit. What it costs to realise that depends on which opportunities you choose to pursue.

For the organisations that proceed to deployment, the measured outcome across delivered engagements has been a 250% average return at project level, with the range running from 114% to 455%. Those are post-deployment figures measured against a baseline captured beforehand, not projections.

SECTION SIX

Is This Right for You?

The Sprint delivers most for a specific kind of organisation.

The Value Creation Sprint is not designed to be universal. Here is an honest assessment of where it fits best, and where it may not.

The Sprint delivers maximum value when

You have enough operational complexity to generate real opportunities. Multiple functions, meaningful transaction volumes, and processes that several people run every week. Scale matters less than complexity: a 200-person business with heavy compliance obligations often has more trapped value than a larger one without them.

You operate in retail, FMCG, financial services, supply chain or a compliance-intensive environment. These are the domains where we have the deepest implementation experience and the most validated value benchmarks.

You have operational bottlenecks you can describe but have not been able to solve. Manual processes your team works around. Information requests that take days. Compliance reviews that create launch delays. These are precisely the problems agents solve.

You have a board or senior team that needs a rigorous business case before approving investment. The business case deliverable is built for that audience. It speaks in financial language, not technology language.

You are under competitive pressure from AI-native players or faster-moving competitors. The Sprint produces a clear response plan rather than a defensive reaction.

The Sprint may not be the right first step if

You are still building basic digital infrastructure. If your business does not yet have consistent, accessible data on its core operations, the Sprint will identify that gap – but deployment will have to wait on infrastructure first. We will tell you this clearly in week one.

You are looking for a proof of concept on a single, already-identified use case. In that case we can move directly to a focused agent build without the full Sprint. It is a faster and cheaper path when the opportunity is already defined.

You are a startup or early-stage business. The Sprint is calibrated for operational complexity. Earlier-stage organisations usually benefit more from a focused strategy session than a four-week diagnostic.

SECTION SEVEN

Why Rising AI Group

Experience deploying AI in the environments where you actually operate.

There is no shortage of advisers willing to talk about AI transformation. There is a significant shortage of practitioners who have built, deployed and embedded AI agents inside operating businesses, and who can show you what happened afterwards.

Rising AI Group was founded on a specific conviction: the reason most AI implementations fail is not the technology. It is the gap between strategic ambition and operational reality. Closing that gap requires people who understand both.

Our background

Twenty years transforming businesses across four continents, including BCG and Bain in Milan, Accenture in Sydney, and NielsenIQ across European markets.

Our founder built and led the Automation & AI Centre of Excellence inside Woolworths Group, Australia's largest retailer.

Deep experience across retail, FMCG, compliance, commercial operations and financial services.

Agents in production today, not in pilot.

Our approach

We deploy inside the workflows you already run – email, shared inboxes, familiar systems. No new platform, no migration, no behaviour change.

We measure from day one. Every engagement includes a benefit tracking framework agreed before deployment begins.

We take no vendor commission and sell no licences. Each agent uses whichever model does its specific job best, at the lowest cost.

Every agent is human-designed and human-reviewed before production. Nothing runs that we cannot explain, decision by decision.

A$1.75M	240+	92–98%
largest annual net benefit	FTE hours reclaimed weekly	adoption
from a single compliance deployment, at 455%	*in one commercial deployment*	*the range across delivered engagements*

What our clients say

> *"I went from one to two weeks of work to a few hours."*
>
> **Buyer, international retailer — on a single analysis, three weeks after go-live. She does not work in AI and never learned a new system.**

> *"Next level game changer for us. We are very excited by what this means for us moving forward."*
>
> **Senior executive, international business — on what the deployment opened up next, rather than what it fixed.**

Clients are not named, by agreement. Sector, scale and process are described accurately, and every metric is drawn from the measured post-deployment result.

SECTION EIGHT

Your Next Step

One conversation. No obligation. Genuine clarity.

It starts with a thirty-minute mapping session. Bring the one workflow you keep coming back to — the one where you know good people are losing time, but nobody has ever put a number on it. We map where the time actually goes and what it is worth over a year.

There is no sales pressure in this conversation. We have built the business on clients who proceed because the value is undeniable, not because they felt they had to. If the Sprint is not right for you, we will say so and point you toward what is.

Book a 30-minute mapping session

Book your session →

Thirty minutes. No preparation required. No obligation.

We will listen to where you are and what has blocked progress so far.
We will share where we have seen the most value in businesses like yours.
We will give you an honest recommendation — the Sprint, a faster path, or something else entirely.
No deck, no discovery questionnaire, and no proposal afterwards unless you ask for one.

risingaigroup.com

Alternatively

The AI Readiness Diagnostic

Twelve questions, about four minutes. It scores your organisation against the five forces that decide whether an AI programme reaches production — vision alignment, implementation discipline, senior ownership, tactical flexibility and adaptive governance — and tells you which of the three failure paths in Section Two your weakest force puts you closest to.

You get your score and the full breakdown immediately, on screen.

risingaigroup.com/assessment | Free, no obligation

The executives who will lead their industries in five years are not the ones who understand AI best.

They are the ones who started with the right diagnosis, and took the first step while others were still deliberating.

The first step takes thirty minutes.



RISING AI GROUP

risingaigroup.com

Australia · Europe

© 2026 Rising AI Group Pty Ltd. All rights reserved. This document is provided for informational purposes. Rising AI Group reserves the right to modify its service offerings at any time.