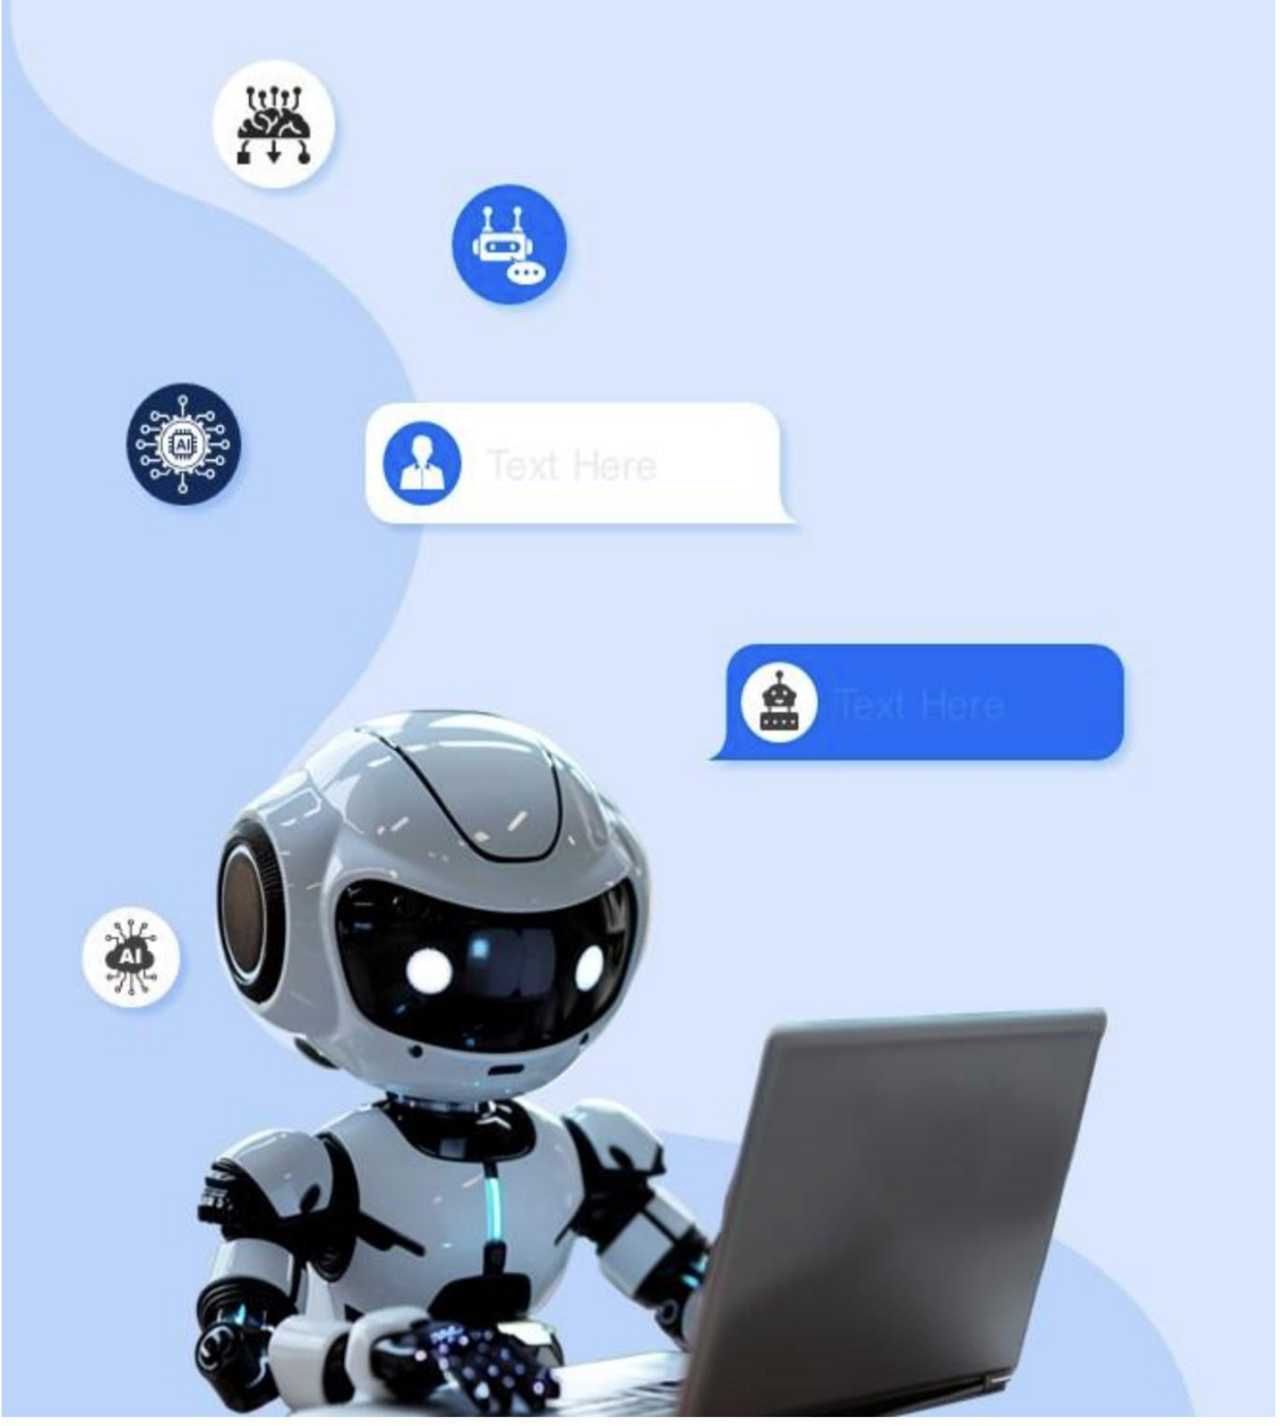

Lessons From Building an AI Chatbot

From NLP to Agentic AI

Over the past several years, I've worked on global AI support ecosystem—building automation capabilities that handle millions of customer conversations every quarter. Today, I'll share how our chatbot evolved, what we learnt, and how AI is reshaping product management.

A long, long time ago (2019) in a galaxy far away…

"We need help! There's too much volume for us to handle!"

Time to call Jarvis!

Jarvis Bat signal!

Let’s get down to business! (to defeat the handoffs)

2019: We started with a simple bot identifying the Cancel intent for Individual customers only.

All requests were connected to agents.

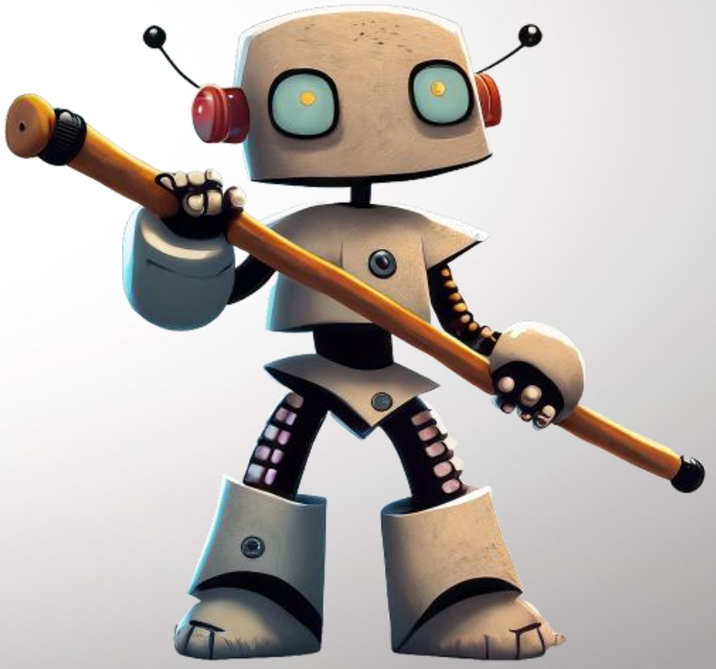

Design Thinking: The Backbone of Our Journey

When we started building Adobe's chatbot, it wasn't an AI problem—it was a human problem. Design Thinking shaped everything we built.

Empathise

Listened to thousands of conversations, studied drop-offs. Customers needed clarity, confidence, and control—not just a chatbot.

Test

Validated with real customers, tracked metrics, and iterated. Testing → learning → iteration became our operating rhythm.

Define

Reframed problems: Users don't trust responses, escalate when not understood, and need resolution, not just answers.

Ideate

Brainstormed bold possibilities: multi-turn reasoning, action capability, contextual responses, and empathetic assistance.

Prototype

Built and tested conversational flows, prompts, reasoning patterns, and automation experiences through limited experiments.

Design Thinking ensured our AI chatbot wasn't just intelligent—it was human-centred, emotionally aware, and built to genuinely solve customer problems.

The Challenge: Supporting Millions at Scale

Global Reach

Millions of users across dozens of products in 190+ countries requiring 24×7 support.

Complex Queries

Billing, cancellations, downloads, errors, how-tos, account issues, and licensing support needed.

The Mission

Make support intelligent, conversational, proactive, and capable of real resolution.

Traditional help channels worked, but they were reactive. Customers couldn't find answers easily, NLP bots hit their ceiling, and support was generic rather than personalised.

	Intents	Target Audience	Languages	ARR
2019	Retention		ENG	$0
2025	80 sub-intents, 200+ products  Adobe ID Download   Sales 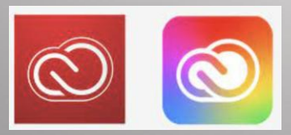  Billing	Individual Teams Enterprise Entitled, Noentitled, Unauth customers + more	English French German Japanese Korean Portuguese Italian Simplified-Chinese Traditional-Chinese	$320M

& more...

Desktop and Mobile app integration

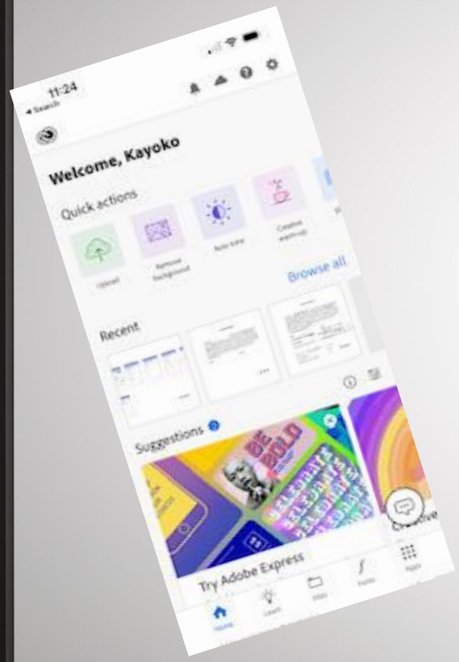

Product / Event specific Bot

Bot Authoring

Personalization

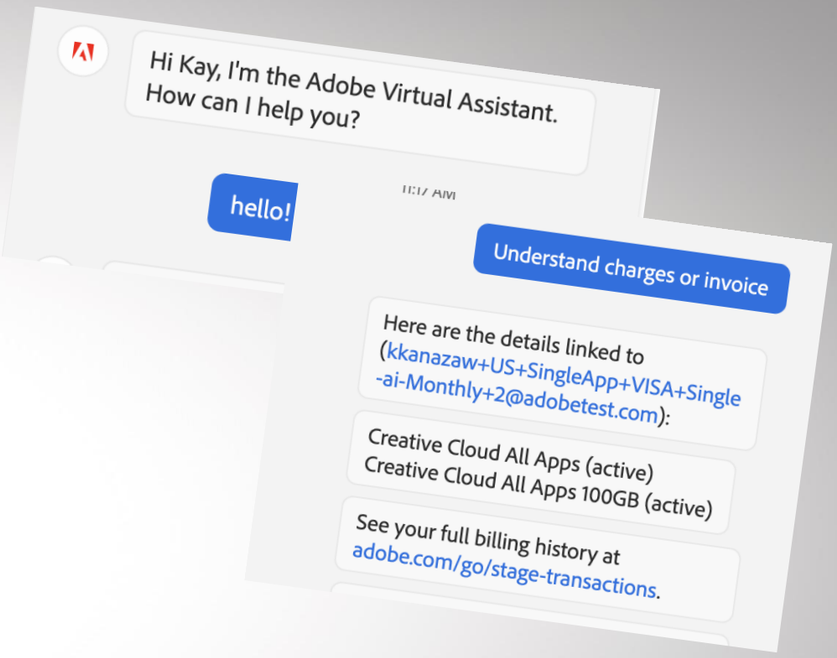

Omni-channel experience

Engagement/sales experiences

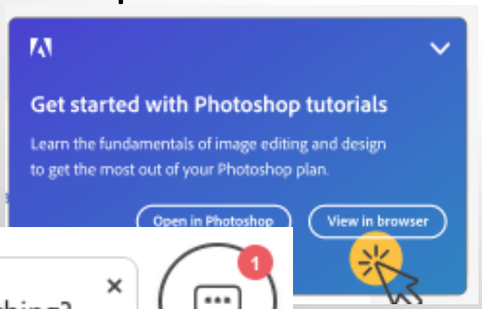

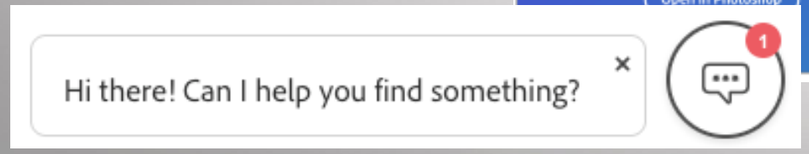

Auto-complete

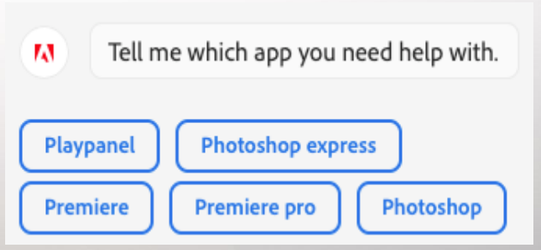

Automated answers with AI

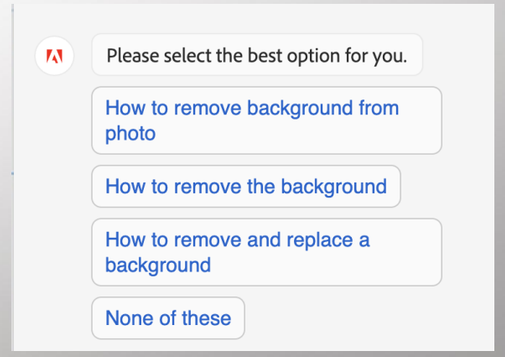

Phase 1: NLP Chatbot

We started with an NLP chatbot built on RASA. It was intent-based and worked well for clear questions, routing users correctly and deflecting some volume.

The NLP Ceiling

Constraints quickly emerged:

- Couldn't handle ambiguous questions
- Broke on multi-turn conversations
- Required constant training data tuning
- Could only answer, not act

Accuracy plateaued, automation stalled, and customer frustration rose. The next improvement needed a step-change in capability.

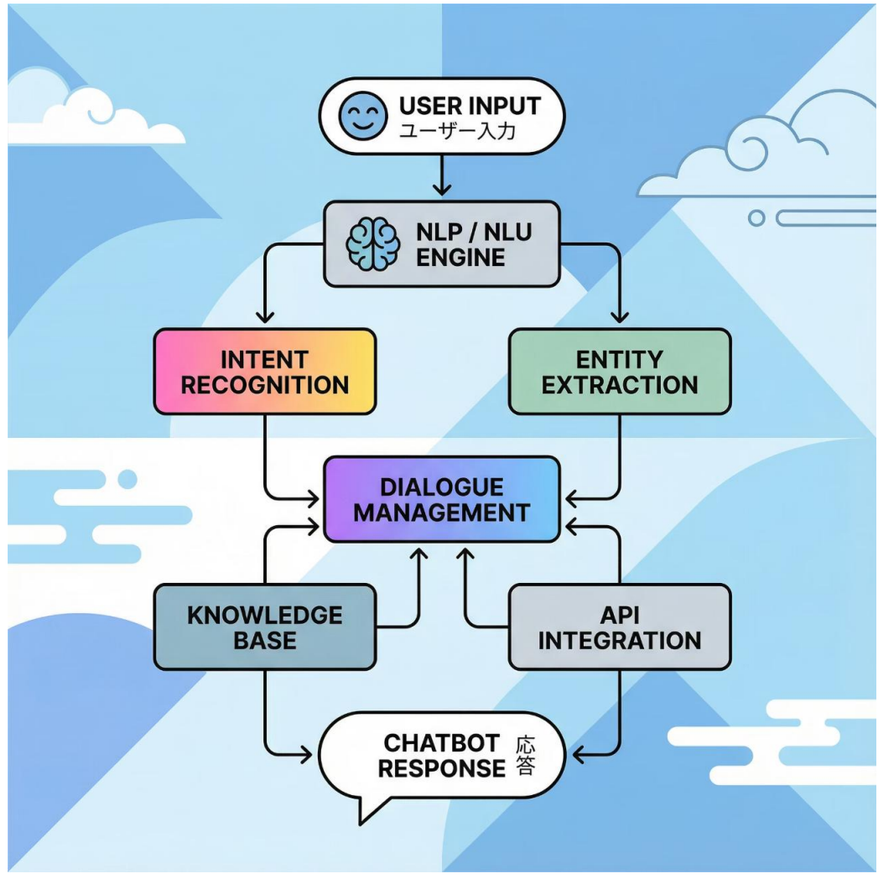

Phase 2: GenAI Transformation

We reimagined the chatbot as a conversational, reasoning-based assistant with three major upgrades:

01

Multi-turn Natural Conversations

GenAI removed strict intent mapping. Customers could speak naturally, and the AI could understand context, ask clarifying questions, retrieve solutions, and adjust based on replies.

02

Knowledge-grounded Responses

Genie as the Knowledge Agent pulls accurate information from updated help content, generates personalised explanations, cites verified sources, and scales across products.

03

Agentic Workflows

The bot can now complete tasks: cancel trials, update payment methods, help change plans, guide installations, identify charge issues, and trigger account actions safely.

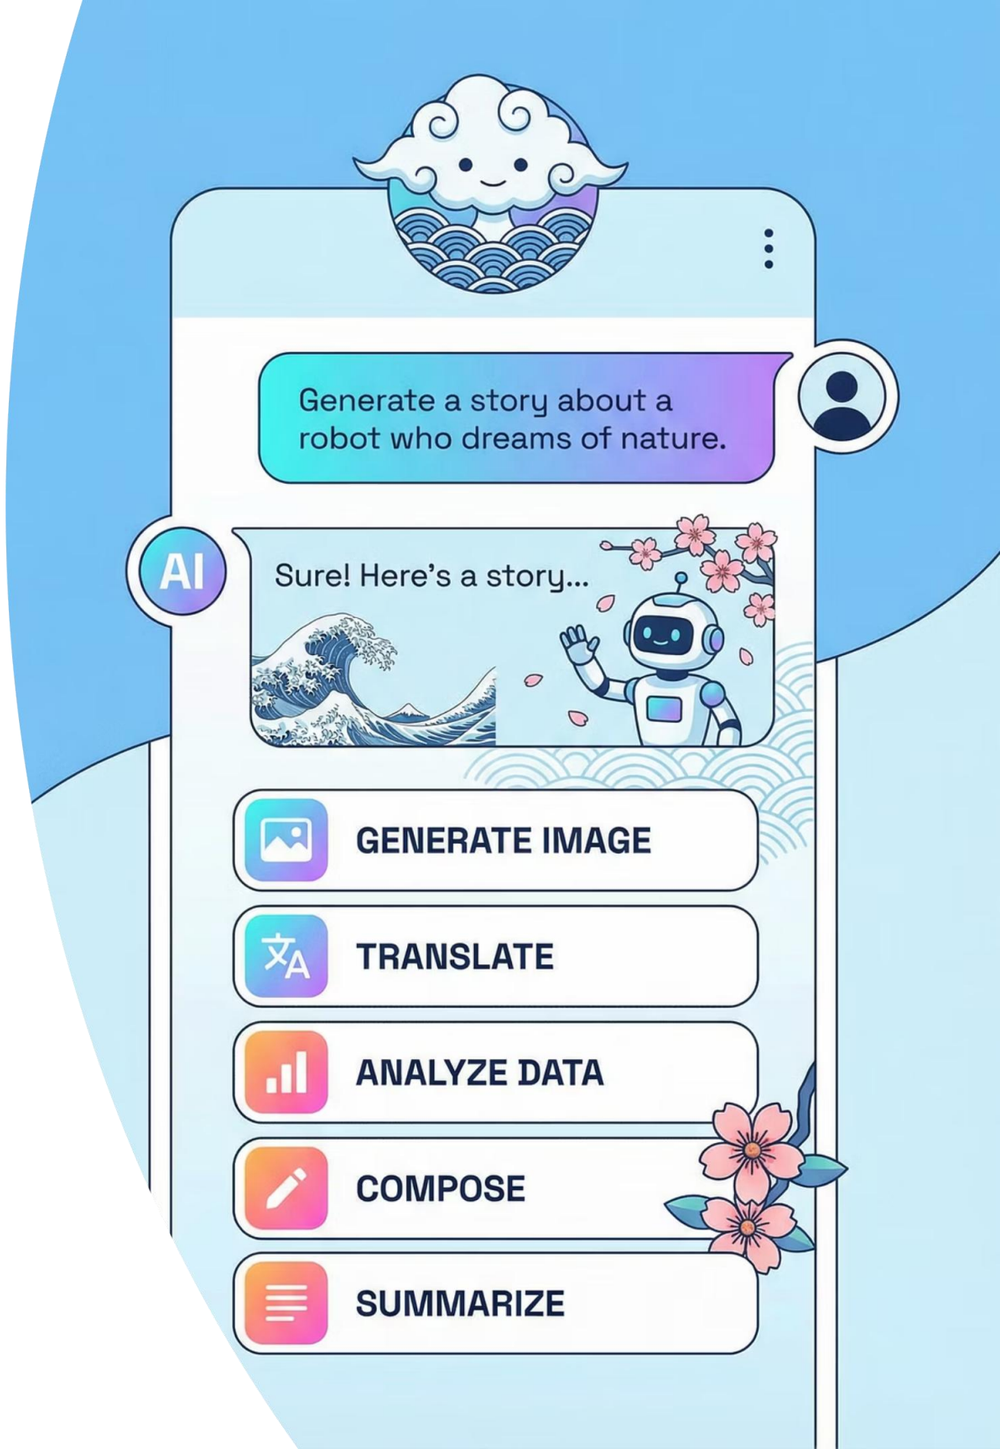

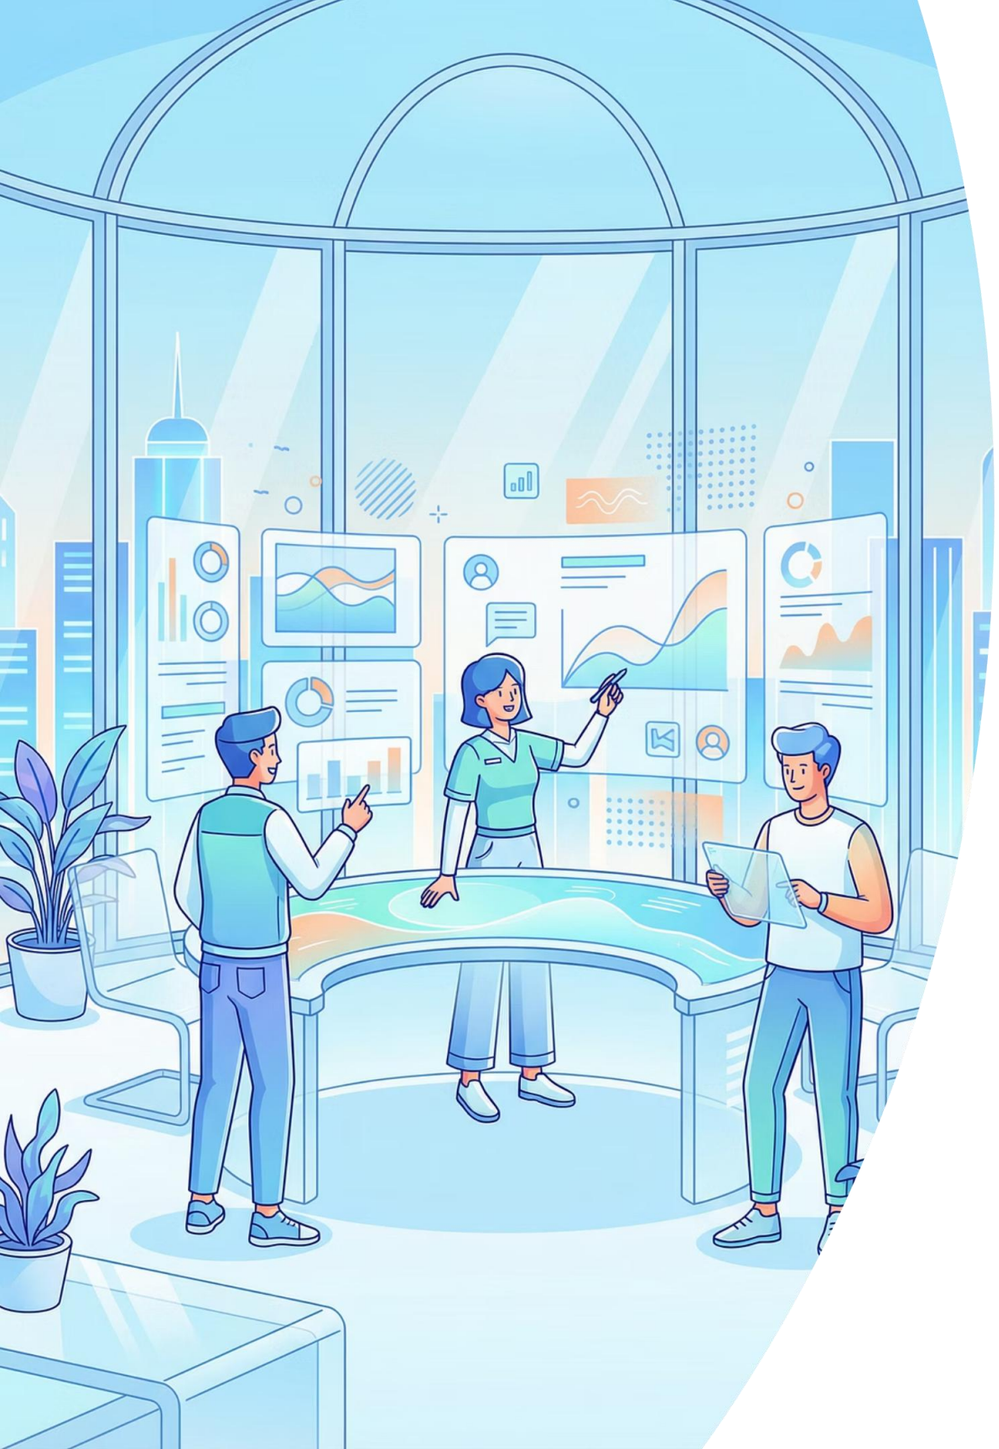

Key Lessons for Product Leaders

These lessons come directly from our automation strategy, rollout guidance, quality frameworks, and PM insights—simplified for practical application.

Foundational Lessons

AI Succeeds Only When It Solves Real Customer Problems

We didn't start with "let's add GenAI." We started with: Why are customers cancelling? Why are they calling agents for self-serve issues? Where does the article experience fail? Problem framing is where PMs make the biggest impact.

Guardrails Are More Important Than the Model

Safety filters, context blocks, high-risk detection, routing rules, personalisation logic, feedback classification, and quality evaluation pipelines make the experience enterprise-ready. This is what makes AI reliable at scale.

Operational Excellence

1

Start Small, Then Scale

5% rollout → Monitor → 25% rollout → Fix issues → 50% → A/B compare → 100% → Stabilise & expand. Track automation rate, resolution rate, transfer rate, CSAT, conversation quality, and failure modes.

2

Design for Systems, Not Just Features

PMs working with AI must think like systems architects: data flow, content readiness, API availability, personalisation models, cross-channel continuity, and experience consistency.

3

AI Changes User Expectations

Customers now expect conversational help, personalised answers, predictive guidance, instant actions, and seamless escalation. Every GenAI improvement raises the bar.

The Future: From Chatbot to Virtual Adviser

Our strategy for the chat bot

- A unified, intelligent support ecosystem
- AI that is proactive, not reactive
- Personalisation for every customer
- End-to-end resolution, not deflection
- Agentic assistants working across channels

The real transformation isn't that AI can talk naturally. It's that AI can finally solve real customer problems. As product leaders, our job is to design intelligent systems—grounded in empathy, powered by data, guided by guardrails, and measured by outcomes.