

USE CASE DISCOVERY + AI READINESS · ILLUSTRATIVE OUTPUT

One pilot, picked. The journey, mapped.

A sample output. Mid-market NBFC. Three vendor proposals on the table. Three stalled POCs behind them. One AI mandate from the board.

This is what three weeks of Use Case Discovery and AI Readiness produced. The pilot to ship now. The journey ahead.

READ TIME

5 min

PAGES

12

ENGAGEMENT

Explore

NEXT MOVE

Prove

Numbers are directional. Names removed.

01 · CONTEXT

The board wanted a decision. The CRO wanted a **pattern.**

INDUSTRY

Mid-market NBFC. ~₹15,000 Cr AUM. Retail lending: two-wheeler, personal, MSME. 3,200 employees across 14 states.

OPERATING ENVIRONMENT

22+ internal systems. LOS, LMS, two bureau feeds, two scorecards, separate DSA portals, a homegrown branch CRM, and an HRMS that doesn't talk to sales reporting. No single view of the RM.

WHAT LEADERSHIP CAME IN WITH

"The board wants AI in production this year. Three POCs. None cleared risk and compliance. We don't know if it's a data problem, a governance problem, or a vendor problem."

WHY DISCOVERY WAS NECESSARY

Three vendor proposals were already in. An underwriting copilot. A disbursement-bottleneck dashboard. A DSA scoring upgrade. None had a regulator-mapped governance plan. None had a P&L hook the CFO would sign.

SCOPE

Three weeks. Use Case Discovery and AI Readiness, run together. Sessions with distribution, branch leadership, credit, ops, IT, risk, HR, compliance. Workflow mapping across the RM, sanction, and disbursement lifecycles. Data audit across six systems. Governance gap map against RBI Master Directions, RBI FREE-AI, Fair Practices Code, DPDP.

2,000

Field RMs across branches + DSAs. ~65% of book is push-sourced.

14 mo

Since the last AI proposal cleared risk and compliance

1:8

Manager-to-RM coverage ratio today. Realistically 1:3 in practice.

₹2.4 Cr

Spent on three stalled POCs. Nothing in production.

02 · PRIORITISED USE CASE STACK

Ranked on four lenses.

Business impact. Feasibility. Time to ROI. Regulatory clearance. The fourth is non-negotiable in BFSI. A use case that won't clear risk and compliance isn't a use case.

USE CASE	IMPACT	FEASIBILITY	TIME TO ROI	REGULATORY	DECISION
RM productivity & coaching AI	High. ~2,000 field RMs. Push products. Mid-tier lift of 20% maps to ~₹45 Cr incremental disbursement in 2 quarters.	High. Activity data lives in CRM + HRMS + sales reporting. Pattern shipped before.	6–8 weeks	Clear. RBAC per RM. Bias scaffold proven. Out of FPC scope, no borrower contact.	PILOT
DSA partner management copilot	High. DSAs source ~65% of two-wheeler + PL. Slippage today is invisible until month-close.	Medium. DSA data sits in three portals plus WhatsApp. Needs unification.	12–16 weeks	Medium. Partner contracts need DPDP refresh.	PHASE 2
Branch performance briefings	Medium. 110 branch managers, 5 dashboards each, 90 min daily lost to manual stitching.	Medium. Reporting layer unified, narrative layer missing.	10–14 weeks	Clear. Internal only.	PHASE 2
Credit memo drafting copilot	Medium. 18 hrs per memo today. ~1,000 memos/month.	Medium. Source data structured. Narrative template stable.	12–16 weeks	Medium. AI drafts, officer signs. HITL on every memo.	PHASE 3
Sales enablement copilot	Medium. Pitch personalisation by segment, product, bureau profile.	Low. Collateral library unsourced. No single content owner.	20+ weeks	Medium.	NOT NOW
Policy & process Q&A for branch staff	Low. Useful, but a productivity tool, not a P&L lever.	Low. Internal SOPs scattered, no single source of truth.	24+ weeks	Clear.	NOT NOW

The obvious pick was the wrong pick.

The DSA partner copilot looked like the obvious starting point. DSAs source two-thirds of the book. Leadership cared about it the most. The board had greenlit a ₹15 Cr envelope.

It was not selected.

DSA performance data sits across three portals and a heap of WhatsApp threads. Building the copilot on top of that would have shipped confident outputs from incomplete signals. Worse than no AI at all.

RM coaching was a closed workflow. Clean upstream data in CRM + HRMS + sales reporting. Push products mean every 1% productivity lift goes straight to disbursement and interest income. A governance pattern Wednesday has already shipped at a 4,500-agent life insurer with 25% mid-tier lift and zero audit findings.

This is the point of Discovery. Without it, the org spends the ₹15 Cr on the wrong use case and loses another year to a fourth stalled POC.

Pilot · RM productivity & coaching AI.

Today

2,000

Field RMs across branches + DSAs

3x

Volume gap between top 20% and bottom 60% of RMs

6 mo

Ramp time for a new RM to mid-tier performance

28%

Annual attrition on field RMs. Costly to rehire and retrain.

60%

Of manager time spent in reporting, not coaching

₹45 Cr

Incremental disbursement in reach. 20% lift on mid-tier RMs in 2 quarters.

Before, after, and the **confidence** around it.

MID-TIER RM PRODUCTIVITY	Baseline	→ +20% in 2 quarters
MANAGER COACHING COVERAGE	1:8 (1:3 in practice)	→ 1:all every RM, personalised
MANAGER TIME FREED	Baseline	→ ~60% back from reporting
NEW-RM RAMP TIME	6 months	→ ~3 months to mid-tier
ATTRITION (AI-USING COHORT)	28%	→ -15 ppt year-on-year

CONFIDENCE

High. Activity + outcome data sits in CRM, HRMS, sales reporting. Pattern shipped at a 4,500-agent life insurer with 25% mid-tier lift.

RISK

Low. RBAC per RM. Bias tested on region, language, tenure cohort before every model update. HITL on compensation- and termination-adjacent signals.

STARTING POINT

Internal-only. No borrower contact. No FPC surface. Not a customer-decision agent.

Top RM patterns, surfaced to every RM.

A multi-agent system surfaces what the top 20% do differently. Personalised next-best-action feed for each RM. Coaching brief for each manager. Out of FPC scope. No borrower contact. AI never touches a customer.

IN SCOPE

- Activity + outcome ingestion from CRM, HRMS, sales reporting
- Pattern-mining agent: what the top 20% do differently by product, segment, geography
- Next-best-action feed per RM, refreshed daily
- Coaching brief per manager, refreshed weekly
- RBAC: each RM sees only their own data and abstracted peer benchmarks
- Bias scaffold tested across region, language, tenure cohort, gender before every model update
- HITL on every compensation- or termination-adjacent signal
- Tested kill-switch back to the existing reporting flow

OUT OF SCOPE

- Any borrower-facing surface
- Automated compensation or performance actions
- Hiring or termination recommendations
- DSA partner data (Phase 2)
- Auto-generated sales pitches or collateral

What good looks like. And **when we measure it.**

MID-TIER RM PRODUCTIVITY

+20% in 2 quarters on the AI-using cohort.

MANAGER COVERAGE

1:all. Every RM gets a personalised coaching feed.

MANAGER TIME FREED

~60% back from reporting. Spent on coaching.

AUDIT FINDINGS

Zero on the pilot scope.

Timeline. 6–8 weeks. Two-region pilot. Productivity lift reviewed jointly with sales leadership and risk at week 6. Below 12%, scope narrows to one region and re-baselined. No silent extensions.

05 · ROADMAP

One pilot. Then the **playbook** spreads.

NOW	RM productivity & coaching AI Starts after Discovery sign-off. Two-region pilot. Owned jointly by Wednesday, sales leadership, HR, the CRO, and the CISO. Week 6 productivity checkpoint.
NEXT	DSA partner management copilot Partner data unification first. Six weeks of internal effort to consolidate three portals and the WhatsApp signal. Then the copilot. Same pattern as RM coaching, harder data.
IN PARALLEL	Branch performance briefings Daily AI-generated briefing per branch manager. Stitches the five dashboards into one narrative. Rides on the same data spine the RM copilot is building. Saves ~90 min/day across 110 branch managers.
NOT NOW	Credit memo drafting. Sales enablement. Policy Q&A. Credit memo drafting waits until the sanction-memo template is standardised. Sales enablement needs a single collateral owner. Policy Q&A is a productivity tool, not a P&L lever. Foundation problems. Not AI projects.

Not everything should be built. **Half the value of Discovery comes from what gets deprioritised.** In BFSI the orgs that ship are the ones that earn regulator trust on one use case at a time.

The journey to maturity.

A company moves from no AI to AI as infrastructure across five stages. Each one builds on the spine of the last. We walk it with you, stage to stage.

00 Shadow AI

WHAT WE FOUND

RMs drafting customer notes in ChatGPT on personal phones. Branch ops asking Claude for policy clarifications. Three POCs spent ₹2.4 Cr with nothing in production. Nothing logged, nothing reviewed.

01 Sanctioned Access

WHAT WE'D BUILD FIRST

One AI gateway in front of every model call. PAN, mobile, account data stripped before anything leaves your network. CISO gets the dashboard. Shadow AI pulled inside the perimeter.

02 First Production Use Case

THE PILOT WE RECOMMEND

RM coaching, two-region scope. Pattern-mining agent surfaces what the top 20% do differently. HITL on compensation-adjacent signals. Risk and compliance sign off before launch.

03 Multiple Agents, Shared Platform

WHAT WE SEQUENCED FOR PHASE 2

DSA partner copilot and branch performance briefings ride the platform the pilot stood up. Shared knowledge base, shared tool layer over CRM, HRMS, LOS. The DSA copilot doesn't pay the cost RM coaching did.

04 Autonomous Workflows

WHAT WE FLAGGED BUT DIDN'T SCOPE

Disbursement lane and sanction memo workflow. Both run 18–22 days of manual recon today. Once the platform exists, multi-agent handoffs turn them into one-day exception queues.

05 Org-Wide Self-Serve

WHAT THE BOARD ASKED ABOUT

When a branch ops team or a credit policy analyst can ship their own agent on the platform without a vendor. Show-back to finance. The board owns AI risk the way it owns credit risk.

Five planes. One spine.

Stages are when you build. Phases are what you build. We build and re-build these five planes with you as the agent count grows.

A Data Plane

WHAT WE AUDITED

Your LOS, LMS, HRMS, sales reporting, two bureau feeds, three DSA portals. Today none reconcile. The unification work that lets the RM copilot see a complete picture, and every agent after it too. AI quality is capped here.

B AI Plane

WHAT WE DESIGNED FOR THE PILOT

Pattern-mining agent, next-best-action agent, coaching-brief agent. Shared knowledge base of FPC, product catalogue, sanction policy. Shared tool layer over CRM, HRMS, LOS so the next agent gets free integration.

C Control Plane

WHAT WE'D BUILD FIRST

The AI gateway. Identity per RM. Audit log every model call. Evaluation harness that gates deployment. Kill-switch tested before launch. Built once for the pilot, reused by every agent that comes after.

D Consumption

WHAT WE SKETCHED WITH SALES LEADERSHIP

The RM mobile feed, refreshed daily. The branch manager's weekly coaching brief. The supervisor's real-time dashboard. Where 2,000 RMs meet the agents.

E Org & Regulatory

WHAT WE MAPPED WITH RISK AND COMPLIANCE


RBI FREE-AI alignment. FPC scope review. DPDP consent posture. Governance runbook. Bias scaffold by region, language, tenure cohort. Inspection pack ready before launch.

08 · OUTCOME

The pilot today. The journey together.

Three stalled vendor proposals replaced by one clear decision. A regulator-mapped governance plan risk and compliance signed off before the pilot started. The RM coaching pilot validated in six weeks.

Phase 2 planning started before the pilot ended. Risk and compliance now have a pattern they trust. The pattern compounds across every future workstream.



EXPLORE → **PROVE** → **SCALE**

This is Explore. The next move is Prove.