# The Last Mile of AI

*Shipping AI Inside Real Organisations*

---

Around 95% of corporate AI pilots never reach production. Four MBA 26Ds who crossed that mile before business school share what it actually took.

Join us for an honest account of the last mile: the challenges, the decisions, and what it means for the organisations you will go on to lead.

**Friday, 5 June 2026**
**3.00 to 4.00pm**
INSEAD Fontainebleau Campus, Amphi de Vitry

*Brought to you by the INSEAD AI Club*

Justin Foo
*Ex-URA*

Sanya Pawah
*ex-CHAI*

Nick Hauck
*Ex-Audi*

Saurabh Bhatia
*Ex-Avataar.ai*

# “It is just a toy.”

## Notes from trying to ship generative AI inside Singapore's planning authority

**Justin Foo**
INSEAD MBA 26D
Ex-Urban Redevelopment Authority Singapore (URA)

# “This is not a useful tool. It is just a toy.”

**A senior leader at URA, on Dream Lab.**
**Six months before we launched.**

He refused to let Singapore's largest greenfield development be included as part of our initiative.

The thing

# Dream Lab.

Singapore's first government deployment of generative AI for urban planning.

Ordinary people describe the future they want for a real Singapore site, and the tool turns their words into vision statements and images, live.

It replaced the feedback forms and the wall of sticky notes at the Draft Master Plan 2025 Exhibition, Singapore's mandated five-yearly consultation.

Still live at URA. Twenty thousand sessions and counting.

*The version that shipped almost never did.*


Dream Lab on display at the URA.

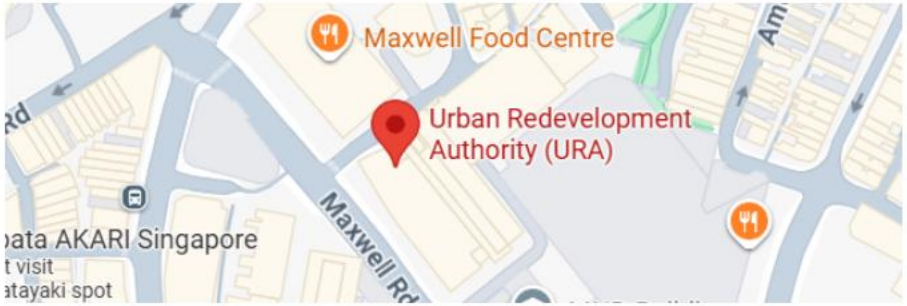


Urban Redevelopment Authority, Singapore.

Source: URA, Draft Master Plan 2025

Before we get back to that gatekeeper

# Five storms. Five lessons.

None of these came from a strategy deck. They came from the friction.

01

**The skeptic**

02

**No budget, no team**

03

**The letter**

04

**The lawsuit**

05

**The safe answer**

Storm 02

# No budget, no team.

## “There is no budget for this. And it is not your job anyway.”

Dream Lab was not what I was hired for. I was an architect, was supposed to focus on the urban design of Jurong Lake District. There was no money to build it, and I could not pull software engineers in full-time. Their managers saw no glory in helping another team win.

So we found a cheaper sandbox. We mentored a university capstone team who pressure-tested the idea for a semester and kept full ownership of what they built. We learnt from their experience, chained off-the-shelf tools together and ran the rest ourselves. The validation cost us almost nothing.

Lesson **Borrow the sandbox, keep the lessons.**

Storm 03

# The letter.

## “Our sister agency just got publicly attacked for using AI. Are you sure about this?”

In January 2025, while we were piloting, 68 Singaporean writers signed a public statement attacking the National Library's uncritical embrace of generative AI. Their fear was displacement: AI sold as a substitute for human writers and illustrators. My management got nervous.

What saved us was framing. We did not receive a single complaint. This was because we never sold Dream Lab as a replacement. We framed it as empowerment, giving people who cannot draw a way to picture their own neighbourhood's future. It adds a voice. It does not remove a job.

Lesson **Ask “who loses if this works?”**

Source: 68 Singaporean writers, January 2025

Storm 04

# The lawsuit.

## “Disney is suing Midjourney. What happens to us?”

In June 2025, weeks before our launch, Disney and Universal sued AI firm Midjourney over copyright. Midjourney was our entire image pipeline, and the lawsuit was out of our hands. So we did two things.

First, we swapped to another model in days, because the model is a commodity: anyone can rent it, anyone can be sued over it. Second, we sat down with Legal, clarified exactly where URA stood on copyright, and walked the CEO through every caveat until he could put his name to it. What you own is your data and your legal ground. The model is just rented.

Lesson **Don't wish the risk away. Govern it.**

Source: Disney & Universal v. Midjourney, June 2025

Storm 05

# The safe answer.

## “Just pre-generate the images. Curate them. Deploy what is safe.”

Leadership proposed removing the live generation: safe, yes, but pointless. And it carried its own risk. If it had ever come out that the public's images were pre-curated, the trust damage would have been worse than any inappropriate output. An engagement tool that fakes it is not engagement.

So we kept it live and drove the risk down in three layers. A match-based filter caught the obvious, an AI filter caught the subtle, and a human could pull anything live from a dashboard. No single layer is enough, so we stacked three. The guardrails were the brakes that let us ship.

Lesson

**Trust is your most valuable resource. Protect it.**

# And the senior leader?

We never converted him, so we stopped trying. You cannot demo your way past a person who has already decided. We went around him, recruiting sites from teams who were already curious, until one refusal stopped mattering.

Then we went over him, all the way to the Prime Minister showing the public's visions at the National Day Rally. Singapore's largest new development is still missing from Dream Lab. Technology does not solve a people problem. People do.


National Day Rally 2025. The Prime Minister presents public visions generated through Dream Lab.

Lesson **The gatekeepers set your ceiling.**

Source: National Day Rally 2025

What you might take with you

# Five lessons. All from the friction.

01 Borrow the sandbox, keep the lessons.

02 Ask “who loses if this works?”

03 Don't wish the risk away. Govern it.

04 Trust is your most valuable resource. Protect it.

05 The gatekeepers set your ceiling.

*None of these were planned. They were what was left after each near-miss.*

If you remember one thing

# The last mile of AI is not a technical problem. It is a series of conversations with people who have very good reasons to say no.

The job is to keep showing up.

---

## Thank you

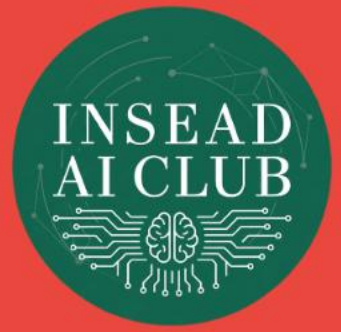
INSEAD
AI CLUB

# Tuberculosis never left India.

Notes from TB Active Case Finding in India.

**Sanya**
INSEAD MBA '26
Ex-Clinton Health Access Initiative

# Tuberculosis never left India.
# COVID made it deadlier.

While the world watched COVID, the world's oldest pandemic surged in its shadow - under-diagnosed, under-treated, and turning drug-resistant.

Symptoms missed › Diagnosis delayed › Drug-resistant TB (MDR) › **Death**

## THE COST OF A MISSED CHEST X-RAY

Oxygen, ferried by auto-rickshaw

A TB patient, wasting away

Mass cremations, 2021

## THE NUMBERS BEHIND THE CRISIS

**2.7M** TB patients in India - 25% of the entire global burden

**42.6%** of TB confirmations would be missed without a chest X-ray

**~15** radiologists per million people to read those X-rays

**63%** with chest symptoms never sought any care

**~55%** of secondary public facilities had a working X-ray

Sources: Global TB Report 2025 . National TB Prevalence Survey 2019-21 . IRIA 2024 . Rural Health Statistics 2022-23

# You can't train your way out of this.

Why AI was the only answer that scaled

### Radiologists don't scale.

~15 per million, years to train - and they don't relocate to villages. You cannot manufacture the specialist fast enough.

### Symptoms miss half the cases.

42.6% of TB would be missed on symptoms alone, and 63% never seek care. Waiting for a cough is waiting too long.

### The bottleneck is reading, not imaging.

An X-ray is useless if no one can interpret it - instantly, on-site, with no internet. That is exactly what AI does.

*The move: turn a scarce specialist skill into a commodity that travels to the doorstep.*

WHAT WE BUILT

## Diagnosis, delivered house-to-house.

**2,000+ ultraportable X-ray devices**
Battery-powered, hand-carried, sourced from South Korea.

**Qure.ai qXR - WHO-endorsed CAD**
AI reads every scan offline in ~30 seconds. No radiologist needed.

**Active Case Finding camps**
We brought the lab to the last village - door to door.

**Integrated screening**
TB + blood pressure, sugar and BMI, all in one visit.

| 1.5M+ | 18,855 | 97% |
|---|---|---|
| screened | camps | got an X-ray |

02 . THE PIVOT

# You can't train your way out of this.

# Nine times fewer tests for every case found.

Number needed to test - people tested to confirm one TB case

WITH AI : NUMBER OF TESTS TO GET 1 CASE

**16**

**AI-suggestive + symptomatic**

6% test positivity

SYMPTOMS ONLY : NUMBER OF TESTS TO GET 1 CASE

**148**

**the old, screen-on-cough way**

0.7% test positivity

**9×**

**leaner**

Every confirmed TB case found with ~9× fewer expensive molecular tests - AI doesn't just find more, it wastes far less.


HOW THE ALGORITHM GETS THERE - THE SCREENING CASCADE

Systematic screen → Ultraportable X-ray → qXR AI score (0 - 1) → **Above threshold?** → Sputum → NAAT test → TB confirmed & treated

***Decision rule:*** *anyone symptomatic, vulnerable, or AI-suggestive of TB advances to molecular testing - the AI casts a wider, smarter net than symptoms ever could.*

03 . THE RESULT

# Nine times fewer tests for every case found.

# The model was the easy part. The threshold was the last mile.

## A SCORE IS NOT A DIAGNOSIS

qXR outputs a TB-risk score from 0 to 1 - not a yes/no. A single number, the threshold, decides who gets a scarce, expensive NAAT test. WHO is explicit: the right threshold differs by population. The vendor default is rarely optimal.


*Set it too low → NAATs wasted (cost). **Too high → TB missed (lives).***

**ITSC**

**Iterative Threshold Selection Calibration**

1. Start at a WHO 90%-sensitivity (or vendor) threshold
2. Pick n - the batch size before each test (e.g. 100 scans)
3. Run n X-rays → AI score → give NAAT to everyone in the batch
4. Measure: unnecessary NAATs vs. additional TB missed
5. Update the threshold and repeat - stop early when optimal

## THREE WAYS TO PICK A THRESHOLD - WE CHOSE THE THIRD

### Set & hope

+ Easy to implement
- Non-statistical - leaves cases and money on the table

### Fixed WHO trial

+ True, validated performance
- Expensive, slow, needs human reads of every scan

### ITSC

+ Statistically informed . can stop early
- Some implementation complexity - worth it

# The last mile wasn't the model.

~95% of corporate AI pilots never reach production. Ours did - and the AI worked on day one. What almost stopped us was everything around it.

01

**Offline, or nothing.**

No connectivity in the last village - the AI had to run on a laptop in a field tent.

02

**The threshold, not the model.**

A risk score is worthless until it's calibrated to the people in front of you.

03

**AI finds; people confirm.**

An ASHA's knock and a medical officer's trust moved more than any algorithm.

04

**The bottleneck just moves.**

AI found so many cases we ran short of NAAT capacity. Fix the system, not the step.

> *The last mile of AI in healthcare isn't a technical problem. It's a series of conversations - and calibrations - with the people and systems that have to act on what the model says.*

**Thank you.**

Sanya . INSEAD MBA '26 . ex-Clinton Health Access Initiative

# The last mile wasn't the model.

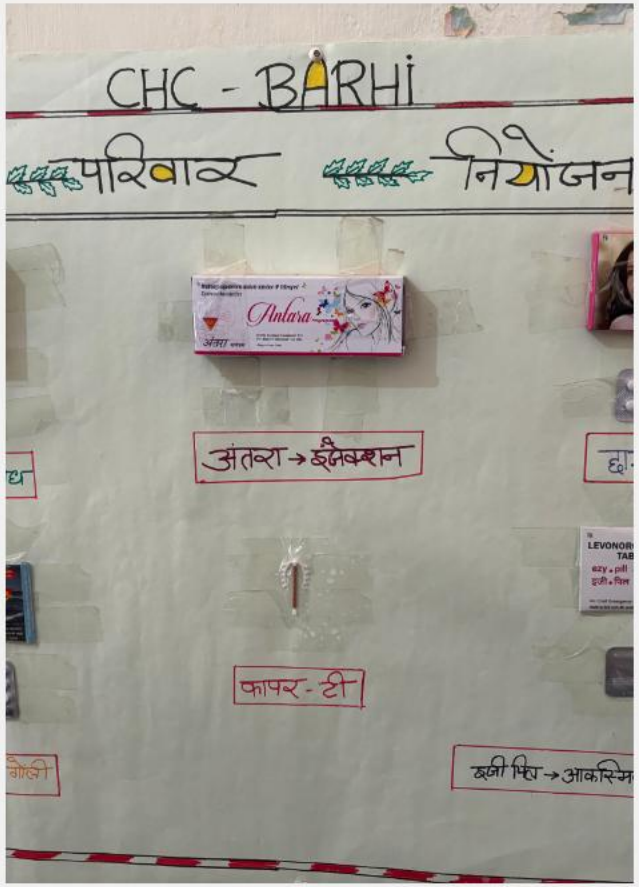



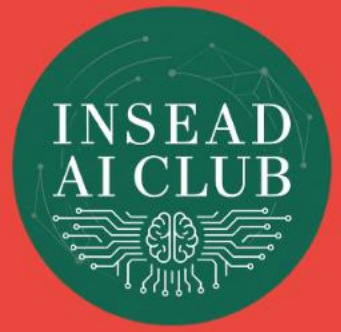
INSEAD
AI CLUB

# AI Strategy Gaps: The Use Case Trap.

**Nick Hauck**
05 June 2026

WHY

# AI as key driver in context of manufacturing.

Automotive in global context.

Automotive Costs
**China**

- Labor Cost: $585 per Car
- Development Cost: 1.0x
- Production Cost: 1.0x

Automotive Costs
**Germany**

- Labor Cost: $3,300 per Car
- Development Cost: 4.0x
- Production Cost: 5.6x

→ Cost advantage of $10,000 for Chinese producers
→ 275,000 - 1,000,000 jobs at risk in Europe
→ AI is main tool to close gap by driving productivity

| 140 Factories | ~ 9 Mio Cars | 700k Employees | Everything becomes a Scaling Problem |
|---|---|---|---|

WHAT

# Fields of interest.

Decreasing the Gap of 5.6x in Production.

| | | | | |
|---|---|---|---|---|
| User | GenAI | Industrial Computer Vision | Advanced Robotics | AI Machinery |
| Frontend | | | | |
| Algorithm | | | | |
| MLOps | GAPS Through Diverse Brownfield | | | |
| Storage | | | | |
| Connectivity | | | | |
| Prod. Line | | | | |

→ The Applications in the End are always in focus, but they rely on the backend

→ If the backend is not developed strategically there is no efficiency win

HOW

# Brute Force without Strategy.

Use Case minded Technology driven.

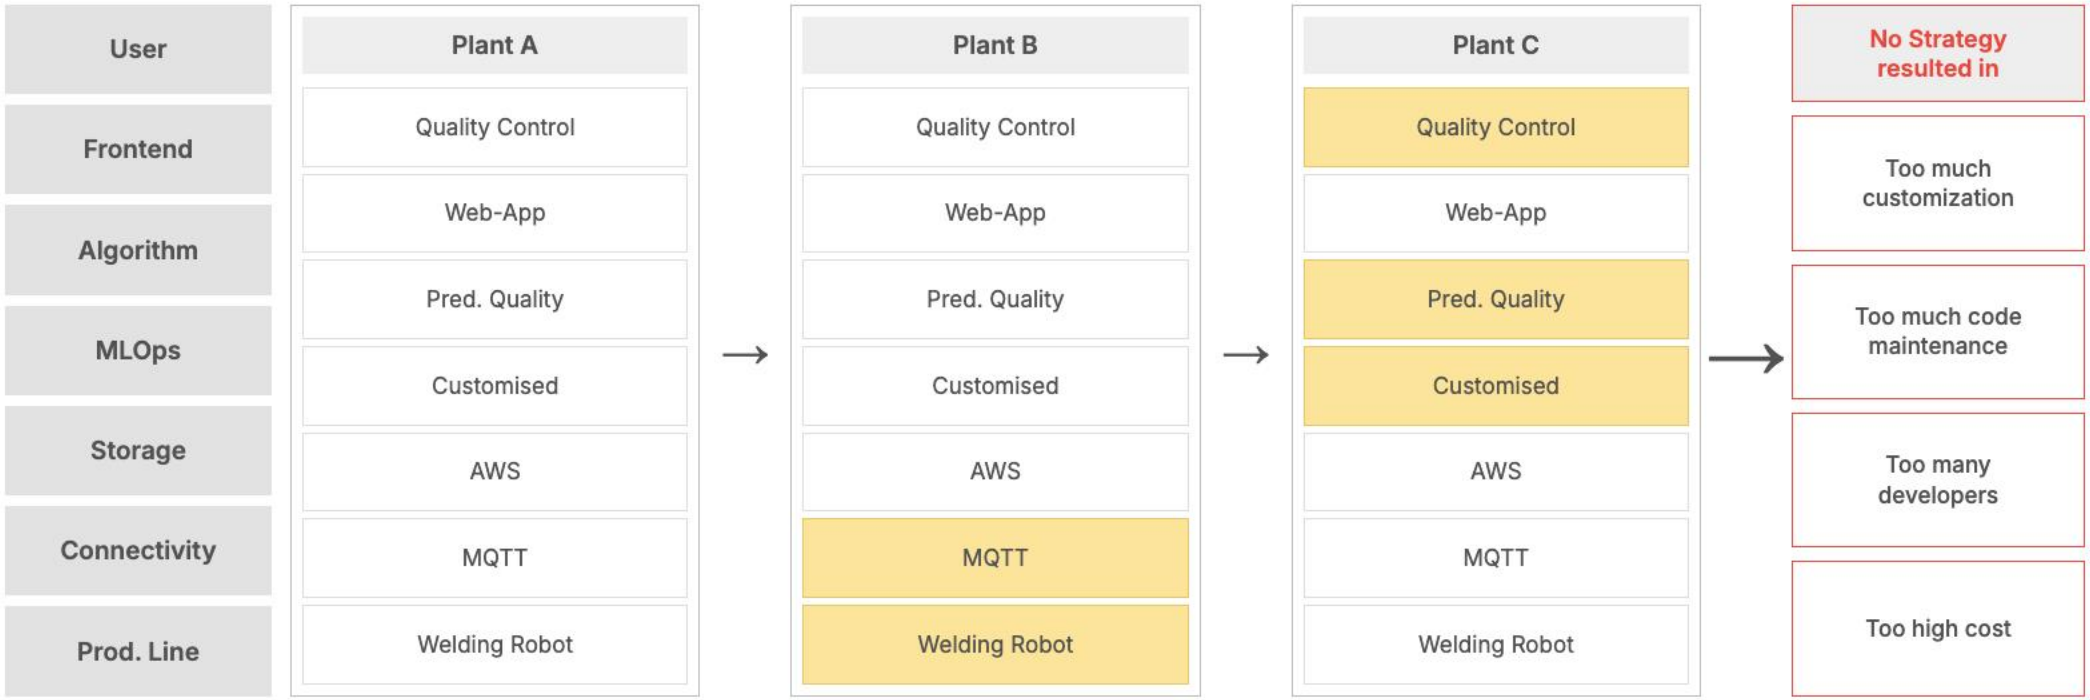


→ Without comprehensive vision and strategy scale up failed

→ Customers are in total equal but totally different in detail

HOW

# Obstacles we faced.

Millions wasted, years lost: here is what blocked us.

## 01 | Organizational Fragmentation

- Plants could opt out of adoption.
- Without mandatory buy-in, cost could not scale.
- Every solution stayed use-case specific and non-standardised.

**Root cause:** Brand-level process standardisation was never mandated, so each plant claimed to be "different".

## 02 | Too much Customer Focus

- Our most advanced use cases were technology driven.
- We self-funded development to prove the ROI.
- Targeted rollout to 8+ European plants.
- Result: blocked by plant-specific requirements at every site.

**Scaled to 3+1 plants only:** The "+1" developed its own custom solution on our blueprint.

## 03 | Cultural Resistance

- Manufacturing management preferred continuous improvement over disruption.
- They downplayed AI efficiency gains to their own engineers to avoid taking on transition risk.

**Key dynamic:** Engineers saw the value. Their managers blocked it.

→ Strategy gap is not a tech problem: it is an organization & business problem.

→ No standardisation mandate = no scale

HOW

# Value Driven with Strategy.

Target Overview & Current Situation.

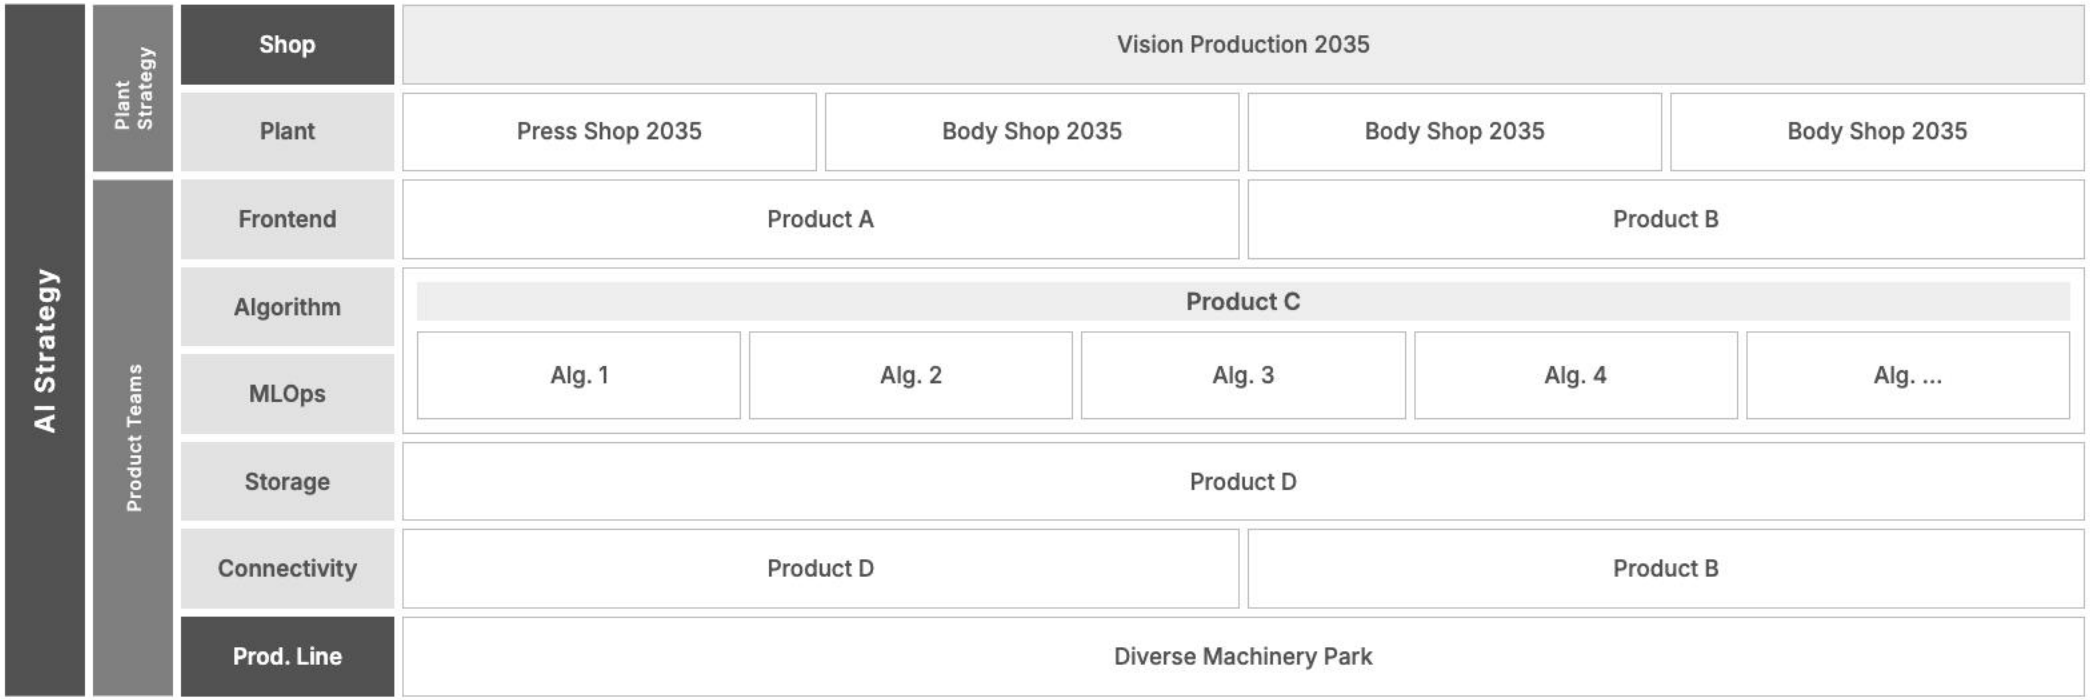


→ Clear AI Strategy enables identify future needs

→ Enables process and product standardization to scale up

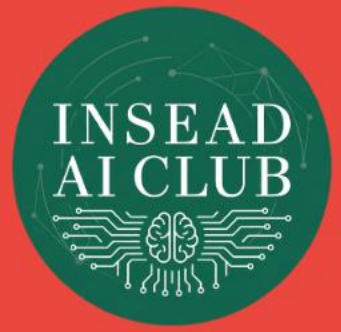
INSEAD
AI CLUB

# My last mile was a legal one.

Notes from shipping AI voice agents into US debt collection.

**Saurabh Bhatia**
INSEAD MBA 26D
Ex-Skit.ai (startup)

# 12,500 calls.
# Every one of them illegal.

Our AI voice bot. One run. $500,000 of exposure - against us, and against our client.

*Hold that thought.*

WHAT IT WAS

# Skit.ai

AI voice agents for US debt collection.

**Not the easy kind of AI.** The kind that calls a stranger who owes a bank money, in the US, and has a real conversation about it - verifying who they are, why it is calling, and taking a payment on the line.

*By every metric you'd put on a slide, the AI worked.*

**20+** banking & insurance clients

**1M+** automated calls every day

**$10M** in annual recurring revenue

THE REALISATION

# A demo and a product are not the same thing.

IN A DEMO

You control everything. One call. Clean audio. A friendly tester who wants it to work. It's a magic trick - and the magic trick is easy.

IN PRODUCTION

A million calls means a million edge cases a day. And in our industry, every single call is a legal event. The demo proves the AI can talk; production asks whether you can talk to a million strangers without breaking the law once.

***That question is where 95% of pilots die.*** *Not because the model isn't smart enough - because nobody costed the last mile.*

THE MAZE

# Five walls.

Debt collection is one of the most regulated conversations in commerce.

| | | |
|---|---|---|
| 01 | **Timing** | No calls before 8am or after 9pm - in the customer's local time zone. The bot must place every number in the right zone, in real time. |
| 02 | **Frequency** | No more than 7 calls about a debt in 7 days. The moment you actually speak to someone, a new 7-day clock starts. Per debt, not per person. |
| 03 | **Identity** | You can't reveal a debt to the wrong human. Before saying why it's calling, the bot verifies the name, then the last 4 of the SSN or date of birth. |
| 04 | **Payment** | The bot can take a payment on the call - but cannot store the card. The compliance line runs through the most valuable moment. |
| 05 | **Geography** | State rules stack on top of the federal ones. Where a state is stricter, the state wins. |

SO WHAT IS THE LAST MILE?

# The model was the easy 80%. This 20% decided whether we had a business.

Timing. Frequency. Identity. Payment. Geography. All of it resolving correctly, automatically, a million times a day, in a live conversation, with no human in the loop.

# Make the rule impossible to break.

### TIME ZONES

**Naive build:** dial when the bot is ready.
**Compliant build:** refuse to dial until the number is mapped to a zone and checked against the 8-9 window - and the stricter state rule if one exists. A hard gate the dialer cannot pass.

### RIGHT-PARTY IDENTITY

**Verification is a locked sequence:** the bot physically cannot disclose why it's calling until name + SSN-last-4 or DOB clears.
Not a guideline the model is asked to follow - a wall in the flow it cannot walk through.

*In a demo, compliance is something you tell the model to do.* ***In production, it's something the system makes impossible to violate.***

THE TURN

# The barrier was the moat.

**Friction** → **Barrier to entry** → **Moat**

Every rule made our entry harder - longer build, slower launch, more cost per call. But that same friction is what stops a generic voice-AI startup from walking in and undercutting us. **Anyone can rent a model. Almost no one will spend 18 months turning it into something a regulated bank will let near its customers.**

*Compliance isn't the tax you pay to ship. **It's the product you're selling - and the wall you build behind you.***

THAT NIGHT

# The wall was there. Someone climbed over it.

**The cause wasn't the AI.** A rushed code deploy, pushed under sales pressure, silently overrode one of those compliance gates - and fired 12,500 calls in restricted hours.

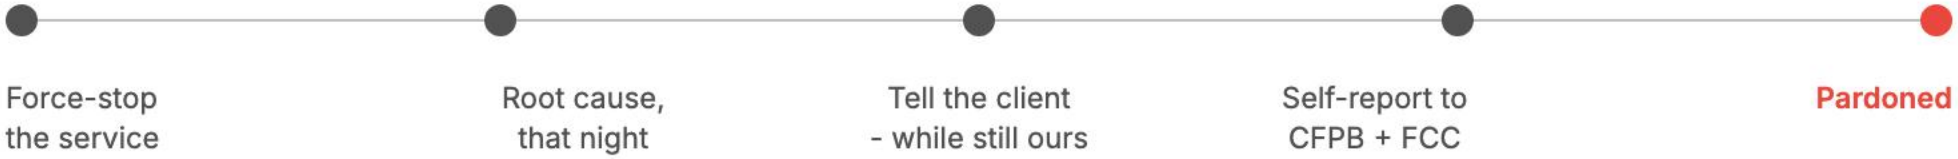


*We almost lost the whole business in one deploy that bypassed a guardrail.* ***We kept it with one honest phone call.***

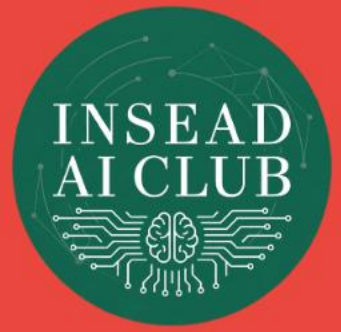
INSEAD
AI CLUB