

THE STATE OF AI IN 2026 · PAPER 5 OF 5

# A practical playbook for SMBs getting started with AI.

Five steps, real metrics, and a framework built on what the research actually says works – not what sounds good in a pitch deck.



Change by Design

PRACTICAL AI. PRACTICED HAND.  
CHANGEBYDESIGN.AI

# Good business + AI = results. Broken business + AI = faster failure.

That's the single most important sentence in this entire five-paper series, and it's the one most AI conversations skip past.

“AI amplifies what you already have. If your processes are clear and your team knows what success looks like, AI will accelerate all of that. If your processes are murky and nobody's defined what a good outcome looks like, AI will scale that mess faster than any human ever could.”

Enterprise generative AI spending hit \$37 billion in 2025 – a 3.2x jump from 2024. The money is moving. The question for SMBs is how to deploy it with discipline, not just enthusiasm.

## The five-step implementation path

This framework is built directly from what the research says works. Not theory – documented patterns from the organisations that beat the failure statistics covered in Papers 1 through 4.

### Step 1 · Baseline your KPIs

Before you touch any AI tool, measure what you're trying to improve. How long does your team spend on customer enquiries today? What's your average response time? How many hours go into monthly reporting?

Most SMBs skip this step because it feels like overhead. It's not. It's the only thing that lets you prove (or disprove) that AI actually moved the needle. Without a baseline, you're guessing – and guessing is how organisations end up in the failure statistics.

### Step 2 · Pick one high-volume process

The research is unambiguous: specificity beats ambition. Pick the one process that's high-frequency, repetitive, and where failure is recoverable. A chatbot that handles routine customer questions is a recoverable experiment. An AI agent making credit decisions without human review is not. Start with the first kind. Earn the right to do the second.

### Step 3 · Run a 30-day pilot

Thirty days is enough to see whether something works without over-committing. Set clear success criteria before you start: a 20% reduction in response time? Handling 50% of routine enquiries without human intervention? Be specific. During the pilot, keep a human in the loop for every output.



ACTIVE TEAM · RUNNING THE PILOT

## Step 4 · Measure and compare

Pull your baseline numbers. Pull your pilot numbers. Compare them. Be honest about what changed and what didn't. AES measured audit cost, audit time, and audit accuracy. Deutsche Bank measured research turnaround. BBVA measured hours saved per employee per week. Follow the same pattern.

## Step 5 · Scale or pivot

If the pilot proved value, expand — but with the same discipline. Add monitoring. Define who's accountable. Build the governance before you need it, not after something goes wrong. If the pilot didn't prove value, diagnose why. A failed pilot that cost you 30 days and taught you something is not a failure. A pilot that "succeeded" but never produced a number is.

# 40–60 min/day

saved per enterprise user (OpenAI 2025). Even the conservative end is worth pricing out against your team's fully-loaded cost.

03

## Six things to avoid

### Key takeaways

- Don't skip the baseline. Without it, you can't prove anything worked.
- Don't try to transform everything at once. That's the failure pattern from Paper 1.
- Don't treat AI outputs as decisions. They're inputs. A human makes the call.
- Don't deploy on bad data. AI on poor data produces poor results at scale. Fix the data first.
- Don't confuse a good demo with a production system. The gap is where budgets disappear.
- Don't assume vendor metrics are independent evidence. Read them carefully. Verify with your own numbers.

04

## The decision in front of you

Across 34 research reports, over 1,000 documented deployments, and findings from McKinsey, Google, Anthropic, BCG, IBM, Deloitte, the OECD, and more, the evidence resolves into one conclusion: organisations that build disciplined, measured, human-centred AI capabilities now will create advantages that compound. Organisations that keep running pilots without the infrastructure to scale them will keep contributing to the failure statistics.

The organisations beating the odds aren't exceptional. They're disciplined. They start specific. They measure before and after. They keep humans in the loop. They build governance before they need it.


“The question isn't whether AI is real. It's whether your organisation is ready to use it well.”



# Let's talk.

Ready to start? We'll help you scope your first pilot.

Change by Design  
HELLO@CHANGEBYDESIGN.AI  
CHANGEBYDESIGN.AI



---

# Change by Design

AI strategy & automation for growing businesses.

Change by Design helps SMBs and enterprise teams turn AI from a buzzword into a business advantage. We don't sell software. We build the strategy, governance, and implementation discipline that makes AI work in the real world.



**Change by Design**

PRACTICAL AI. PRACTICED HAND.