MOXIE'S AGENTIC AI DISCOVERY APPROACH

Moxie

Deciding Where Agentic AI Actually Pays Off

A structured, technology-agnostic, evidence-based way to decide which of your use cases deserve an AI system that can plan and act on its own, before you commit budget to building one.

What is agentic AI for you?

Less autonomy → More autonomy

1 Assistant

AI drafts a recommendation. A person carries out every action.

REAL-WORLD ANALOGY
GPS navigation

2 Approve each step

AI proposes one action at a time. A person signs off before anything happens.

REAL-WORLD ANALOGY
Adaptive cruise control

3 Act, then review

AI acts within an agreed scope. A person checks the work afterward.

REAL-WORLD ANALOGY
Lane-centering assist

4 Autonomous with exceptions

AI completes the task on its own, only pausing for defined high-risk cases.

REAL-WORLD ANALOGY
Hands-off highway autopilot

5 Fully autonomous

No human checkpoint at all. Rare in practice, and high-risk.

REAL-WORLD ANALOGY
Driverless robotaxi

Discovery decides the right level for each action, not the highest one by default.

THE CHALLENGE

Why most agentic AI projects stall

AI is moving faster than most teams can absorb

New models, tools, and promises of transformation arrive weekly, faster than most teams can learn and adopt without dropping everything else.

The result: many IT leaders are stuck between urgency and uncertainty, unwilling to commit before the landscape settles and a clear path forward emerges.

Enthusiasm for AI is outrunning proof that it works

40%+

of agentic AI projects are expected to be shut down by 2027

Source: Gartner

Gartner points to three recurring causes: **costs that run higher than expected, unclear business value, and risk controls that were never designed in.**

None of these are technology problems. They are decision-making problems.

Not every agent should get the same freedom

Every agentic AI deployment sits somewhere on a scale of how much a person needs to stay involved. Left undefined, autonomy defaults to whatever is easiest to build, usually higher than intended, not what the risk actually calls for.

A workflow that reads and drafts is a different bet than one that acts and spends on its own, and the two should never be approved the same way.

THE RISK

Autonomy isn't a default setting.

Every action an agent can take needs its own deliberate autonomy level and a named approver, decided during discovery, not left to chance.

Both problems are solved the same way: **a staged, evidence-based discovery process.**

A staged, evidence-based path from idea to costed plan

Run the workshop with an optional ideation step if you don't have named use cases yet, or start directly at Step 1 if you do. Only the use cases that earn a Go move into a separate, scoped Detailed Discovery stage.

No use cases identified yet? Stage 1 can start with an optional Step 0 ideation session to surface and gate candidate ideas first.

Stage 2 is a separate, scoped stage. You invest only in the use cases that earned a Go in Stage 1.

Every step above has a defined output. **The next page breaks down all nine.**

EVERY STEP, AND WHAT IT PRODUCES

Moxie

Stage 1 is a short, focused commitment. It tells you, with evidence, which of your use cases deserve the deeper investment of Stage 2, and which don't.

STAGE 1 The Workshop

3 STEPS

Step 0 (optional): Ideation. → a gated shortlist of named use cases

Step 1: Business goal. → problem statement and success metrics

Step 2: Agent fit test. → a Go, Hold, or Redesign call

STAGE 2 Detailed Discovery

6 STEPS

Step 3: Map the workflow. → end-to-end process map

Step 4: Inventory tools & data. → tool/action register and data-access map

Step 5: Decide autonomy. → autonomy map with named approvers

Step 6: Guardrails & governance. → threat model and guardrail requirements

Step 7: Model the cost. → per-task cost model and break-even volume

Step 8: Plan for failure. → evaluation plan and rollback procedures

Ready to find out where agentic AI actually pays off for you? Let's talk and scope your Stage 1 workshop.