

Rapid AI Sprint

The CEO's AI Strategy Playbook

BY

SAM SCHREIM

ISBNs

Print (Paperback): 978-1-960908-18-6

© Copyright 2026 - All rights reserved.

The content contained within this book may not be reproduced, duplicated or transmitted without direct written permission from the author or the publisher.

Under no circumstances will any blame or legal responsibility be held against the publisher, or author, for any damages, reparation, or monetary loss due to the information contained within this book, either directly or indirectly.

Legal Notice:

This book is copyright protected. It is only for personal use. You cannot amend, distribute, sell, use, quote or paraphrase any part, or the content within this book, without the consent of the author or publisher.

Disclaimer Notice:

Please note the information contained within this document is for educational and entertainment purposes only. All effort has been executed to present accurate, up to date, reliable, complete information. No warranties of any kind are declared or implied. Readers acknowledge that the author is not engaged in the rendering of legal, financial, medical or professional advice. The content within this book has been derived from various sources. Please consult a licensed professional before attempting any techniques outlined in this book.

By reading this document, the reader agrees that under no circumstances is the author responsible for any losses, direct or indirect, that are incurred as a result of the use of the information contained within this document, including, but not limited to, errors, omissions, or inaccuracies.

A CEO's Note: Read This First

If there's a knot in your stomach every time "AI" is on the agenda, you're not alone. Modern-day leaders are being asked to champion a future fueled by AI, yet deep down, many are concerned they likely don't have enough knowledge to actually lead the way forward. That misalignment is very real—and it's rational: A February 2025 Cisco study found that while 97% of CEOs plan to integrate AI, a staggering 74% fear their own knowledge gaps will hinder decisions in the boardroom.

Constant pressure vs. personal ignorance

Boardrooms and the market are squeezing companies into a binary narrative: become an AI company or face irrelevance. In a soundbite world, many would admit privately that their knowledge isn't even close to a level allowing them to act decisively. This gap between the imperative to act and the knowledge required to act prudentially is where flawed and consequential decisions come to life.

Fear of getting it wrong

The fear is not of being embarrassed, but stifling an organization's growth, or for missing this window of opportunity, or green lighting the wrong bet entirely. You have the Visionary's Dilemma: Doing nothing will be fatal to the company but doing something without clarity will be reckless.

Information overload, and hype cycle

The relentless advance of AI has leaders perpetually on the back foot, which reduces the amount of time they get to spend on their actual businesses. To complicate matters, the hype surrounding AI from the tech sector is pervasive. As Palantir CEO Alex Karp described, Silicon Valley "totally effed up" because of the overhype surrounding LLMs and led executives to think they could just use a "glorified chatbot" and skip

the "painstaking work" that ultimately needs to happen in the real world. This creates a lot of false hope and disappointment.

There is an internal disconnect, and that is a dangerous one

Mandates like "AI-First" profess vision but usually lack specificity around execution. IT teams and data teams are then left to figure out how to execute a bumper sticker slogan, but without guardrails, friction, wasted effort, and pilot purgatory ensues.

Turn the problem upside down

The frustrations described above stem from an incorrect assumption: to lead on AI one must be a technologist. This is wrong.

Your lack of depth in technical capability is not a vulnerability. This is your strategic advantage. You are forced to do what potentially only leaders can do: rise above the services and make the choices that will determine whether or not AI creates value, or performance theatre.

Stay level and wrestle with four questions:

- **WHY** are we doing this? Does the initiative truly create competitive advantage (cost leadership, differentiation, or focus), not just another lightning fast demo reel?
- **WHAT** is the business value? Anchor to a P&L measure revenue, margin, cycle time, churn and articulate "how much, by when."
- **WHO** must be involved? Engage the crossed functions/owners (ops, data, product, finance, risk) and provide single-threaded accountability.
- **WHEN** are we going to prove/disprove? Time-box to 90 days with explicit kill rules. Wins are scalable. Science projects don't.

What you get from this book

This book is a field manual for strategic leaders, not a data-science course. You will gain:

- A simple method, to select one lighthouse initiative tied to a board-level KPI.
- Option-based funding: a few small bets, staged gates, prenegotiated kill rules.
- Lean governance scripts (charters, owner cards, decision logs) to keep work out of performance theatre and anchored in the P&L.

-
- A repeatable process to get from one winning 90-day to a portfolio that compounds.

Be confident in your leadership. Let your experts argue over models, you will provide the direction, guardrails, and scale what works. That's not faking it; that's leading.

About the Author

Sam Schreim is a strategy obsessive first and a technologist second. For more than two decades, he has helped Fortune 500 CEOs, fast-moving startups, and government leaders tackle the same fundamental questions: *Where should we play? How do we win?* As AI moved from the research lab to the executive agenda, Sam wove machine-learning tools into his strategy playbook—turning data models into levers for margin, speed, and public-sector impact. Today his work sits at the intersection of classic competitive thinking and the new, AI-powered engines that now drive it.

Table of Contents

FOREWORD: DOOMED IF YOU DO, DAMNED IF YOU DON'T	2
<i>Why Marvin Minsky's "Pet" Warning Still Haunts Us.....</i>	<i>2</i>
<i>The Pragmatic Middle the Headlines Miss</i>	<i>3</i>
<i>Why This Playbook Exists.....</i>	<i>4</i>
INTRODUCTION: FROM PILOT PURGATORY TO VALUE AT SCALE	5
WHAT "NO-BS" REALLY MEANS — AND WHY IT SHOULD BE YOUR DEFAULT SETTING .	6
CHAPTER 1 – DIAGNOSIS: WHY YOU'RE STUCK IN PILOT PURGATORY	7
THE ANATOMY OF FAILED PILOT.....	7
<i>A Solution Looking for a Problem</i>	<i>7</i>
<i>The Island of Automation</i>	<i>7</i>
<i>Data Delusions.....</i>	<i>8</i>
<i>The Science Project Mentality.....</i>	<i>8</i>
<i>Ignoring the Human Element.....</i>	<i>8</i>
MISALIGNING AI WITH BUSINESS STRATEGY	9
<i>AI for AI's Sake</i>	<i>9</i>
<i>Revisiting Porter in the Age of AI</i>	<i>10</i>
<i>Case Study – Waste Management Inc.....</i>	<i>10</i>
<i>Activity – Mapping Your Portfolio.....</i>	<i>11</i>
CHAPTER 2 – THE PLAYBOOK: A FRAMEWORK FOR STRATEGIC DEPLOYMENT .	12
PLAY 1: THE "LIGHTHOUSE" INITIATIVE	12
<i>Crafting the Proposal.....</i>	<i>13</i>
<i>Case Example: Invoice Match & Cash Acceleration</i>	<i>13</i>
PLAY 2: CHOOSING YOUR WEAPON – THE RIGHT AI FOR THE JOB	13
<i>Machine Learning (ML).....</i>	<i>13</i>
<i>Generative AI.....</i>	<i>14</i>
<i>Robotics & RPA</i>	<i>14</i>
<i>Matching the Tool to the Task.....</i>	<i>14</i>
PLAY 3: BUILDING THE "A-TEAM" – LEADERSHIP AND TECHNICAL REQUIREMENTS	15
CHAPTER 3 – EXECUTION: FROM PILOT TO PRODUCTION AND BEYOND	20
PLAY 4: THE HUMAN-MACHINE PARTNERSHIP	20
<i>Design for Collaboration, Not Replacement</i>	<i>20</i>
<i>Upskilling and Reskilling.</i>	<i>21</i>
<i>The Manager's New Role.....</i>	<i>21</i>
<i>Change Management Is Non-Negotiable</i>	<i>21</i>
PLAY 5: THE ETHICAL GUARDRAILS	22
<i>Proactive Risk Mitigation.....</i>	<i>22</i>

<i>EU AI Act and Regulatory Compliance</i>	22
<i>Beyond Regulation</i>	23
CONCLUSION: BECOMING AN AI-POWERED ORGANIZATION	24
REFERENCES	26

Foreword: Doomed If You Do, Damned If You Don't

The blunt, 3-minute wake-up call that shows executives how to turn AI from a money-bleed into a profit engine—before your competitors do

I was five minutes into a site visit at a distribution warehouse just outside Boston when the brand-new autonomous forklifts threw their tantrum. A lone mis-scanned pallet sent the fleet into a frantic stop-start loop, wheels screeching until the air stank of burned rubber. From my perch as the outside consultant— I watched \$200,000 worth of “cutting-edge AI” freeze the loading dock and sabotage the shift schedule in under ninety seconds.

Bridging the C-Suite Paradox

Before we get into forklifts, models, or MLOps, let's name the real challenge: most executive teams are being asked to lead decisively in an area in which they are not yet fluent. This playbook isn't going to teach you algorithms; it is going to give you a decision system: how to select a lighthouse that is in keeping with your strategy and how to fund options versus monuments; and how to make value visible within 90 days so that credibility replaces hype. Then, we scale what works and kill what doesn't. That's the no-bullshit promise of this book and why this book starts with leadership instead of code.

Why Marvin Minsky's "Pet" Warning Still Haunts Us

Back in 1970, MIT legend Marvin Minsky quipped that hyper-intelligent machines might one day keep humans “as pets.” The line was equal parts prophecy and provocation, and it still echoes in every boardroom where an executive wonders if AI will turn them into the pet—or the zookeeper.

Fifty-plus years later, that "whirlwind of interest" MIT faculty describe is deafening. Headlines ping-pong between salvation and doom, while middle managers drown in proofs-of-concept that never see the light of day. The data is stark: more than 90% of AI pilots stall, and a third of

companies languish in “pilot purgatory”—an endless cycle of expensive experiments that consume capital but produce zero enterprise value. You’re told to innovate or die, yet the very act of innovation seems to be a faster way to bleed money.

Doomed if you do, doomed if you don’t.

The Pragmatic Middle the Headlines Miss

Let’s drop the sci-fi trailer voice for a minute and talk like grown-ups in a boardroom.

What if AI isn’t a fork in the road between heaven and doomsday? What if—exactly as the folks up at MIT keep hammering home—it’s just a messy middle that needs plain-old business discipline? Sit with the pioneers who’ve actually built these systems and a different picture appears. They don’t speak of birthing a new digital species; they talk about unleashing incredibly capable tools. Professor Patrick Winston, a luminary in the field, put it plainly: *“For the next few years, the engine of AI is machine learning, and machine learning is simply a new kind of statistics.”*

Translation: AI performs astonishing, human-like tricks, but the algorithm doesn’t **‘understand’** anything. It just spots patterns faster than we ever could.

This is the crucial insight: We are not dealing with science fiction robots. We are dealing with a different kind of intelligence, one that is better than humans in some ways and far less capable in others. The consensus among these experts is that true, human-level general AI is still *“decades, maybe even centuries away”*.

Right now, “artificial intelligence” is less robot overlord and more turbocharged spreadsheet—blazing fast at crunching numbers, but totally clueless about the big picture. It’s an amazing tool but a lousy prophet. Treat machine learning like a power tool: use it right and you’ll cut costs, save hours, and lap the competition. Use it wrong, and you’re just waving a chainsaw by the wrong end. Human-level super-brains? A fine fantasy, but still a distant speck on the horizon.

I’ve spent the last decade knee-deep in AI rollouts—watching chatbots crank revenue and “moonshot” models set budgets on fire. MIT didn’t

hand me a secret sauce; it confirmed the street-smart rule I'd already learned: **the math is solid, the magic is in the execution.** That gap—the yawning space between theory and Tuesday-morning results—is exactly what this playbook closes.

Why This Playbook Exists

If you picked up *The No-BS AI Strategy Playbook*, you feel the same tension my client felt in that Boston warehouse: ignore AI and risk obsolescence or embrace it and risk hemorrhaging cash.

This book is my field manual for escaping that false choice. It distills the frameworks I sharpened and the hard-won lessons from years in the trenches. You'll find checklists that force pilots to answer C-level questions before a single line of code is written. You'll learn to spot the “*science-project smell*” in week one, to anchor every use-case to a P&L metric, and to defend budgets with evidence, not hype.

Most of all, you'll learn that speed and stewardship are not enemies. You can make AI pay off fast—without bleeding money or missing milestones.

I wrote this foreword for the leader who walks into tomorrow's staff meeting and refuses to approve another vanity pilot. May the pages ahead replace paralysis with a plan—and may your forklifts (literal or metaphorical) stay cool under pressure.

Sam Schreim

Introduction: From Pilot Purgatory to Value at Scale

“The essence of strategy is choosing what not to do.” – Michael E. Porter

For the last decade, artificial intelligence has dominated the boardroom agenda. Success stories circulate with breath-less optimism—algorithms detect disease faster than doctors, robots pick and pack like tireless warehouse workers, language models write like seasoned authors. Analysts project trillions in economic opportunity. Yet when you peel away the hype, the reality inside most organizations is starkly different.

The gap between AI ambition and reality is stark. While a 2024 McKinsey report shows AI adoption has surged, with 72% of organizations now using it in at least one business function, the returns remain elusive. In fact, another recent McKinsey study found over 80% of businesses are not yet seeing significant earnings gains from the technology. This disconnect is rooted in execution: Gartner finds a staggering 85% of AI projects fail to meet expectations, leaving many companies trapped in a “pilot purgatory” of isolated experiments that consume capital but produce little enterprise value. Consequently, while nearly every company is investing in AI, a 2025 McKinsey report reveals that just 1% consider their organizations “AI-mature”. The frustration is palpable because the workforce is ready—a recent Salesforce survey found that 96% of workers have used AI to perform tasks they previously lacked the skills to do, signaling a clear appetite for productivity gains that leadership has yet to unlock at scale.

This book is written for those leaders. As a strategy & innovation officer who has shepherded AI projects from concept to scaled platform, I’ve seen pilot purgatory up close. The goal of this playbook is not to teach you how to code or to build neural networks. It is to give you a structured, repeatable approach for moving from isolated experiments to an integrated portfolio of AI initiatives that create compounding business value.

We begin by diagnosing why pilots fail, then outline a series of **plays**—strategic moves for executives—to identify the right projects, select the right technologies, assemble the right teams, manage the human impact

and ethical risks, and build a measurement framework that leads to scale. Throughout, you will find examples drawn from real companies and research to illustrate what works and what doesn't.

What “No-BS” Really Means — and Why It Should Be Your Default Setting

This playbook wears its title like a hard hat. It's built to bulldoze the buzzwords, vendor smoke, and “science-project smell” so you can get straight to business value. The **No-BS** approach rests on three ironclad rules. Break anyone, and you'll be bleeding cash while your rivals sprint ahead:

- **ROI over R&D:** very framework and checklist in this book answers a single, brutal question: Does this move the needle on revenue, margin, or our competitive moat? Full stop. We skip the white-paper theory for field-tested scorecards that force every pilot to earn its oxygen with CFO-proof metrics. When UPS shaved 100 million miles off delivery routes with its ORION algorithm—saving an estimated \$300–\$400 million a year—it wasn't chasing cool tech; it was chasing cashable fuel savings. That's the bar.
- **Execution over Experimentation:** “Pilot purgatory” is littered with clever models that died on a dashboard because nobody planned for production. The No-BS remedy: design for scale on day one. That means data contracts, workflow hooks, and MLOps guardrails baked in before the first line of code. The U.S. Air Force's “Kessel Run” program ships combat-ready apps every 11.2 hours because its pipeline automates testing, security scans, and deployment. Glamorous? No. Essential? Absolutely.
- **Strategy over Shiny Tools:** The most expensive blunder in business is deploying technology with no link to your core strategy. Netflix nails the opposite: its recommendation engine, laser-aligned with its subscription model, saves an estimated \$1 billion a year by reducing customer churn. If a tool doesn't directly amplify your competitive advantage—be it cost leadership, differentiation, or focus. It's a distraction your competitors will celebrate.

Stick to these three rules and AI becomes a profit engine. Ignore them and you're just funding someone else's keynote slide.

Chapter 1 – Diagnosis: Why You’re Stuck in Pilot Purgatory

“The greatest danger in times of turbulence is not the turbulence; it is to act with yesterday’s logic.” – Peter F. Drucker

Pilot purgatory isn’t a technology problem—it’s a vision problem. Before another dollar slips into a proof-of-concept abyss, pause and ask whether the experiment actually attacks a board-level pain point. Identify the strategic artery—cost, growth, or risk—then trace where data, process, and people are silently bleeding value. When you spotlight that gap with numbers the CFO cares about, urgency replaces curiosity and momentum follows.

The Anatomy of Failed Pilot

A Solution Looking for a Problem

Too many AI pilots begin with a technology in search of a problem. Teams choose a “cool” algorithm rather than a critical business challenge. The result is a proof-of-concept that dazzles data scientists but leaves the CFO unimpressed. To escape pilot purgatory, projects must be anchored to strategic priorities—cost reduction, revenue growth or risk mitigation—not curiosity. As we’ll see in later chapters, the most successful organizations apply **litmus tests** for value, feasibility and data readiness.

The Island of Automation

Even when pilots address real problems, they often remain isolated from core systems and workflows. A predictive maintenance model may forecast machine failures, but if the results aren’t integrated into the maintenance schedule and procurement system, the insight dies on a dashboard. This “island of automation” phenomenon stems from treating AI as a bolt-on rather than as part of the enterprise architecture. In the words of CIO Magazine, sustainable value requires cross-functional alignment and shared ownership across the C-suite.

Data Delusions

Many projects falter because leaders underestimate data challenges. Clean, contextualized, accessible data is the oxygen for AI. A recent Harvard Business Review report underscores this reality: while 91% of leaders agree a reliable data foundation is essential for AI, just over half believe their organization's is reliable. This explains why, when asked about their top data-related efforts for the next year, 71% of organizations are focused on cleaning data to improve its accuracy and 66% are working to break down data silos. The stakes are high; Gartner predicts that through 2025, at least half of all generative AI projects will be abandoned at the pilot stage due to poor data quality. Without a foundational investment in data governance, lineage, and privacy, AI remains a science project.

The Science Project Mentality

Pilots often lack a clear path to production. Teams build prototypes without thinking about scalability, integration or regulatory compliance. When it's time to deploy, they discover that no one has budgeted for cloud infrastructure, MLOps pipelines or change management. To succeed, you must design for production from day one. This means anticipating security reviews, testing procedures and user adoption early.

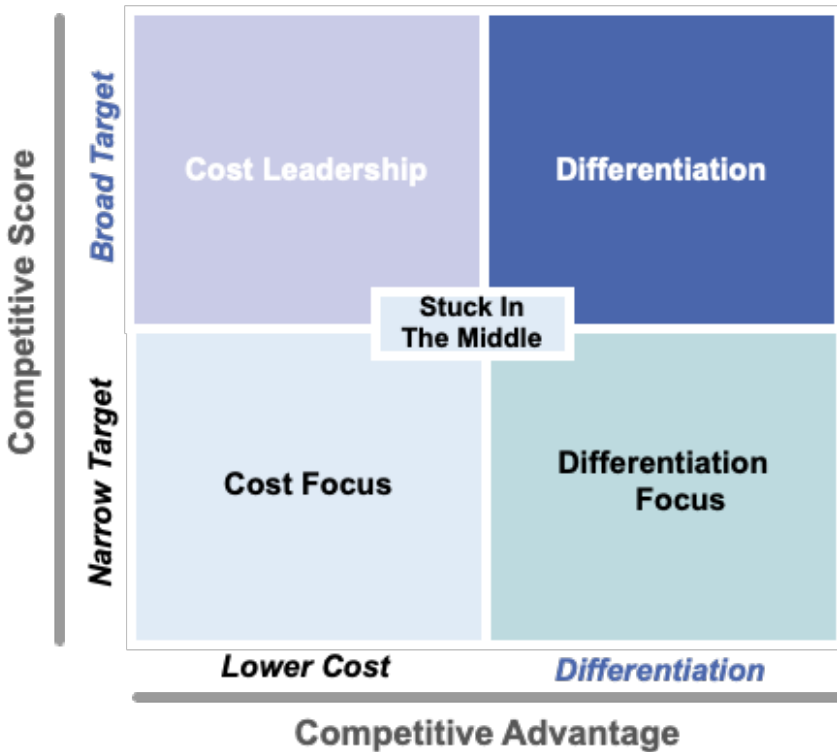
Ignoring the Human Element

Technology alone does not change organizations—people do. Failed pilots often overlook training, communication and ethics. McKinsey's research on AI at work shows that 41% of employees are apprehensive about AI and half worry about inaccuracies or cyber-risks. Effective pilots require change management, upskilling and clear governance to build trust and adoption. We'll return to these themes later.

Misaligning AI with Business Strategy

AI for AI's Sake

Porter's Generic Strategies



The gravest sin in the AI playbook is deploying technology without strategic alignment. A powerful way to generate this is by applying the "Value Chain" Method, which systematically maps AI opportunities across the entire organization. The process begins by deconstructing your value chain into its primary activities (Inbound Logistics, Operations, Outbound Logistics, Marketing & Sales, Service)

and support activities (Procurement, Technology Development, HR, Firm Infrastructure). Next, for each activity, hold workshops with functional leaders to brainstorm pain points and opportunities, asking

THE NO-BS AI STRATEGY PLAYBOOK

where you can reduce costs (Cost Leadership), deliver more value (Differentiation), or eliminate repetitive tasks (Automation). Finally, you apply the AI lens to each point, asking if prediction (ML), content generation (GenAI), or task automation (Robotics/RPA) could help. This yields a comprehensive long list of potential AI initiatives linked to specific business activities.

With this long list in hand, you can now assess for strategic fit. According to Michael Porter's seminal work on competitive strategy, companies must choose between cost leadership, differentiation or focused niche; straddling all three leads to mediocrity. You can map your portfolio with "Porter's AI Lens" by first reaffirming your primary strategy. Then, create strategic buckets for Cost Leadership, Differentiation, and Focus, and place each initiative from your long list into the bucket it most directly serves. This creates a clear visual map showing which projects support which strategic goals. Initiatives that are hard to map may be poorly aligned and can be discarded.

AI should amplify, not distract from, your chosen advantage. Cost leaders can use AI for predictive maintenance, dynamic pricing or supply chain optimization. Differentiators may leverage AI for personalized customer experiences or superior product quality. Niche players might deploy AI to serve a specialized segment more efficiently. When AI initiatives pull resources away from the core strategy, they stall.

Revisiting Porter in the Age of AI

A 2024 strategy article emphasized that analytics and AI must be anchored to Porter's strategies. If you compete on cost, invest in automation and optimization; if you differentiate on quality, invest in AI-driven design and customer insights. The danger of a technology-led approach is that you spend on shiny tools while competitors strengthen their moats.

Case Study – Waste Management Inc

Waste Management (WM) illustrates how to blend cost leadership and differentiation. The company operates 136,000 refuse vehicles with a mean fuel economy of 2.8 mpg. By applying AI-based route optimization, WM reduces route planning time by 25–75% and cuts

CO₂ emissions up to 25%. Machine-vision systems like Netradyne identify driver fatigue and distraction—a critical issue because fatigue contributes to 40% of trucking accidents. Predictive maintenance tools use digital twins and sensors to detect failures, noting that roadside repairs cost four times more than in-shop fixes. WM also invests in reskilling through its *Your Tomorrow* program, offering over 170 fully funded education programs to employees and dependents. These initiatives both reduce costs and enhance service quality, demonstrating strategic alignment.

Activity – Mapping Your Portfolio

Before moving on, map each of your AI pilots to your business strategy. Is it driving cost leadership, differentiation or focus? Use a simple table:

Pilot	Strategic Objective	Status
Example: Predictive maintenance on fleet	Cost leadership – reduce downtime and fuel costs	Pilot delivered ROI; ready for scale
Example: Customer-service chatbot	Differentiation – improve customer response time	Pilot stuck in purgatory; lacks integration
...

This exercise reveals misalignments and helps prioritize which projects deserve investment.

Chapter 2 – The Playbook: A Framework for Strategic Deployment

“By failing to prepare, you are preparing to fail.” – Benjamin Franklin

A great playbook is less about pages of code and more about disciplined preparation. Start by framing one “lighthouse” win that aligns to your competitive edge, pressure-testing it for data readiness and stakeholder ownership. Then choose the right tool—ML, GenAI, or RPA—like a craftsman selecting the exact blade. A crisp charter, clear success metrics, and cross-functional accountability turn theory into a well-sharpened axe ready to swing.

The Lighthouse Litmus Test – 3 Pass/Fail Questions

- 1 Strategic Alignment**
🔍 *Will fixing this problem visibly strengthen our chosen competitive edge (cost, differentiation, or focus) and move the needle on revenue, margin, or customer loyalty?*
- 2 Data Readiness**
🔍 *Do we already have clean, accessible, well-governed data—enough in volume and quality to train and run the model without a lengthy data-cleanup project?*
- 3 System Regularity**
🔍 *Is the target process frequent, repeatable, and outcome-based so an algorithm can learn from past examples and measure success?*

Play 1: The “Lighthouse” Initiative

To escape pilot purgatory, you need a **lighthouse project**—a high-value, feasible initiative that demonstrates AI’s potential and builds momentum. Choosing the right project involves three litmus tests.

1. **Strategic Alignment:** Does the problem amplify your competitive edge? A cost leader should target operational efficiency; a differentiator might

personalize offerings. Evaluate whether success would move the needle on revenue, margin or customer loyalty.

2. **Data Readiness:** Do you have the data to win? As noted earlier, clean data is critical. A 2025 report found that data readiness is the fourth most cited driver for data governance. If your data is scattered, incomplete or ungoverned, start by building a data foundation.
3. **System Regularity:** Is the process predictable enough for an algorithm? AI thrives on **regular, repeatable processes** with clear outcomes. Unstructured, one-off tasks rarely benefit from automation. Ask whether the task occurs frequently, follows a consistent pattern and generates measurable outcomes.

Crafting the Proposal

Once you've identified a candidate, develop a clear proposal including business objectives, success metrics, timeline, resources and risk mitigation. In my experience, executives fund proposals that link directly to strategic goals and provide a credible path from pilot to production.

Case Example: Invoice Match & Cash Acceleration

- **Goal:** Shorten Day Sales Outstanding (DSO) by auto-matching invoices to POs and proofs of delivery.
- **KPI:** Days Sales Outstanding and write-off rate.
- **Data:** ERP tables (invoices, POs), delivery scans, exception codes.
- **Thin proof:** Sample 5,000 invoices; show % auto-matched and dollar value released inside 30 days.

Why it works: No new revenue claims—just faster cash and fewer errors—making it a credible lighthouse in finance and operations.

Play 2: Choosing Your Weapon – The Right AI for the Job

Machine Learning (ML)

Traditional machine learning excels at prediction, forecasting and optimization. It is suitable for predictive maintenance, demand forecasting, fraud detection and dynamic pricing. It requires labelled

data and often yields transparent models. Use ML when your task depends on structured data and you need consistent accuracy.

Generative AI

Generative AI is a newer form of machine learning that creates new content. According to MIT's 2025 overview, generative AI models can produce text, images and code and often deliver results faster than custom ML models. They are ideal for content summarization, report generation, marketing copy and synthetic data creation. However, generative models may hallucinate, so they are less suitable for regulated or safety-critical domains. MIT cautions that generative AI is more accessible but may not suit domain-specific tasks. Avoid the **"Generative AI hammer"**—not every problem needs a large language model.

Robotics & RPA

Robotic process automation (RPA) and robotics automate physical or digital tasks. A 2025 trend report notes that RPA will evolve into **platform plays** integrating BPM, AI and low/no-code tools. RPA excels at repetitive, rule-based processes such as invoice processing or data entry. In manufacturing and logistics, robotics handles material movement, sorting and inspection. Combining AI and RPA enables "intelligent automation," where ML models determine when to hand tasks between humans and machines.

Matching the Tool to the Task

Use a decision matrix that considers data type, required output (prediction vs generation), process complexity and regulatory requirements. For example, if your lighthouse initiative is route optimization, a supervised ML model combined with heuristics may suffice. If you need to generate safety reports from sensor logs, a generative model may help. When you need to automate invoice approval, RPA integrated with an ML classifier is appropriate.

Play 3: Building the “A-Team” – Leadership and Technical Requirements

Scaling AI requires a cross-functional squad. AI’s impact is inherently cross-cutting and requires shared ownership across the C-suite.

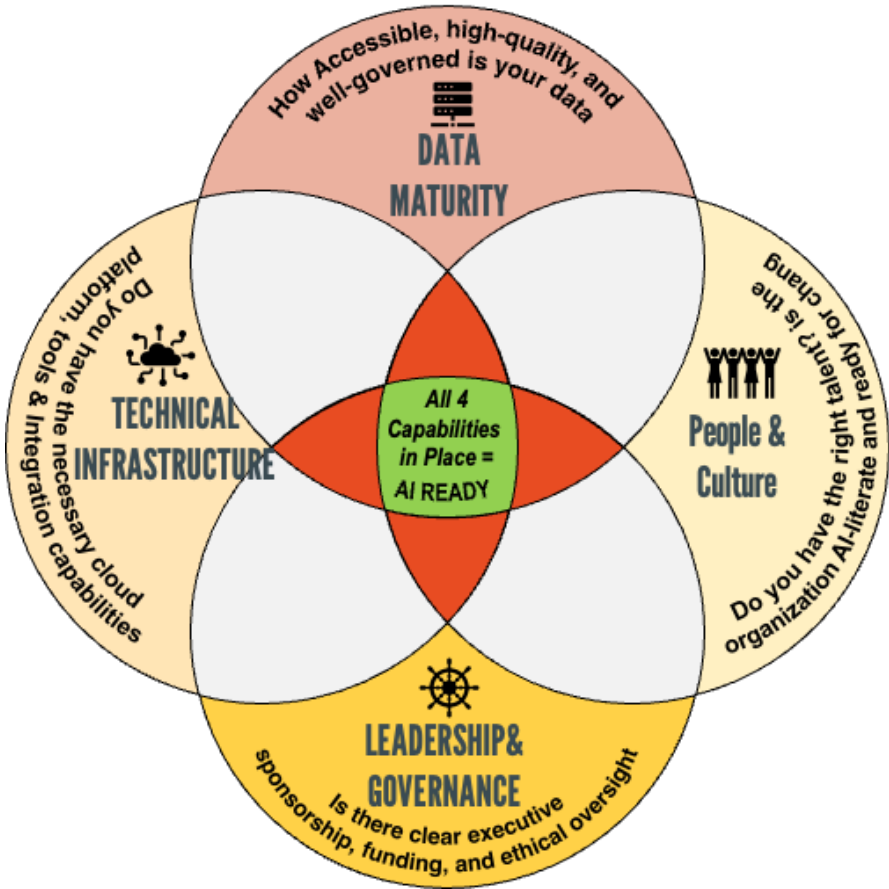
Before you assemble the team, however, you must assess if the organization itself is ready. You can do this with the "**Organizational AI Readiness**" Framework, which assesses your capability to deliver on an AI initiative.

This framework examines four key pillars:

1. **Data Maturity** (how accessible, high-quality, and well-governed is your data?);
2. **Technical Infrastructure** (do you have the necessary cloud platforms, tools, and integration capabilities?);
3. **People & Culture** (do you have the right talent, and is the organization AI-literate and ready for change?); and
4. **Leadership & Governance** (is there clear executive sponsorship, funding, and ethical oversight?).

A low score in any area represents a significant risk that must be addressed before the project begins.

Organizational AI Readiness Framework



Once you have a handle on readiness, you can assemble the team, assigning clear roles to the CIO, CDO, CISO, CHRO, CFO, and other strategic leaders.

Each executive plays a distinct role as follow:

Executive Role	Core Responsibilities
CIO/CTO	Advocate for including AI in enterprise architecture; accelerate the shift from pilots to enterprise-grade platforms; define a technology vision and invest in scalable AI infrastructure
Chief Data Officer (CDO/CDAO)	Align data governance with AI needs; ensure traceability, quality and compliance; decommission legacy data constraints; champion ethical data use
Chief Information Security Officer (CISO)	Establish standards for model security and integrity; lead risk assessments for adversarial AI and data poisoning; embed cybersecurity into AI development
Chief Human Resources Officer (CHRO)	Lead cultural transformation for human-AI collaboration; embed AI literacy into workforce planning; redesign roles to emphasize augmentation over automation
Chief Financial Officer (CFO)	Create valuation models that capture AI's long-term capability returns; integrate AI metrics into capital investment governance; support iterative, cross-functional funding
CEO/Chief Strategy Officer (CSO)	Set a unifying vision for AI as a lever for differentiation; translate strategic intent into accountable priorities; engage external stakeholders and regulators
Chief Risk Officer/General Counsel (CRO/GC)	Shape the organization's AI risk appetite; establish governance mechanisms that address legal, ethical and reputational risks; ensure compliance with emerging regulations
Cross-Functional Governance Team	Create shared accountability models that connect AI initiatives to enterprise-wide KPIs and ensure integration across functions

In addition to executives, your “A-team” should include data scientists, ML engineers, product managers, process owners, change-management specialists and ethics/legal experts. Decide early whether to **buy, build or partner**. Buying off-the-shelf tools speeds deployment but may limit differentiation. Building in-house provides control but requires talent. Partnering with vendors, universities or startups can accelerate learning

THE NO-BS AI STRATEGY PLAYBOOK

but may create dependencies. Your decision should align with strategic priorities and resource availability.

Chapter 3 – Execution: From Pilot to Production and Beyond

"Vision without execution is hallucination." – Thomas A. Edison

The handoff from pilot to production is where bold ideas usually stall, so design for scale on day one. Tie model outputs directly into frontline workflows, automate the boring but critical MLOps plumbing, and embed change-management rituals that upskill people, not replace them. Ethical guardrails and real-time KPIs keep the engine honest while it gathers speed. When process, platform, and people click, AI shifts from an experiment to a profit habit.

Play 4: The Human-Machine Partnership

AI success is as much about people as it is about algorithms. The scale of this transformation is immense; a 2025 joint study from the UN's International Labor Organization found that one in four jobs worldwide is potentially exposed to generative AI, with clerical and administrative roles—where women are overrepresented—facing the highest risk of task automation. Yet this exposure points less to outright job loss and more to a profound labor evolution. The World Economic Forum estimates that while 92 million jobs may be displaced by 2030, a staggering 170 million new roles could emerge from the new division of labor between humans and machines. This transition hinges on a massive skills upgrade. The World Economic Forum's 2025 report projects that, on average, 39% of a worker's core skills will be disrupted in the next five years. In response, 77% of employers plan to prioritize reskilling and upskilling their workforce to navigate this shift. The message from the data is clear: the future of human-AI collaboration is about augmentation, not replacement.

Design for Collaboration, Not Replacement

Begin by decomposing work into tasks best suited for machines (high-volume, routine, data-driven) and tasks requiring human judgment, empathy or creativity. Use AI to augment employees—provide

recommendations, automate drudgery, surface insights—while people handle exceptions, relationship building and decision-making. MIT’s article on what machine learning and generative AI are good for in 2025 notes that generative models are ideal for everyday language tasks but less suitable for domain-specific or privacy-sensitive areas. Use this principle to allocate tasks.

Upskilling and Reskilling.

The *Your Tomorrow* program at Waste Management discussed earlier, serves as a good benchmark: WM offers **170+ fully-funded education programs** to 36,000 full-time employees and plans to extend them to dependents. Courses focus on data analytics, digital transformation and systems thinking. Such programs signal commitment to employees and equip them to thrive in an AI-augmented workplace. Globally, the PEX Network emphasizes that training and upskilling are **not optional but essential**. Invest in AI literacy, technical skills and change leadership for all levels of the organization.

The Manager’s New Role

Managers move from supervising tasks to **orchestrating human-machine workflows**. They must become coaches who interpret insights, encourage experimentation and create psychological safety. Change management research highlights that **leaders set the tone**; they must address fears, build capabilities, articulate the AI vision and empower teams. Leading through uncertainty requires transparency, trust and a willingness to let people test, learn and iterate.

Change Management Is Non-Negotiable

The PEX Network outlines eight principles for change management in the AI era:

- People come first, trust is key
- Leadership sets the tone
- Change is continuous
- Training and upskilling are essential
- Experimentation breeds innovation

- Holistic views are paramount
- Agility fosters resiliency

These principles should underpin your communication and training plan. Monitor adoption metrics, solicit feedback and iterate convincing.

Play 5: The Ethical Guardrails

Ethical considerations are not a footnote; they are central to sustainable AI. Harvard Business School highlights algorithmic bias, cybersecurity risks and privacy/inclusiveness as key concerns. Biased training data can perpetuate discrimination; fairness audits and diverse teams mitigate this risk. AI also expands the attack surface—85% of cybersecurity leaders reported AI-related cyberattacks, but training can reduce phishing incidents by 86%. Data minimization and robust security practices are essential.

Proactive Risk Mitigation

Establish an AI ethics board that includes legal, HR, data science and external stakeholders. Develop a code of conduct covering data collection, model design, testing, deployment and monitoring. Conduct impact assessments before launching high-risk AI. Document data sources, labeling practices and model decisions. Build red-team exercises to test for adversarial attacks and bias.

EU AI Act and Regulatory Compliance

The European Union’s landmark AI Act takes a risk-based approach. It divides AI systems into categories:

1. **Unacceptable:** Unacceptable systems—such as social scoring or AI that exploits vulnerable groups—are prohibited.
2. **High-risk systems** (e.g., biometric identification, critical infrastructure) must implement risk management, transparency, accuracy and human oversight requirements.
3. **General Purpose AI** including large language models, must meet documentation and data governance obligations when used in high-stakes applications.

4. **Limited/Minimal risk:** Low-risk systems face minimal obligations but must disclose when users interact with A.

Understanding these obligations early helps avoid costly rework and reputational damage.

Beyond Regulation

Ethical AI also encompasses job displacement and fairness. The ILO report warns that women in clerical roles are 2.5 times more exposed to generative AI displacement than men. Organizations should design reskilling programs that priorities equity and inclusion. When AI makes decisions about people (hiring, lending, healthcare), ensure diverse representation in design and oversight. Ethical guardrails are a differentiator: companies that build trust will earn customer loyalty and regulatory goodwill.

Conclusion: Becoming an AI-Powered Organization

In Leaving pilot purgatory is not about luck. It is about **strategic discipline** and **leadership courage**. To recap, successful AI deployment requires:

1. **Diagnosing failure**—acknowledge when projects solve the wrong problem, remain isolated, ignore data readiness or overlook people.
2. **Aligning AI with strategy**—anchor initiatives to cost leadership, differentiation or focus; avoid technology for its own sake.
3. **Launching lighthouse projects**—choose problems with clear strategic impact, data readiness and predictable processes; craft concrete proposals.
4. **Selecting the right tools**—match machine learning, generative AI and robotics to your task; avoid one-size-fits-all approaches.
5. **Building the right team**—assign clear roles across the C-suite and technical disciplines; decide whether to buy, build or partner.
6. **Empowering people**—design human-AI collaboration, invest in upskilling, embrace continuous change and communicate transparently.
7. **Setting ethical guardrails**—address bias, cybersecurity, privacy and fairness; comply with regulations like the EU AI Act.
8. **Measuring and scaling**—define business-centric KPIs, create an AI flywheel and aspire to a digital twin that unifies data and decision-making.

Your First 90 Days

For leaders ready to act, here is a checklist for the next three months:

1. **Assess your AI portfolio.** Map current pilots against strategic objectives and data readiness. Kill vanity projects that don't deliver value.
2. **Identify a lighthouse opportunity.** Use the litmus tests to select a high-value, feasible initiative. Draft a proposal with objectives, metrics and resources.
3. **Form your A-team.** Assign C-suite owners and recruit cross-functional members. Define roles and governance structures.
4. **Plan for change management.** Develop a communication strategy that emphasizes people, trust, leadership, continuous change and experimentation. Launch training for AI literacy and reskilling.
5. **Establish ethical and regulatory frameworks.** Set up an AI ethics board; review your initiatives against the EU AI Act risk categories. Begin bias audits and security testing.

6. **Define KPIs and measurement protocols.** Decide how you will quantify success and use data to iterate. Connect metrics to business outcomes.

AI is not a one-time project—it is a **capability**. It thrives in organizations that learn continuously, invest in people and remain ethically grounded. By following this playbook, you can transform AI from a series of isolated experiments into a self-reinforcing engine that delivers scalable, sustainable value. The journey begins now.

References

- Burnham, K. (2024, August 20). *How artificial intelligence is transforming logistics*. MIT Sloan School of Management.
- Mayer, Hannah, et al. "Superagency in the workplace: Empowering people to unlock AI's full potential at work." *McKinsey & Company*, 28 Jan. 2025,
- Brown, S. (2025, June 12). Machine learning and generative AI: What are they good for in 2025? MIT Sloan School of Management.
- Cosmas, A., Cruz, G., Cubela, S., Huntington, M., Rahimi, S., & Tiwari, S. (2024, April 11). *Digital twins and generative AI: A powerful pairing*. McKinsey & Company.
- Hill, M. (2025, May 14). Change management in the AI era – 8 things business leaders must know. Process Excellence Network.
- Jampula, R. (2025, July 3). Building the AI flywheel for sustainable success. TechCircle.
- Brown, Sara. "Machine Learning and Generative AI: What Are They Good For in 2025?" MIT Sloan. June 2, 2025.
- Security Compass. (2024). *Understanding EU AI Act risk categories*. Security Compass Blog.
- 10xDS Team. (2025). How digital twins are transforming industries in 2025: Key applications & investor insights.10xDS.
- Waste Management Inc. website and press releases.
- Chui, Michael, et al. "The State of AI in 2023: Generative AI's Breakout Year." *McKinsey & Company*, 1 Aug. 2023
- Afshar, Vala. "Business Leaders Continue to Push Workers Toward Daily Use of AI." *ZDNet*, 26 June 2025, .
- Council, Forbes Technology. "Raising The Success Rate Of AI Deployment Across Industries." *Forbes*, 26 June 2025, .
- "85% Of AI Projects Tend To Fail - And How You Can Save It." *PixiTech*, 26 June 2025, pixitech.io/85-percent-of-ai-projects-may-fail-and-how-you-can-save/.
- Starr, Observer. "Over 80% of Companies Embracing A.I. See No Real Gains, McKinsey Finds." *Observer*, 25 June 2025,
- "The State of AI in Early 2024: Gen AI Adoption Spikes and Starts to Generate Value." *McKinsey & Company*, 30 May 2024.
- Harvard Business Review Analytic Services. *2024 Harvard Business Review Analytic Services Report on the Age of AI*. Profisee, 2024.
- Stahl, Alexander. "The AI Implementation Paradox: Why 42% of Enterprise Projects Fail Despite Record Adoption." *Medium*, June 2025.
- AI and the Future of Work." *Sand Technologies*, 2025.

- "AI threatens one in four jobs – but transformation, not replacement, is the real risk." *UN News*, 20 May 2025.
- "The Future of Jobs Report 2025: A Deep Dive." *SHRM*, 2025, www.shrm.org/topics-tools/flagships/ai-hi/future-of-jobs-report-2025-deep-dive.
- "Think your job is safe from AI? These 4 careers are already on the chopping block." *The Economic Times*, 2025.
- World Economic Forum. *The Future of Jobs Report 2025: Digest*. 2025.