PLAYBOOK

Four decisions **that determine AI success**

Scaling AI successfully requires more than the right technology. It requires the right foundations. This playbook explores four critical decisions that help organisations adopt, govern and scale AI with confidence.

advania

The tech company
with people at heart

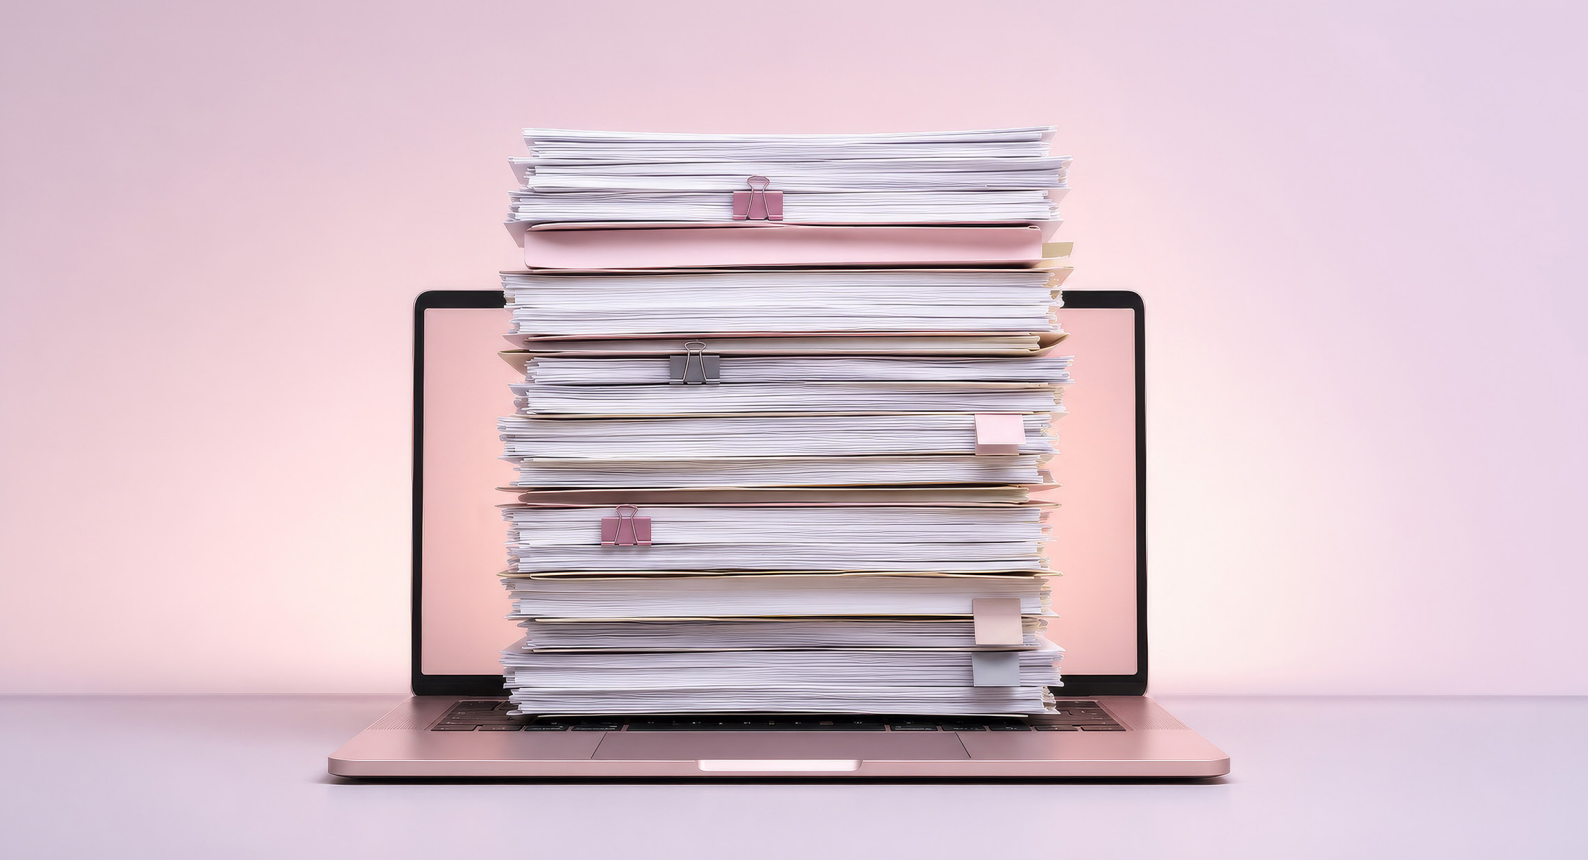

CHAPTER 1 – COMPLIANCE PARALYSIS

Create the conditions for AI adoption

Many AI initiatives start with promising ideas, gain momentum through a pilot – and then stall. Not because the technology falls short, but because no one feels confident enough to say yes. Questions around data, ownership and compliance continue to grow, decisions are postponed, and the initiative remains stuck in limbo. The phenomenon has a name: compliance paralysis.

This is rarely about resistance to AI. Business teams want to create value, security teams want to reduce risk, legal teams want to ensure compliance, and IT wants solutions that can be operated and maintained over time. The problem is that there is no shared framework for how AI should be used. Everyone has the authority to say "wait", but no one has the confidence to say a safe "yes."

The answer is not less governance – it is clearer governance. When the organisation knows where it is safe to experiment, what data can be used and how AI usage is monitored, it becomes possible to learn on a small scale before scaling with confidence.

TWO ZONES INSTEAD OF UNCLEAR YES-OR-NO DECISIONS

The framework can be simpler than you think: two zones. Instead of evaluating every new idea from scratch, the question becomes not "Can we use AI?" but "Which zone does this use case belong in?"

Zone	Typical use cases	Requirements
Safe testing zone	Internal processes, synthetic or approved test data, internal summaries, decision support and learning in a sandbox.	Data classification, audit logging and a clearly assigned owner. The more sensitive the data, the stronger the controls.
Restricted zone	Customer-facing content, branding, sensitive customer data and decisions affecting customers or individuals.	Do not experiment in this zone. AI use requires an approved environment, clear controls and explicit ownership.

The distinction is based on potential impact. In internal processes, using test data and controlled sandboxes, mistakes can be identified and corrected without affecting anyone outside the organisation. That is where people build confidence and experience. In customer interactions, brand communications and external processes, the consequences of mistakes are much harder to reverse. Those situations require established controls rather than experimentation.

READINESS CHECKLIST

Area	Key question	Why it matters
Ownership	Who owns the business value and can approve the pilot?	Without ownership, initiatives become everyone's priority, and no one's responsibility.
Process	Which specific workflow should be tested first?	A clearly defined workflow makes both risk and value easier to evaluate.
Data	What data is required, and how sensitive is it?	Data classification determines what can be tested safely.
Sandbox	Where can non-technical users experiment without risking customer data, brand reputation or production decisions?	AI only becomes a learning journey when people have a safe environment to practise.
Zone	Does this initiative belong in the safe testing zone or the restricted zone?	Clear boundaries make it easier to say yes without losing control.
Governance	What needs to be logged, monitored and approved?	Traceability creates confidence and enables responsible scaling.

SIX QUESTIONS BEFORE SAYING YES

- What problem or business task is AI expected to solve?
- Does this initiative belong in the safe testing zone or the restricted zone?
- What is the potential impact if the output is incorrect, shared with the wrong audience or used in the wrong context?
- What data is required, and can it be used in the intended environment?
- Is there a sandbox where non-technical users can experiment safely?
- What should success be measured by: time saved, quality, risk reduction, cost or user value?

KEY TAKEAWAY:
Compliance paralysis is not solved by less governance, but by clearer guardrails for when and how AI can be used. When the organisation knows where it is safe to experiment, what data can be used and who is responsible for follow-up, it becomes easier to say yes, start learning and scale what works with control.

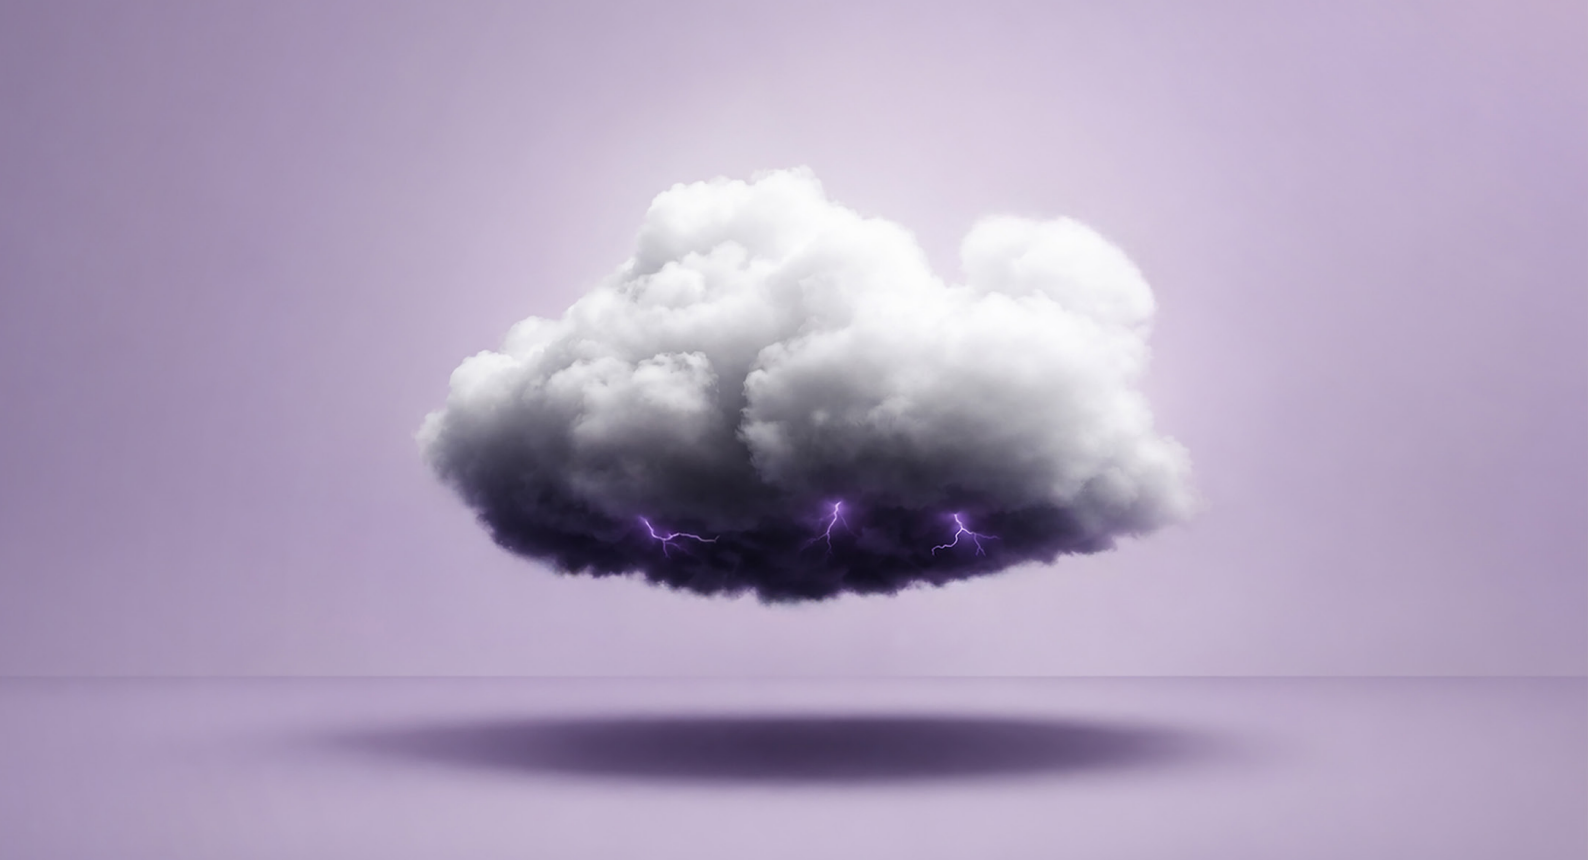

CHAPTER 2 – SHADOW AI

Make the right path **the easiest path**

When employees turn to their own AI tools, it is easy to see the issue as a security problem. In reality, Shadow AI is usually a symptom of something else: the organisation's approved path does not meet business needs. Reducing the risk starts with understanding why Shadow AI emerges, and how to make the approved path the easiest one to follow.

WHEN SHADOW AI BECOMES AN ORGANISATIONAL CHALLENGE

Shadow AI does not appear because a particular tool is irresistible. It appears when employees have a genuine need, but the organisation has not provided an approved, practical way to meet it. Most people have good intentions: summarising a document, drafting a response, analysing a report or understanding complex information. The problem is that sensitive information can leave the organisation without anyone knowing what data was shared, where it went or how the answer was generated.

This becomes especially clear when organisations encourage wider AI adoption without providing the right tools. When people are expected to work with AI but lack the right environment, approved tools and clear boundaries, doing the right thing becomes difficult. The path of least resistance wins.

PRACTICAL EXAMPLES: WHERE SHADOW AI TYPICALLY APPEARS

Area	Employee need	What the organisation should provide
Customer service	Summarise cases or draft responses.	An approved environment for handling customer data, or clear restricted zones.
HR & internal support	Draft policy responses, onboarding content or summaries.	Guidance for personal data and approved internal knowledge sources.
Sales & procurement	Interpret requirements, draft proposals or compare documents.	Support for confidential documents and clear review before sharing.
Analytics & reporting	Summarise data or generate draft reports.	Traceability, source references and clear data classification.

WHAT THE"RIGHT TOOL" SHOULD PROVIDE

An approved AI tool that does not work in everyday situations will not solve Shadow AI. The approved alternative should provide:

- A trusted environment where data is handled according to organisational requirements.
- Clear rules for what data may be used.
- Traceability of usage, sources and outputs.
- Simple access for the people who need AI in their daily work.
- Support for working with documents, knowledge and business processes without exposing sensitive information.
- A safe learning environment where employees can experiment without risking customer data or production decisions.

This is where the principle of "making the right thing easy" becomes essential. If the approved path is complicated, unclear or too limited, people will choose another one. If it genuinely helps them do their jobs better, much of the incentive for Shadow AI disappears.

QUICK ASSESSMENT: IS SHADOW AI A RISK?

- Do employees have access to an approved AI tool or environment that genuinely helps them in their daily work?
- Is the approved option slower, weaker or less intuitive than the alternatives?
- Do employees know what data they may and may not use with AI?
- Is there a safe sandbox where they can learn without risking sensitive data?
- Can you see which AI workflows are being used and what business needs they address?

HOW TO CREATE AN APPROVED PATH

The way out of Shadow AI is not stricter bans but giving people what they actually need. Start with clear guidance for data classification, approved environments and the types of work AI is allowed to support. Make it easy to ask for guidance whenever the boundary is unclear. The goal is to make the right behaviour the easiest behaviour. Shadow AI does not disappear because people stop being curious, it disappears because the approved alternative works better.

KEY TAKEAWAY:
Shadow AI is not solved by more restrictions. It declines when the organisation provides an AI environment that is just as easy to use as the tools people would otherwise choose for themselves. The goal is not to limit AI adoption, but to make it trusted, traceable and scalable.

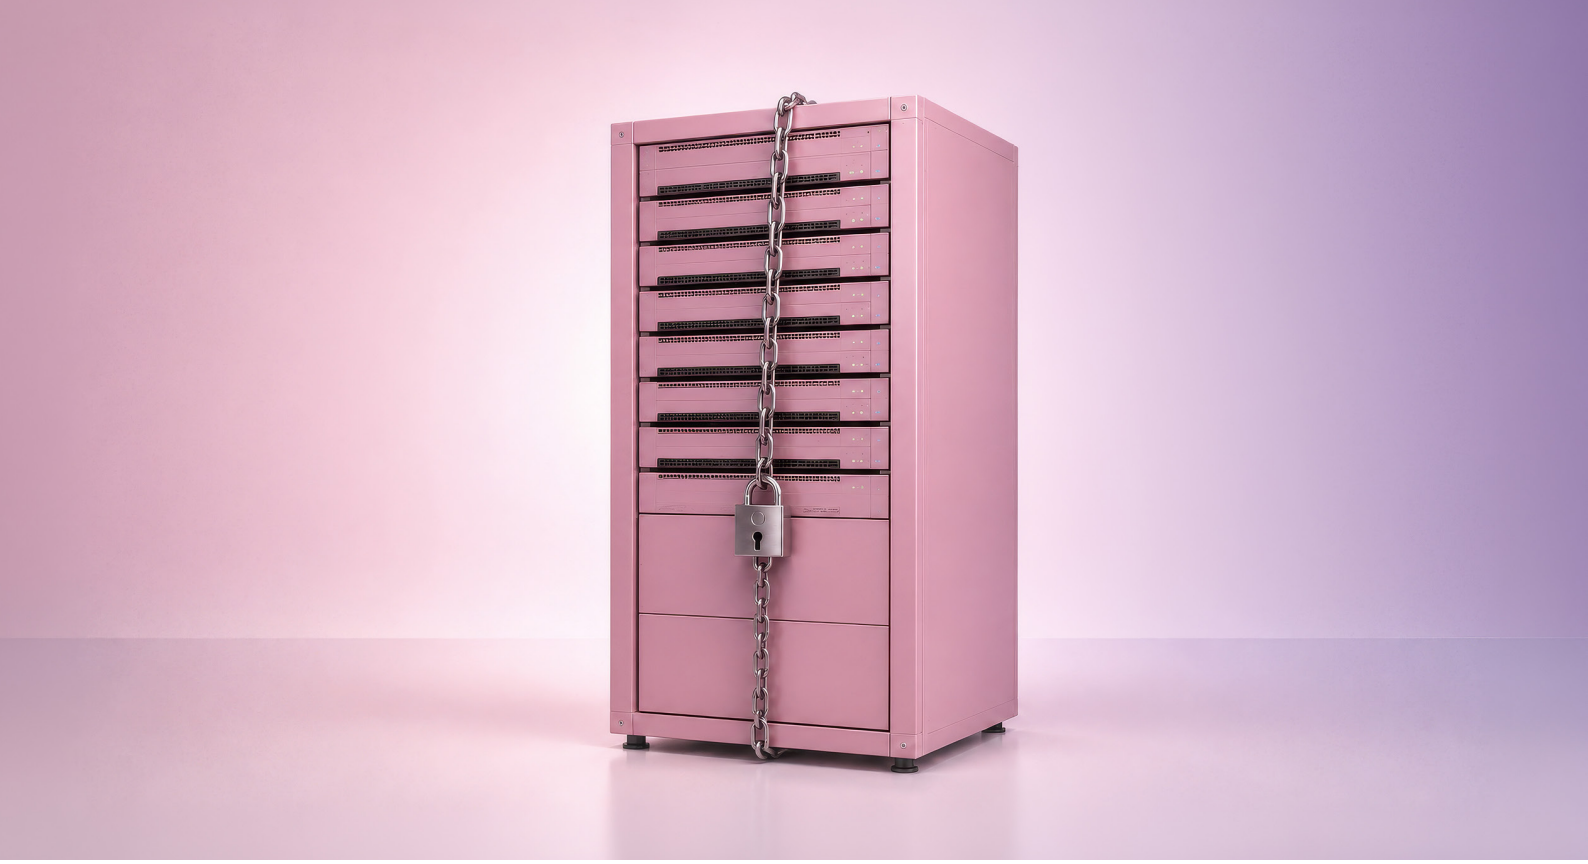

CHAPTER 3 – CAPABILITY LOCKOUT

Stop choosing between **quality and control**

Many organisations believe they have to choose between the most advanced AI models and full control over their data. In reality, that trade-off is often false. As AI solutions mature, the model they require changes too. Success is not about always choosing the biggest model, it is about choosing the right model for the right job.

THE FALSE TRADE-OFF

It is easy to assume that better AI models mean less control over your data. In practice, the opposite is often true. As an AI use case matures, more of the intelligence moves from the model into the solution itself. Smaller models can then deliver the same outcome while giving the organisation greater control over data, cost and operations.

UNDERSTANDING MOVES INTELLIGENCE FROM THE MODEL TO THE SOLUTION

At the beginning of an AI initiative, much is still unknown. What is the real task? What questions will users ask? What data is required, and in what condition? Where are the boundaries?

At this stage, a frontier model is often the right choice, not because the task itself demands it, but because the model's broad capabilities compensate for everything the organisation has not yet learned. That is why frontier models are invaluable during experimentation: they allow teams to iterate quickly, test ideas and discover what the problem really is before building a production-ready solution.

With every iteration, however, something changes. Knowledge sources improve. Prompts become more precise. Business rules are documented. Interfaces to surrounding systems are defined. Step by step, intelligence moves from the model into the solution itself. What the organisation learns becomes embedded in the process instead of being improvised by the model every time it is called.

As the solution moves closer to production, the embedded knowledge does more of the work, and the model has less to figure out on its own. At that point, a smaller model can often deliver the same quality, not because expectations have been lowered, but because the task has become better defined. Quality depends less on the model and more on the solution surrounding it.

CHOOSE THE MODEL BASED ON THE TASK AND THE DATA

Type of work	Example	What should drive the decision?
Exploration and testing	Understanding requirements, validating ideas, identifying workflows	Start with a powerful model to maximise learning.
Low sensitivity, high quality requirements	External content creation, structuring public information, summarising public material	Model quality and cost.
Internal information	Internal procedures, project documentation, meeting notes	Access control, traceability and where data is processed.
Sensitive or protected data	Customer data, personal data, contracts, classified information	Data classification, access control and approved environments.
Repetitive business processes	Case handling, procurement support, reporting	Operational cost, governance and long-term ownership.
Mature production workloads	Clearly defined prompts, knowledge sources and business rules	Evaluate whether a smaller model is sufficient. It often is.

THREE EXAMPLES OF MODEL MATURITY

Use case	During exploration	Moving into production
Talk to your documents (RAG)	A powerful model helps identify question types, quality issues and gaps in the knowledge base.	As metadata, retrieval and instructions improve, smaller models often become sufficient.
Customer service assistant	A capable model helps discover customer behaviour, escalation paths and operational risks.	Once workflows and approved responses are defined, smaller models can handle most requests.
Agentic workflows	Frontier models are valuable while agents reason, plan and navigate ambiguous tasks.	As workflows become modular and specialised, smaller models can take over more steps.

CHECKLIST BEFORE SELECTING A MODEL

- What data will the model need to process, and how sensitive is it?
- How well is the task understood today? Are you still exploring, or is the workflow well defined?
- Which parts of the process are repetitive enough to become narrow, well-defined tasks?
- Does the output need to be traceable to its sources?
- Is this a one-off task or a recurring business process?
- Is the solution still being explored, or is it ready for production?

A COMMON CHALLENGE

Many organisations start by comparing models instead of understanding the task. The discussion quickly becomes technical: Which model is the fastest? The smartest? The most capable? A better starting point is to ask what the use case actually requires, and how well it is already understood.

The goal is not to start with the most expensive model. If the task is already well understood, a smaller model may be the better choice from the outset. During experimentation, powerful models accelerate learning. As the solution matures, that knowledge should become part of the solution itself. A model that makes sense in a proof of concept may become unnecessarily expensive when the same workflow runs thousands of times every month. At that stage, a carefully chosen smaller model can often deliver the same outcome at a fraction of the cost.

KEY TAKEAWAY:
Choosing an AI model is not about selecting the most advanced option. It is about understanding the task, the data and the maturity of the solution. As AI solutions evolve, intelligence shifts from the model into the process, making it possible to combine high quality with stronger control and lower cost.

CHAPTER 4 – INVISIBLE AI SPEND

Make AI visible and governable

As AI adoption scales, the key question is no longer just what AI costs, but what value it creates. To make better investment decisions, organisations need visibility into usage, costs and outcomes.

WHAT NEEDS TO BECOME VISIBLE?

Metric	Question	Why it matters
Team	Who is using AI?	Makes AI costs visible and accountable.
Use case	What business task is being solved?	Connects AI usage to business value.
Model or service	Which AI is being used?	Shows whether the right tool is being used for the right task.
Cost	What does the workflow cost per task or reporting period?	Supports budgeting and prioritisation.
Outcome	What value is being created?	Helps determine whether an initiative should be scaled, improved or stopped.

A NEW WAY TO THINK ABOUT KPIS

AI evolves quickly. That means organisations need flexible measurement as well as flexible technology. A modern AI programme should be able to introduce, adjust and retire KPIs as AI adoption matures.

A SIMPLE WAY TO ESTIMATE VALUE (ROI LITE)

You do not need a complex business case to get started. Begin with three questions:

1. **What value should AI create?**
   This could mean shorter processing times, higher-quality decision support, reduced operational risk, or improved service for customers and employees.
2. **What does the solution cost?**
   Consider both implementation and operations, including data preparation, integrations, licences, infrastructure, change management, training and ongoing monitoring.
3. **What assumptions need to be tested?**
   Most AI initiatives are based on assumptions: the value they will create, how often they will be used and whether they actually improve the work being done. Validate those assumptions at a small scale before expanding, and be prepared to adjust your KPIs as you learn.

SIX QUESTIONS BEFORE SCALING

- Is it clear which team and use case each AI cost belongs to?
- Can you measure cost per task or workflow?
- Are you measuring both consumption and business value?
- Is there clear ownership for monitoring progress over time?
- Is it clear which AI workflows should be scaled, paused or improved?
- How quickly can KPIs evolve as models, workflows and AI usage change?

KEY TAKEAWAY:
AI does not become valuable simply because it is widely used. It becomes valuable when it is used intentionally. When organisations can connect AI usage, costs and outcomes, AI becomes an investment that can be prioritised, measured and continuously improved. Visibility is therefore not about limiting AI, it is about creating the confidence to keep scaling.

Next steps

These four decisions come together to create the foundation for successful AI adoption. A trusted foundation gives people the confidence to start using AI. Clear governance helps them use the right tools. A sovereign architecture removes the false choice between quality and control. And when AI usage becomes visible, organisations can make better decisions about what to scale, improve or change. **That is how AI moves from isolated initiatives to a long-term organisational capability.**

Would you like to discuss what this could look like for your organisation? Talk to our AI experts about your next steps.

TALK TO AN AI EXPERT

advania

The tech company *with people at heart*