How I'd Scale AI in a Mid-Size SaaS Company in 90 Days

Disclaimer

This scenario breakdown is a fictionalized, illustrative case study created for educational and strategic thinking purposes. While inspired by real-world patterns and organizational challenges, all details—company context, team structure, and suggested approaches—are generalized and do not represent any specific employer, client, or confidential situation.

The content is designed to demonstrate strategic problem-solving, not to prescribe one-size-fits-all solutions. Readers are encouraged to adapt ideas and frameworks to suit their unique organizational needs, capabilities, and compliance contexts.

Context: The Situation

A mid-size SaaS company (approx. 300–500 employees) has a leadership team excited about AI but no concrete implementation roadmap. They've experimented with ChatGPT internally, a few developers built side projects, and execs keep mentioning AI in town halls — but there's no structure, budget, or alignment.

The product org is fragmented — PMs are unsure where AI fits in the roadmap, engineers are unclear on approved tools or processes, and data privacy concerns are creating friction between innovation and compliance. There's potential, but no orchestration.

Root Problems

- **No Shared AI Vision**: Leadership mentions AI in aspiration, but it hasn't translated into prioritized roadmaps, funding, or team charters.

- **Capability Gaps Across Roles**: Engineers lack access to models or experimentation space. PMs lack fluency to define realistic AI features.

- **Missing Infrastructure**: No secure, sandboxed environments for testing; no GPU access; fragmented toolchains.

- **Risk Aversion**: Legal, security, and privacy concerns create a culture of "better not try."

- **Isolated Prototyping**: Early experiments aren't visible, shared, or reused — most die in the dev's notebook.
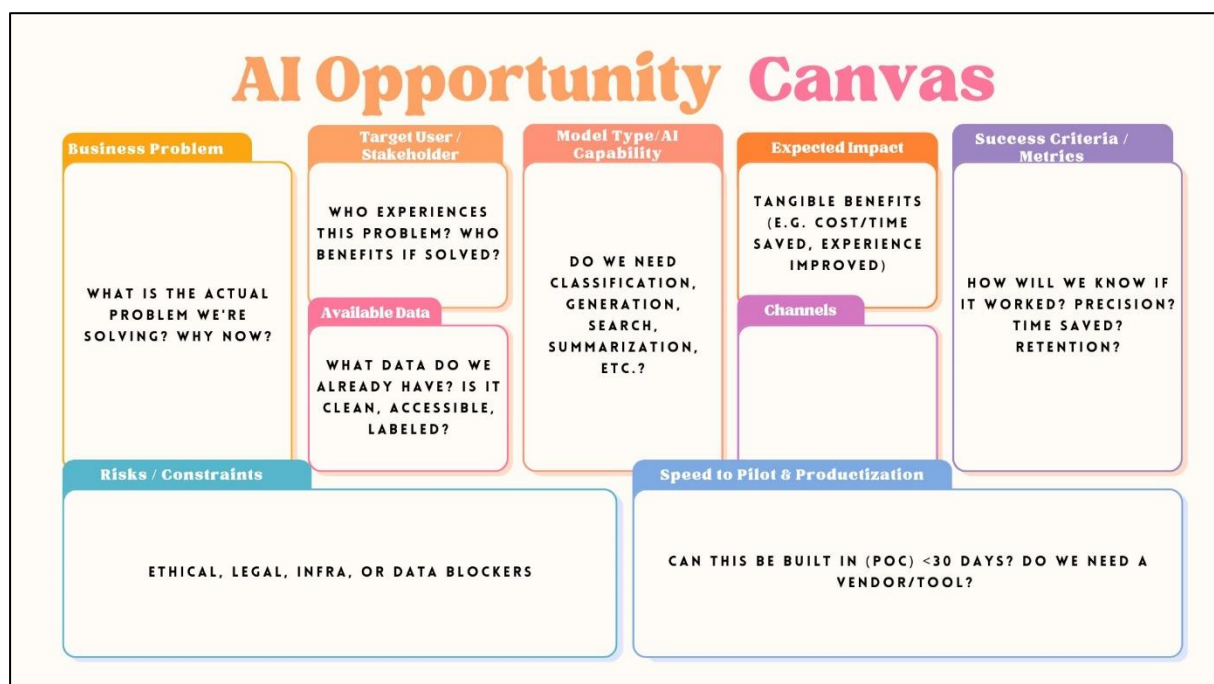
Approach & Framework

To bring coherence to this chaos, we use the **"Enable → Align → Build"** strategic rollout model. It's designed for speed *and* structure.

◈ **Phase 1: Enable**

- Run an **AI Readiness Audit** across roles (skills, infra, mindset).

- Set up secure sandboxes with OpenAI API keys, HuggingFace endpoints, or internal inference clusters.

- Create an **Internal AI Hub** — includes AI playbooks, case studies, working code templates, best practices.

- Host a cross-functional **AI Bootcamp** tailored for PMs, developers, QA, and design (mix async + live).

◈ **Phase 2: Align**

- Conduct **AI Ideation Workshops** — structured sessions where PMs, Designers, and Engineers co-create use cases.

- Introduce the **AI Opportunity Canvas** to qualify and vet ideas against feasibility, data, and ROI.

- Shortlist top 3–5 use cases tied to existing business goals (reduce churn, automate manual workflows, etc.).

- Pull legal & security into the loop early — define red/yellow/green data zones and logging requirements.

AI Opportunity Canvas

Business Problem	Target User / Stakeholder	Model Type/AI Capability	Expected Impact	Success Criteria / Metrics
WHAT IS THE ACTUAL PROBLEM WE'RE SOLVING? WHY NOW?	WHO EXPERIENCES THIS PROBLEM? WHO BENEFITS IF SOLVED?	DO WE NEED CLASSIFICATION, GENERATION, SEARCH, SUMMARIZATION, ETC.?	TANGIBLE BENEFITS (E.G. COST/TIME SAVED, EXPERIENCE IMPROVED)	HOW WILL WE KNOW IF IT WORKED? PRECISION? TIME SAVED? RETENTION?
	Available Data WHAT DATA DO WE ALREADY HAVE? IS IT CLEAN, ACCESSIBLE, LABELED?		**Channels**	

Risks / Constraints	Speed to Pilot & Productization
ETHICAL, LEGAL, INFRA, OR DATA BLOCKERS	CAN THIS BE BUILT IN (POC) <30 DAYS? DO WE NEED A VENDOR/TOOL?

Other Optional Sections you can add: -

- **Strategic Alignment:** Which OKRs or business goals does this align with?
- **Ownership:** Who leads this pilot? Who supports it?
- **Baseline / Status Quo:** What is the current state? (manual work? legacy system?)
- **User Journey Touchpoint:** Where does this fit in the customer or internal workflow?
- **Data Sensitivity Level:** Mark Red / Yellow / Green zone directly in canvas
- **Reusability:** Can the model/code be reused for other use cases or teams?
- **Dependencies:** Needs a new API, new logging, buy-in from Legal?

◈ Phase 3: Build

- Launch **AI Squads** — short-lived, sprint-based teams with clear scope and enablement.

- Build lightweight AI MVPs using tools like LangChain, Streamlit, vector DBs, and cloud functions.

- Create structured feedback loops with support/sales/customer-facing teams.

- Define **Exit Criteria**: what qualifies a pilot to graduate to full roadmap integration.

30-60-90 Day Execution Plan (The Core Blueprint)

This section breaks down the entire journey across 3 months with tangible actions, rituals, artifacts, and outcomes.

▦ Days 0–30: ENABLE

Objective: Prepare teams with tools, knowledge, and space to begin the journey.

- Run a 1-week asynchronous AI fluency self-assessment

- Conduct interviews across product, design, engineering, and QA to identify blockers

- Provision sandbox access with approved APIs (OpenAI, Hugging Face) and trackable logs

- Launch the Internal AI Hub: past experiment writeups, prompt templates, FAQs, Slack help threads

- Host live workshops: AI Basics for Product Teams, LLM 101 for Developers, Ethics in AI for PMs

- Output: Baseline skill maps, first round of internal success stories, tooling cheat sheets

🖼 Days 31–60: ALIGN

Objective: Surface high-leverage use cases and match them to real business priorities.

- Run design thinking-led ideation sprints: "How might we use AI to ___?"
- Introduce AI Opportunity Canvas: Teams fill this collaboratively with PM/Eng/Design
- Shortlist use cases by ROI, feasibility, and speed-to-pilot
- Involve legal, security, and infra teams to pre-approve data pathways
- Publish a use case repository with ownership, next steps, and risk grading
- Output: 3–5 validated AI use cases with team buy-in, no-go zones mapped, success metrics defined

🖼 Days 61–90: BUILD

Objective: Convert validated ideas into working pilots with measurable impact

- Form 2–3 dedicated AI squads with leads, goals, and defined sprint charters
- Enable building via LangChain, Streamlit, Gradio, FAISS, Azure/AWS functions
- Embed pilot reviews into demo days; invite sales/support for usability feedback
- Track metrics: time saved, effort reduced, end-user feedback, accuracy/relevance
- Run pilot retrospectives: What worked, blockers, what we'd productize vs archive
- Output: Working demos, go/no-go productization reports, updated team capability heatmaps

Success Metrics

Adoption & Enablement

- % of teams with full access to AI environments
- Completion rate of AI bootcamp or onboarding

Execution Health

- of AI pilots launched within 90 days
- Pilot quality score (measured via checklist: performance, safety, design quality)

Business Alignment

- % of pilots tied to strategic goals
- Executive sponsorship index (measured through engagement, approvals, demo attendance)

Risks & Trade-Offs

Risk	Mitigation
AI fatigue or scepticism from previous failed experiments	Share early wins weekly, make efforts visible with storytelling
Legal friction around customer data usage	Involve privacy/security early; build clear guardrails into sandbox
Understaffed pilots	Use guilds or chapters to provide part-time support
Budget freeze after initial demos	Build clear ROI narratives in every pilot report
Shadow experimentation	Create Slack AI channel + working groups to share progress openly

Try This (Interactive Simulation)

Imagine you're leading this AI scale-up sprint at your company.

Scenario Challenge: You've been asked to brief your executive team and secure buy-in. Build a 3-slide deck:

1. Top 3 blockers to AI momentum in your company today

2. Your 30-60-90-day rollout plan based on Enable → Align → Build

3. One high-ROI pilot you could greenlight immediately (what it solves, how fast it can be built, who it helps)

Let your inner strategist shine — and make AI implementation feel real, fast, and meaningful.

Thank you

Happy Learning!