# From Mandate to Momentum

## How to Deliver AI Projects Without Blowing the Budget

An executive guide for leaders who need real AI results, without runaway costs.

Presented by

RŌNIN CONSULTING

# Table of **contents**

# AI is no longer "nice to have."

It's a race between companies that pull ahead and those that slip behind.

The catch?

Early adopters are already encountering costly missteps and unfulfilled promises. As time has passed, it has become clear that the road to implementing AI isn't paved in gold; it's littered with stalled pilots, scope creep detours and costly breakdowns.

But it doesn't have to be this way. There are ways to implement AI projects successfully, and this guide can help you achieve that. Inside, we'll show you what keeps projects moving forward, what sends them off course, and how to steer your business toward real results without blowing the budget.

01

# Eight AI red flags

AI projects are more complex than the boardroom pitch decks make them look. The learning curve is costly, and early AI adopters have begun to find this out the hard way.

The good news?

Most of the failures we are seeing coming from these early projects share the same set of root causes, which means they're avoidable. In the pages ahead, we'll break down eight common red flags that derail AI projects and share field-tested counter-moves you can put into practice right away:

- **fuzzy goals**
- **tech in search of a problem**
- **boiling the ocean**
- **data neglect**
- **skill gaps & wishful thinking**
- **siloed teams & misalignment**
- **no change management**
- **ignoring governance & risk**

# 1 Fuzzy goals

**"We'll know it when we see it."**

**Why it burns budget**:
Without a single, measurable business outcome, you'll chase features rather than results. The scope will never stay contained, and success will be subjective.

**Smart strategy: define success early.**

- Shoot for one business KPI and one technical KPI to lock in the project scope. Example: Reduce average support response time from 2 hours to 30 minutes by Q4 (Business KPI); achieve <2% hallucination rate on retrieval-augmented answers under load. (Technical KPI).

- Apply SMART criteria (Specific, Measurable, Achievable, Relevant, Time-bound). Clarity up front is a huge budget saver.

- Agree on a baseline and measurement plan before building anything.

**Pro tip:**
Start with a one-page "outcome charter" signed by business and tech sponsors that will not change without a formal scope review.

# 2 Tech in search of a problem

**"Let's try your vision first."**

**Why it burns budget:**
Tools become the center of gravity; you end up solving an interesting problem instead of the right one.

**Smart strategy: start with the pain points.**

- Ask first: Where are we bleeding money? Where are customers frustrated? Where are we missing revenue?
- Score your use cases on a combination of value, data availability, complexity, and risk.
- Choose the smallest problem with the highest net impact, and the tech will follow.

**Pro tip:**
Start with neutrality and evaluate multiple approaches (rules, analytics, retrieval-augmented generation, and fine-tuning) against the same business test.

# 3 Boiling the ocean

**"We can solve everything in this single project."**

**Why it burns budget**:
Complex projects grow non-linearly. Every new feature or integration point multiplies complexity. The more you try to cover up front, the faster you burn time and money on rework.

**Smart strategy: Start small, prove value, then begin to scale.**

- Pick a slice of your vision. Choose a quick win that can deliver visible value in a few months.
- Work in short cycles. Deploy → measure → refine. Momentum builds faster with quick iterations.
- Use proof of concept the right way. A POC isn't about perfection; it's about validating assumptions. Treat your POC as a learning lab: celebrate what works, and use failures to adjust your approach before scaling it up.

**Pro tip:**
Anchor your POC around a single workflow that has clear success metrics. This measurable scope makes it easier to win buy-in for the next phase.

# 4 Data neglect

**"We have lots of data, so we're fine."**

**Why it burns budget**:
Late discovery of complicated data issues causes redesigns, security delays, and potential model under-performance.

**Smart strategy: Focus on your data first.**

- Conduct a small data audit before you begin any project. Identify the required sources, gaps, and latency needs.
- Allocate 50–70% of your effort to data extraction, cleaning, integration, and governance.
- Involve security and compliance from day one; they should always be involved in the process.

**Pro tip:**
More data isn't always better. What matters is whether it's clean, relevant, and usable. A few well-curated sources will outperform terabytes of messy, duplicate, or non-compliant records every time.

# 5 Skills gaps and wishful thinking

**“Let’s rely on one AI architect."**

**Why it burns budget:**
Loading every responsibility onto one person creates bottlenecks, and the gaps seem to surface at the worst, and most expensive, moments.

**Smart strategy: Conduct an honest skills gap analysis.**

- Map out what technology your project actually needs: data engineering, ML modeling, MLOps, integration, etc.
- Identify what you have in-house vs. what you need for this project.
- Fill the gaps intentionally: consider time-boxed contractors, strategic consultants, or a delivery partner with experience in your industry.

**Pro tip:**
Don't set someone up to fail by making them the "AI person." Beyond the technical blind spots, one-person strategies almost always lead to burnout. Build a balanced team so responsibilities are shared and no one burns out.

# 6 Siloed teams & mismangement

**“It's an IT project; the business will see it when it's done."**

**Why it burns budget:**
You might build the wrong thing beautifully. Then spend cycles rebuilding it for buy-in or adoption.

**Smart Strategy: Align early and align often.**

- Make your project cross-functional from day one. Involve key stakeholders across departments to ensure everyone agrees on what success looks like.
- Hold joint demos every sprint; revisit your outcome charter to stay within scope.
- Set dual success metrics (business + technical) and require sign-off at each stage gate.

**Pro tip**:
Consider your AI projects as co-created products, not IT handoffs. The more your business and technical teams collaborate early and often, the less you'll spend on rewrites, and the faster adoption will come.

# 7 No change management

**"If we build it, they will use it."**

**Why it burns budget:**
AI adoption doesn't happen automatically. Without training and clear incentives, employees are liable to stick to old habits.

**Smart strategy: Champion user adoption.**

- During the process, keep the users informed and clearly define their responsibilities.
- Identify "change champions" – respected folks on the front lines who are excited about the AI and can advocate to their peers.
- Provide role-specific training and on-ramp playbooks; celebrate wins.
- Establish clear escalation paths when the AI is not received well or fails.

**Pro tip:**
Adoption is as emotional as it is technical. Acknowledge the anxiety AI can create, and frame AI training as empowerment, not a replacement. People buy into tools that make them feel more capable, not less secure.

# 8 Ignoring governance & risk

**"We'll figure out privacy, bias, and model risk later."**

**Why it burns budget:**
Waiting until deployment to address governance guarantees late-stage surprises. Security and legal teams will halt launches, regulators will raise flags, and user trust will evaporate.

**Smart strategy: Build in guardrails.**

- Bring legal and security stakeholders into the early planning process, not post-production. They can define policies, data usage boundaries, and offer third-party review requirements.
- Clarify what happens when the AI is incorrect and who is responsible for the response.
- Track for bias, hallucinations, drift, and privacy leaks, and establish a process to act on them quickly.

**Pro tip:**
Governance can be an accelerator instead of a blocker. When you build trust from the start, your approvals move faster, your adoption rate climbs, and your team avoids the costly missteps that could come from ignoring risk until the end.

02

# AI wins from the field

Winning with AI isn't about chasing the flashiest tools or the boldest promises; it's about turning innovation into real outcomes. Companies that succeed aren't just experimenting with AI; they're applying it with focus, discipline, and a clear eye on business value.

What follows are two examples from the field. These examples show how organizations moved past the hype, navigated the pitfalls, and came out ahead with measurable wins. They're proof that AI success is possible and a preview of what's ahead for businesses ready to make AI work for them.

# Wins from the field

## AI saves thousands of physician hours

When a large healthcare group, The Permanente Medical Group (TPMG), rolled out ambient AI (1), its goal was to decrease the documentation burden, but the results were much more impactful than that. After a single year in, the data revealed:

- Physicians saved 15,700 hours of documentation in one year (about 1,794 workdays). Each appointment cut about 1 minute of after-hours charting.
- 84% of surveyed physicians said it improved patient connection, and 82% felt greater job satisfaction as a result of reduced documentation burden.
- Most patients reported the AI scribe had a neutral or positive impact on care quality.


Credit: Permanente Medical Group

## WHY THIS PRODUCT WORKED

**Chose the right pain point**: By targeting documentation, a universal, high-friction task, AI scribes reduced administrative burden and freed up clinicians for more patient interaction.

**Integrated with workflows:** Rather than disrupting how physicians worked, the tool slotted into existing routines, making adoption seamless.

**Scaled thoughtfully:** A slow, measured rollout ensured reliability and sustained satisfaction, allowing the program to grow without losing momentum.

# Wins from the field

## Generative AI boosts wealth management productivity

When Morgan Stanley introduced its GPT-powered assistant for wealth management advisors, Debrief, the goal was to reduce administrative overhead and enable advisors to focus more on clients. Within months, adoption and impact exceeded expectations. By mid-2024, the AI assistant had been adopted by 98% of advisor teams (2), making it one of the fastest enterprise rollouts in the firm's history. The results speak for themselves:

Credit: Morgan Stanley home office

- 98% adoption across advisor teams.
- Document access increased from 20% to 80%, reducing search time and enhancing efficiency.
- Follow-ups that once took days now happen within hours, thanks to Debrief auto-summaries.
- Advisors reported stronger client engagement and higher satisfaction, along with reduced administrative work.

## WHY THIS PRODUCT WORKED

**Chose the right pain point:** By targeting high-friction tasks like information retrieval and meeting notes, AI freed advisors for more valuable client engagement.

**Integrated with workflows:** The tool was embedded directly into the knowledge base. Users could integrate it from meetings directly into Salesforce, which eliminated learning curves and encourages daily use.

**Scaled into a platform:** The success of Debrief created a foundation for future "super app" style AI services across the firm (3).

03

# AI losses from the field

Not every AI story ends in a win.

For every successful deployment, there are projects that collapse under the weight of their own ambition.

The truth is, AI failures aren't just about bad technology, they often stem from misaligned goals, rushed execution, or scaling before the foundation is ready. These two examples show how even well-funded, headline-grabbing AI initiatives can stumble and what lessons businesses can take to avoid repeating the same mistakes.

# When AI projects fail

## AI ambition outpaces reliability

Forward Health, a startup founded by former Google and Uber executives, set out to disrupt primary care with AI-enabled "CarePods." (4). These futuristic, kiosk-like stations promised a new model of healthcare delivery: patients could step inside, have vitals measured by AI, and receive a diagnosis in minutes. Backed by more than $650 million in funding, the company scaled quickly and aimed to deploy 3,200 CarePods nationwide.

But by November 2024, Forward abruptly shut down all its clinics and canceled patient appointments. They had built 19 tech-enabled primary care clinics across the U.S. and rolled out AI-based CarePods that automated health check-ups in several cities (5).

Credit: Forward

### WHY THIS PRODUCT FAILED

**Chose the wrong pain point:** Instead of tackling a clearly defined, high-friction problem, CarePods tried to reinvent the entire clinical experience, an overly ambitious target prone to failure.

**Did not integrate with workflows:** Rather than easing clinicians' workload, unreliable AI outputs added extra steps and slowed care delivery.

**Scaled too quickly without validation:** Forward raced to expand without proving reliability at a small scale first, leading to wasted investment.

# When AI projects fail

## AI inspiration turns harmful

Pak'nSave, a major New Zealand supermarket chain, launched an AI-powered meal-planning app in mid-2023 designed to help customers use up leftovers by entering ingredients and generating creative recipes. The app, nicknamed Savey Meal-bot, was powered by OpenAI's GPT-3.5 model. Initially, it's an engaging tool, but then things took a dangerous turn (6).

In August 2023, users discovered the app suggesting truly horrifying "recipes" when non-food home items were entered into the app. These included a concoction dubbed "aromatic water mix," which is chemically equivalent to chlorine gas. Yet, the app described it as "the perfect non-alcoholic beverage to quench your thirst and refresh your senses".


Credit: Pak'nSave

## WHY THIS PRODUCT FAILED

**Chose a volatile input domain:** Allowing users to type arbitrary ingredients, including hazardous substances, led to the AI's misbehavior. The interface blurred the line between food items and dangerous chemicals.

**Neglected critical safety controls:** The tool lacked necessary content filtering and vetting layers to block unsafe outputs, a fatal oversight for public use.

**Underestimated user testing:** The app was launched without sufficient trials for edge cases or malicious use, allowing dangerous outputs to slip into release.

04

# Your AI playbook

We've looked at both sides of the AI journey: the red flags that derail projects, and real-world examples of initiatives that have succeeded or fallen short.

What becomes clear is that success isn't about luck or limitless budgets. It's about recognizing the warning signs early, learning from the missteps of others, and applying proven strategies to keep projects on track.

This playbook distills these lessons from the field into three rules. Think of them as your guardrails: focus on what matters first, build on a strong foundation, and design for long-term adoption.

Follow these rules , and you'll give your AI projects the best chance to deliver results that last.

# Three rules for AI success

## 1 start with clarity and focus

- Lock the one KPI that matters for the first release; express it as a before/after statement with a date.
- Define the user, the moment of need, and the decision or task that the AI will assist with.
- Write the non-goals: what you will not do in your first version.

## 2 build on a strong foundation

- Treat data as the product: invest early in access, quality, lineage, and security.
- Staff the whole stack: data → model → platform → integration → UX → change management.
- Align the business and tech folks at every sprint; no surprises at launch.

## 3 design for adoption & longevity

- Put the human first. Build trust and create feedback loops.
- Bake in guardrails and incident response from day one.
- Measure outcomes continuously and iterate; retire what isn't working and double down on what is.

05

# About Rōnin Consulting

At Rōnin Consulting, we don't just follow trends, we help shape them. Our team has been developing and deploying real AI solutions that address significant business challenges across various industries.

Whether it's streamlining operations, automating time-consuming workflows, or empowering your team with AI-powered solutions, we bring practical, high-impact AI to the table.

Let's make your vision a reality, with AI that works in the real world.

**LET'S START THE CONVERSATION**

RŌNIN CONSULTING

06

# Sources


1) Bandaru, Vaishnavi. "Ambient AI Scribes: Kaiser Permanente's 7,000-Physician Study Illuminates the Future of Clinical Documentation." Future Medicine AI, 9 Apr. 2025, www.fmai-hub.com/ambient-ai-scribes-kaiser-permanentes-7000-physician-study-signposts-the-future-of-clinical-documentation/

2.) "Morgan Stanley Uses AI Evals to Shape the Future of Financial Services." OpenAI, 2025, openai.com/index/morgan-stanley/

3.) Morgan Stanley Wealth Management. "Launch of AI @ Morgan Stanley Debrief." Morgan Stanley, 26 June 2024, www.morganstanley.com/press-releases/ai-at-morgan-stanley-debrief-launch

4.) Burns, Matt. "Forward Health Launches CarePods, a Self-Contained, AI-Powered Doctor's Office." TechCrunch, 15 Nov. 2023, techcrunch.com/2023/11/15/forward-health-carepod-ai-doctor/

5.) McKay, Chris. "AI Healthcare Startup Forward Shuts Down." Maginative, 13 Nov. 2024, www.maginative.com/article/ai-healthcare-startup-forward-shuts-down/

6.) McClure, Tess. "Supermarket AI Meal Planner App Suggests Recipe That Would Create Chlorine Gas." The Guardian, 10 Aug. 2023, www.theguardian.com/world/2023/aug/10/pak-n-save-savey-meal-bot-ai-app-malfunction-recipes


RŌNIN CONSULTING