

IMPLEMENTATION HANDOFF

90-Day AI Pilot Roadmap

A practical route from Discovery findings to readiness, controlled pilot execution and an evidence-based investment decision.

SAMPLE / PREVIEW DOCUMENT Illustrative sample showing the format a client receives after an AI Discovery Engagement. Names, dates, figures and outcomes below are fictional and for demonstration only.

PREPARED FOR	Apex Precision Components Pvt. Ltd. (fictional sample)
ENGAGEMENT	AI Discovery Engagement — full-day onsite Discovery + post-visit decision package
RECOMMENDED FIRST WORKFLOW	Customer Enquiry Assistant
PREPARED BY	Byond Boundrys Consulting
ROADMAP HORIZON	90 days from readiness through controlled pilot evaluation
DOCUMENT STATUS	Illustrative sample — replace with client-specific findings after Discovery

WHAT THIS ROADMAP IS FOR This roadmap converts the Discovery recommendation into a staged 90-day path. The first 30 days cover readiness and preparation. The remaining period covers controlled pilot execution and evaluation.

INCLUDED VS OPTIONAL The 90-Day Pilot Roadmap itself is included in the ₹55,000 + GST AI Discovery Engagement. Execution or project management of this roadmap, readiness support, technical scoping, development and implementation are separate engagements.

Example first use case

Customer Enquiry Assistant classifies routine enquiries arriving by email and WhatsApp, retrieves approved information, prepares a draft response, routes sensitive or uncertain cases to a person, and logs the outcome for measurement.

01 / ROADMAP OVERVIEW

How the 90 days are organised

The roadmap is deliberately staged so the client does not move from an interesting idea directly into an uncontrolled build.

Phase	Days	Purpose	Decision at the end
Phase 1 — Readiness	1–30	Confirm the problem, baselines, source data, controls, integrations, pilot scope and acceptance criteria.	Proceed / revise / pause the pilot plan
Phase 2 — Controlled Pilot	31–60	Build the smallest usable pilot, evaluate offline, then expose it to a limited user group under human control.	Continue / redesign / pause live pilot
Phase 3 — Evidence & Decision	61–90	Stabilise, compare results with baseline, test limited expansion only where evidence supports it, and prepare the next investment decision.	Expand / revise / pause

What Day 90 should produce

- A measured comparison between the current workflow and the controlled pilot.
- Evidence on quality, speed, staff effort, adoption, escalation and control effectiveness.
- A record of recurring failures, source-data gaps and operating constraints.
- A clear recommendation to expand, revise or pause the pilot.
- An updated scope, dependencies and indicative commercial range for the next phase, if leadership chooses to proceed.

Core roles

Role	What this person is responsible for	Typical person
Executive sponsor	Removes blockers, confirms priority and owns major stage-gate decisions.	Business leader / director
Workflow owner	Explains the current process, approves changes and owns adoption.	Sales / service / ops lead
Content or data owner	Confirms which information is current, approved and safe to use.	Product / ops / knowledge owner
IT / system owner	Confirms access, integrations, security and environment constraints.	IT / application owner
Pilot users	Test the assisted workflow with real work and record corrections.	3–8 representative users
Delivery lead	Coordinates scope, architecture, evaluation and build planning.	Byond Boundrys / client lead

COMMERCIAL BOUNDARY This roadmap describes what should happen. It does not mean Byond Boundrys is contracted to run all 90 days. Any supported readiness, development, implementation or project management is separately proposed and priced.

02 / DAYS 1-5

Week 1 — confirm the problem and measure today

Do not design the AI solution yet. First make the current workflow measurable and agree exactly what the pilot will and will not cover.

Day	What happens	Owner	Output / proof
Day 1	Kick-off. Confirm the business problem, first user group and sponsor. Write a one-sentence pilot objective.	Sponsor + workflow owner	Signed-off pilot objective and named owners
Day 2	Map the current enquiry workflow from arrival to resolution. Mark hand-offs, repeated searches, approvals and re-entry.	Workflow owner + users	Current-state workflow map
Day 3	Capture baseline numbers: enquiry volume, response time, staff handling time, common categories and escalation rate.	Workflow owner	Baseline worksheet with source/date for each metric
Day 4	Define scope boundaries: which channel, enquiry types and customer group are in scope, and what must remain human.	Sponsor + workflow owner	In-scope / out-of-scope list
Day 5	Review whether the problem is large enough and measurable enough to justify pilot preparation.	Sponsor	Gate 1: proceed / revise / stop

WEEK 1 GATE Proceed only if one workflow owner is accountable, the current process can be measured and the pilot scope is narrow enough to test safely.

BY FRIDAY, EVERYONE SHOULD BE ABLE TO ANSWER What exact customer problem are we improving? Who owns the workflow? How often does it happen? Where is staff time currently spent? Which cases must remain human from the start?

03 / DAYS 6-10

Week 2A — prepare trusted knowledge and examples

The AI cannot be more reliable than the information and examples used to test it.

Day	Task	How to do it	Output
Day 6	List source knowledge	Name the exact folders, documents, product sheets, policies and FAQs currently used to answer enquiries.	Source inventory with owner for each item
Day 7	Remove weak sources	Mark duplicates, outdated versions, drafts and information that should not be exposed to the pilot.	Approved-source list
Day 8	Assign content owners	For each source, name who can confirm accuracy and how often it should be reviewed.	Ownership and freshness register
Day 9	Collect real enquiry examples	Sample common, difficult and risky enquiries. Remove personal or sensitive data where necessary.	Representative test set
Day 10	Write expected good outcomes	For each test enquiry, record the correct answer, source and whether human escalation is required.	Evaluation set v1

SIMPLE RULE If the business cannot agree what a good answer looks like for a test case, the AI system cannot be fairly evaluated on that case.

04 / DAYS 11-15

Week 2B — define controls, permissions and integrations

This converts “we have the data” into “we know how the pilot is allowed to use it.”

Day	Task	Key questions	Output
Day 11	Define approval rules	Which responses can be sent after normal review? Which need senior approval? Which must never be generated automatically?	Approval matrix
Day 12	Define exception rules	What happens when a source is missing, confidence is low, a customer is unhappy, or pricing/contract issues appear?	Exception and escalation flow
Day 13	Confirm user permissions	Which users can see which knowledge? Are any customer or product records restricted?	Access matrix
Day 14	Confirm integration route	Can the pilot read email, WhatsApp or CRM data? Can it write notes or drafts back? What API/export constraints exist?	Integration notes + system-owner sign-off
Day 15	Week 2 review	Are source data, test cases, controls and access sufficiently clear to design the pilot?	Gate 2: ready for pilot design / gaps to close

Minimum control rules for the first pilot

- Use approved sources only.
- Make the source used for an answer visible to the reviewer.
- Escalate low-confidence or unsupported answers rather than guessing.
- Keep pricing, contractual commitments, complaints and other high-impact cases human-controlled.
- Log every correction so the team can learn why the draft failed.

05 / DAYS 16-20

Week 3 — design the controlled pilot

Now the team has enough evidence to design the smallest useful pilot.

Day	Task	Decision to make	Output
Day 16	Design the future workflow	Agree the path: classify → retrieve approved knowledge → draft → human review → log.	Future-state workflow
Day 17	Choose the pilot interface	Decide whether users work inside email, WhatsApp Business, a simple web tool or another controlled interface.	User-flow decision
Day 18	Choose the technical pattern	Confirm model access, knowledge retrieval approach and how CRM/ERP updates will work.	Architecture v1
Day 19	Set acceptance criteria	Agree minimum quality, escalation behaviour and error tolerance for pilot release.	Written acceptance criteria
Day 20	Estimate build & dependencies	List build tasks, integration dependencies, environment needs and expected pilot duration.	Pilot estimate + dependency register

KEEP IT SMALL A good pilot proves one workflow with a limited user group. It should not require a company-wide migration, full knowledge-base cleanup or autonomous customer communication.

Week 4A — test the plan before the build

Use the evaluation examples and walkthroughs to find design problems before development effort increases.

Day	What to test	Who participates	Evidence produced
Day 21	Walk through 10-20 real examples manually using the proposed future workflow.	Workflow owner + pilot users	Missing rules, missing sources and unclear escalations
Day 22	Review difficult cases: ambiguous requests, missing information, complaints, pricing and unusual products.	Senior reviewer + users	Exception catalogue
Day 23	Review data, privacy and security design.	IT / security / management	Control sign-off or open actions
Day 24	Review integration feasibility with system owners.	IT / system owners	Confirmed API/export path and constraints
Day 25	Update the pilot brief based on evidence.	Delivery lead + workflow owner	Pilot brief v2 ready for decision

What should cause a redesign

- Too many test cases have no trusted source answer.
- The pilot depends on integrations that are not available or approved.
- The workflow requires autonomy in decisions leadership is not comfortable delegating.
- There is no baseline or no way to measure whether the pilot improved anything.

07 / DAYS 26–30

Week 4B — make the readiness decision

The goal of Day 30 is a clear evidence-based decision on whether the recommended pilot is ready to build.

Day	Action	What leadership should receive
Day 26	Finalise scope and architecture	One-page scope, future workflow and architecture diagram
Day 27	Finalise measurement plan	Baseline, pilot measures, thresholds and reporting cadence
Day 28	Finalise delivery plan	Build activities, dependencies, owners, timeline and indicative commercial scope
Day 29	Pre-decision review	Open risks, assumptions and unresolved questions explicitly listed
Day 30	Decision meeting	Proceed / revise / pause decision with reasons and named next actions

DAY 30 OUTCOME Leadership should be able to say: “We know the exact workflow, owner, data, controls, success measures, technical approach and cost/timeline assumptions for the pilot — and we know why we are proceeding or not proceeding.”

Readiness pathways

Readiness outcome	Next step
Ready to proceed	Prepare / approve a controlled AI pilot proposal.
Important gaps remain	Close the gaps internally or use a separately scoped Pilot Readiness Sprint.
Client will prepare internally	Use this roadmap and the included progress-review call to review readiness progress.

08 / DAYS 31–45

Phase 2A — build the smallest usable pilot

If Day 30 says “proceed,” the next step is a controlled build and offline evaluation — not immediate enterprise rollout.

Period	Focus	What should happen	Evidence produced
Days 31–35	Build the agreed workflow	Implement the agreed classification, approved-knowledge retrieval, drafting, review and logging flow.	Working pilot path in a controlled environment
Days 36–40	Connect approved sources and required systems	Use the approved knowledge/data set and only the confirmed integration routes.	Connected pilot environment with access controls
Days 41–43	Run offline evaluation	Test against the approved representative examples before exposing the pilot to live work.	Evaluation results, failures and correction log
Days 44–45	Fix blocking issues	Resolve missing rules, source gaps and critical control failures that would make live use unsafe or misleading.	Gate 3: ready / redesign / pause live pilot

OFFLINE-FIRST RULE Do not expose the pilot to live work until the evaluation set shows that the workflow, sources, human review and exception handling behave as agreed.

For the illustrative Customer Enquiry Assistant

- Start with email and WhatsApp — the two highest-volume channels in the sample.
- Retrieve from approved knowledge only and show the reviewer the source used.
- Keep pricing, complaints, commitments and unsupported answers human-controlled.
- Log category, response, source, correction and resolution status for measurement.

09 / DAYS 46–60

Phase 2B — controlled live pilot

Expose the pilot to a limited user group and measure it under normal operating conditions without removing human control.

Period	Focus	Operating rule	Evidence to collect
Days 46–50	Pilot-user onboarding	Confirm access, user responsibilities, exception rules and how corrections are recorded.	Named pilot users + operating checklist
Days 51–55	Limited live use	Use only the agreed enquiry types and channels. Keep required human review in place.	Response time, handling time, approvals, corrections, escalations
Days 56–58	Failure review	Group recurring failures by source issue, workflow rule, integration issue or model behaviour.	Failure taxonomy + priority fixes
Days 59–60	Acceptance checkpoint	Compare early results with the agreed acceptance criteria.	Gate 4: continue / revise / pause

WHAT NOT TO DO Do not expand channels, content, users or autonomy simply because the pilot is functioning. Expansion should follow evidence, not enthusiasm.

Illustrative success measures carried from the Findings Report

Measure	Illustrative baseline	Pilot target
First response time	Average 6.4 hours; median 3.1 hours	Under 1 hour for routine enquiries
Staff handling time	Approx. 9 minutes per enquiry	30–40% reduction in staff minutes
Draft acceptance	No AI-assisted drafting today	≥70% accepted with minor/no edits by week 6
Draft correction	No baseline	≤15–20% need substantial correction by week 6
User adoption	No assisted workflow today	≥80% of in-scope enquiries by week 8

10 / DAYS 61-75

Phase 3A — stabilise and measure

The pilot now needs to become reliable enough to judge. The priority is fixing repeat failures and comparing results with the baseline.

Period	Focus	What should happen	Output
Days 61-65	Stabilise recurring failures	Address repeated source gaps, unclear rules, user friction and integration defects identified during the live pilot.	Prioritised fix log with owners
Days 66-70	Retest difficult cases	Re-run the most common and highest-risk cases after fixes; verify that human escalation still works.	Regression / exception test results
Days 71-73	Measure business outcomes	Compare response speed, staff effort, quality, adoption and escalation with the baseline.	Measurement checkpoint
Days 74-75	Sponsor review	Confirm whether the evidence is strong enough to continue, whether more stabilisation is required, or whether the pilot should pause.	Gate 5: continue / revise / pause

AVOID VANITY METRICS Do not judge the pilot by the number of prompts or AI responses. Judge it by workflow quality, time, adoption, controlled error rates and business outcomes.

ROI hypothesis for this illustrative case

- The sample report estimates 55-70 staff-hours per week are currently lost to manual search and drafting across related work.
- Recovering even 30% of that time would return roughly 18-21 staff-hours per week to higher-value work.
- Faster, more consistent responses may also reduce service risk for the highest-value customer accounts, although that benefit should not be treated as guaranteed financial ROI.

11 / DAYS 76-90

Phase 3B — decide what to do next

Use evidence from the controlled pilot to decide the next investment rather than committing to expansion in advance.

Period	Action	Decision logic / output
Days 76-80	Test limited expansion only if evidence supports it	Add a small amount of content, users or workflow coverage only where the existing pilot is stable and measurable.
Days 81-85	Refresh the business case	Update success measures, remaining risks, dependencies, effort and indicative budget assumptions using actual pilot evidence.
Days 86-89	Prepare the Day 90 decision pack	Summarise results, failures, controls, adoption, economics and recommended next scope.
Day 90	Leadership decision	Expand / revise / pause — with reasons, owners and the next decision date.

If the decision is “expand”

- Expand gradually based on measured performance.
- Preserve the controls that proved necessary during the pilot.
- Define the next scope, owner, success measures and commercial proposal before increasing investment.

If the decision is “revise”

- Record which assumption, source, integration or operating rule failed.
- Assign actions to close the gaps and set a new decision date.

If the decision is “pause”

- Document why the opportunity is not ready.
- Keep the Discovery and pilot evidence so the opportunity can be revisited later.
- Redirect effort to a higher-readiness opportunity rather than forcing the original idea.

12 / PILOT BUSINESS CASE

Success measures, ROI hypothesis and indicative investment

These are sample planning assumptions for the fictional Apex case. They are not a quotation and must be refreshed with client-specific evidence.

Recommended first pilot

Customer Enquiry Assistant for email + WhatsApp, one team, human-reviewed.

Business-case element	Illustrative sample
Current problem	Approx. 640 enquiries/month; average first response 6.4 hours; knowledge spread across multiple sources.
Pilot objective	Use approved knowledge to classify enquiries, draft responses, route exceptions and log outcomes.
Success measures	Response time, handling effort, draft acceptance/correction, escalation behaviour and user adoption.
ROI hypothesis	Recover part of the staff time currently lost to search/drafting and improve response consistency.
Indicative Phase 1 range	₹8-14 lakh for the controlled Customer Enquiry Assistant pilot in this illustrative case.
Typical illustrative duration	6-8 weeks build + pilot, subject to readiness, integrations and scope.

COMMERCIAL NOTE The indicative range above belongs to this fictional sample and is not a standard offer price. A binding proposal follows confirmed scope, integrations, data volumes and readiness. Readiness, technical scoping, development and implementation are separately proposed and priced.

13 / RESPONSIBILITY MATRIX

Who owns what across the 90 days

A task is not complete until one person is accountable for the outcome.

Workstream	Business sponsor	Workflow owner	IT / system owner	Delivery team
Pilot scope	Approve	Own	Consult	Support
Baseline metrics	Review	Own	Support	Support
Knowledge / data	Review	Own / assign owner	Support access	Prepare / test
Controls & exceptions	Approve	Own	Review security	Design / support
Integration	Informed	Consult	Own	Design / build
Offline evaluation	Review	Own business quality	Support	Design / run
Controlled live pilot	Review	Own adoption	Support operations	Support / monitor
Day 90 decision	Own	Recommend	Advise	Present evidence

Suggested weekly decision meeting agenda

#	Question	Why it matters
01	What changed since last week?	Completed tasks, new evidence and resolved questions.
02	What is blocked?	One owner and due date for each blocker.
03	What did we learn that changes the scope?	Do not hide assumptions that proved wrong.
04	Are we still solving the same business problem?	Avoid scope drift into unrelated AI ideas.
05	What decision is needed this week?	Make decisions explicit rather than carrying ambiguity forward.

14 / 90-DAY DECISION PACK

What should physically exist at the end of the roadmap

Use this as the evidence checklist before approving wider deployment or the next investment.

Required item	What “complete” looks like
Pilot objective	One sentence explaining the business problem and expected outcome.
Current-state workflow	How work happens today, including hand-offs and delays.
Future-state workflow	How AI assists the process and where people remain in control.
Baseline metrics	Current volume, speed, handling effort, errors/escalations and other agreed measures.
Approved knowledge/data set	Named sources, owners and exclusions.
Evaluation set	Representative real examples with expected good outcomes.
Approval & exception rules	Written rules for human review, low confidence and sensitive cases.
Integration plan	Systems, access route, owner and known constraints.
Architecture v1	Major components and data flow in plain language.
Acceptance criteria	What must be true before users can rely on the pilot.
Pilot evidence	Measured results, failures, corrections, adoption and controls from the live pilot.
Decision record	Expand / revise / pause — and why.

OWNER REMINDER If an item has no named owner or no evidence behind it, it is not complete.

15 / COMMERCIAL BOUNDARY

What the ₹55,000 AI Discovery Engagement includes — and what it does not

This page prevents the roadmap from being mistaken for 90 days of included delivery support.

Included in the AI Discovery Engagement	Separately proposed and priced
AI Discovery Findings Report	Pilot Readiness Sprint, if required
AI Opportunity Scorecard	Technical scoping beyond the Discovery deliverables
90-Day Pilot Roadmap	AI solution development or implementation
Leadership readout	Execution or project management of the 90-day roadmap
One progress-review call of up to 45 minutes within 30 days of the leadership readout	Data cleaning, migration, labelling or system integration
Executive direction within 48 hours and complete decision package within 5 business days, subject to the agreed terms	Ongoing consulting, support or advisory work beyond the included follow-up

How this roadmap can be used after Discovery

- If the client is ready, use the roadmap to support a separately scoped controlled pilot proposal.
- If important gaps remain, use the roadmap to define a separately scoped Pilot Readiness Sprint.
- If the client prefers to prepare internally, use the roadmap as the working guide and review progress during the included follow-up call.

FINAL PRINCIPLE The purpose of the roadmap is not to force implementation. It is to make the next investment decision explicit, measurable and evidence-based.