BYOND BOUNDRYS

BYOND BOUNDRYS CONSULTING

FINAL CLIENT DELIVERABLE

AI Discovery Findings Report

What we found, what it means, and the practical AI changes we recommend in language a leadership team can act on.

SAMPLE / PREVIEW DOCUMENT Illustrative sample showing the format, depth and structure a client receives after an AI Discovery Engagement. All company names, figures and quotes below are fictional and for demonstration only.

PREPARED FOR	Apex Precision Components Pvt. Ltd. (fictional sample)
INDUSTRY	Precision-machined components for automotive & industrial OEMs
COMPANY SIZE	₹184 crore annual revenue · approx. 260 employees · 3 sites (Noida, Manesar, Pune)
ENGAGEMENT	AI Discovery Engagement — one full onsite Discovery day + post-visit analysis, complete decision package and leadership readout
PREPARED BY	Byond Boundrys Consulting — AI Discovery Team
REPORT DATE	Illustrative sample — dated on issue to client
DOCUMENT STATUS	Illustrative sample — replace with client findings after discovery

Purpose of this report

This report is written for business leaders, not AI specialists. It explains the current problems in plain language, shows where AI can realistically help, and describes a practical, low-risk path from today's workflow to a safer and more efficient future workflow. It intentionally avoids technology hype: every recommendation is tied to a specific business problem, a named owner and a way to measure whether it worked.

Where the visit did not produce enough evidence to make a confident recommendation, we say so directly rather than filling the gap with a generic AI idea.

CONTENTS

What is in this report

Sixteen sections plus two appendices, organised so a leader can read only Section 01 and still understand the recommendation, or read end-to-end for full supporting evidence.

01 / EXECUTIVE SUMMARY

What we learned, in one page

This is the page a leadership team should be able to read on its own and still make an informed decision.

MAIN CONCLUSION The business does not need "AI everywhere." The strongest first move is to improve the customer-enquiry workflow, because it is frequent (roughly 640 enquiries a month across email, WhatsApp and phone), measurable, and already supported by existing business knowledge that simply needs to be organised.

The five things leadership should know

- Customer questions arrive through three separate channels, so response quality and speed depend heavily on which employee happens to receive the enquiry average first response is 6.4 hours, ranging from 20 minutes to 3 working days.
- Useful product, pricing and policy knowledge exists, but it is spread across roughly 40 shared-drive folders, personal WhatsApp chats and the experience of 4–5 senior staff, rather than being easy to search.
- Teams spend an estimated 55–70 staff-hours a week re-entering information from purchase orders and enquiries that follow repeatable patterns.
- The best first AI project is one with a clear owner, clear source information and a human approval step not a fully autonomous agent making unsupervised decisions.
- A controlled 6–8 week pilot should prove quality, time saved and staff adoption before any wider rollout is funded.

Recommended order

01	**Start with customer-enquiry assistance** Use approved business knowledge to draft accurate answers, classify enquiries and route exceptions to people. Estimated impact: response time cut from ~6.4 hours to under 1 hour for routine questions.
02	**Then add an internal knowledge assistant** Give staff one place to search approved product, process and policy information with source references, reducing dependence on 4–5 senior staff for routine questions.
03	**Then automate selected document work** Extract structured data from purchase orders and other repeatable documents, validating before system updates targeting ~110 POs processed manually each week.

WHAT WE ARE NOT RECOMMENDING YET A company-wide autonomous agent, a large-scale document migration, or automated pricing/contract decisions. Each of these was raised during the visit, scored on the companion Opportunity Scorecard, and parked until the first pilot has produced evidence.

02 / DISCOVERY METHODOLOGY

How we reached these conclusions

A recommendation is only as trustworthy as the process behind it. This section explains exactly how the AI Discovery Engagement was run so leadership can judge the evidence, not just the conclusion.

Engagement structure

Stage	Format	What happened
1. AI Fit Call	20-min online call	Mandatory qualification to confirm a suitable workflow, sponsor, access and intention to act.
2. Preparation Call	Remote preparation call	Confirm up to 3 business functions, 6–8 stakeholders, candidate workflows, systems and access required.
3. Onsite Discovery	1 full working day, onsite	Interviews and live workflow, document and system review across the agreed scope.
4. Executive Direction	Within 48 hours of onsite day	Brief headline direction and likely first workflow, issued after the balance is received.
5. Complete Decision Package	Within 5 business days	Findings Report, Opportunity Scorecard and 90-Day Pilot Roadmap, subject to timely payment and access.
6. Leadership Readout & Follow-up	Remote readout + one review call	Readout after package delivery; one 45-minute progress-review call is available within 30 days.

TIMELINE, IN PLAIN TERMS Executive direction is provided within 48 hours of the onsite day after balance payment. The complete written decision package follows within five business days, subject to timely payment and access. The leadership readout follows delivery, with one 45-minute progress-review call available within 30 days.

What the on-site visit covered

- Structured interviews with 8 stakeholders across customer service, sales and operations; IT and Quality supported the systems, controls and documentation review (Appendix B).
- Live shadowing of two customer-service agents handling real enquiries, plus review of 14 sample documents (POs, spec sheets, the FAQ and two policy documents).
- A walkthrough of the CRM, shared email inbox and shared-drive folder structure with the IT lead.
- A short, anonymous 11-response team survey on where employees feel the most time is lost.

WHAT WE DID NOT DO We did not review payroll, HR or finance systems out of scope for this engagement and did not access customer PII beyond what was already visible in the CRM during normal work; no customer data left the client's systems.

03 / COMPANY & MARKET SNAPSHOT

The business context behind the recommendation

AI recommendations that ignore business context tend to fail. This snapshot captures the scale and shape of the business so the recommended pilot is sized correctly.

Dimension	Illustrative sample detail
Business	Precision-machined metal components for automotive and industrial OEMs
Scale	₹184 crore annual revenue · approx. 260 employees · 3 sites
Customer base	Approx. 95 active B2B accounts; top 10 accounts ≈ 60% of revenue
Enquiry channels	Email (55%), WhatsApp Business (30%), phone (15%)
Enquiry volume	≈ 640 customer enquiries per month; peaks 40% higher in Q1 and Q3
Core systems	Zoho CRM, Tally ERP, shared Google Drive, WhatsApp Business, email
Customer-facing team	6 customer-service staff, 4 regional sales managers
Seasonality	Order volume rises ahead of automotive OEM production cycles

Why this matters for the recommendation

- A 260-person, 3-site business needs a workflow that works consistently across sites — not a solution tuned to one office.
- With 85% of enquiries arriving through email and WhatsApp, the first pilot should focus there rather than the lower-volume phone channel.
- Seasonal peaks mean staff time savings matter most in Q1 and Q3 — the pilot should be running and stable before the next peak.
- A concentrated customer base (top 10 accounts ≈ 60% of revenue) means response consistency for a relatively small number of named accounts has outsized business impact.

04 / DISCOVERY CONTEXT

What we reviewed during the visit

This section explains the business context behind the recommendations. In a real client report, these details are replaced with the actual departments, systems, volumes and observations from the on-site visit.

Area reviewed	What we looked for	Illustrative sample observation
Customer enquiries	Channels, response time, repeated questions, escalation	Email, phone and messaging are handled separately. Many questions repeat, but answers are recreated manually each time.
Internal knowledge	Where staff find product/process answers	Information is split across shared folders, spreadsheets and the memory of experienced staff.
Documents	High-volume files that are read and re-keyed	Purchase-order data is checked and entered manually before downstream processing — about 110 POs a week.
Systems & data	CRM/ERP/email/file repositories and access	Core systems exist, but information is not connected in a way that supports a single AI workflow.
Controls	Approval, privacy, access and exception handling	Human review exists informally; it should be made explicit in any AI-assisted process.

What we did not assume

- We did not assume that every manual step should be automated.
- We did not recommend a specific model or vendor before confirming data, integration and security requirements.
- We did not treat estimated productivity gains as guaranteed financial ROI — baselines must be measured first.
- We did not assume the sample scores in this document reflect a real client — in a live engagement, every score in the companion Scorecard is traced back to a specific interview or observation.

05 / FINDINGS

What is slowing the business down

Each finding below is written in three parts: what we observed, why it matters, and what we recommend doing about it.

Finding 1

Answers depend on the person handling the enquiry

WHAT WE OBSERVED

Similar questions are answered differently across email, phone and messaging. Staff search old messages or ask experienced colleagues for help. In our sample of 20 recent enquiries, the same question produced three meaningfully different answers.

WHY IT MATTERS

Response time varies (20 minutes to 3 working days), customers can receive inconsistent information, and experienced staff become bottlenecks two employees account for roughly 45% of all knowledge-related interruptions reported in the team survey.

WHAT WE RECOMMEND

Create an AI-assisted response workflow that uses only approved business knowledge and requires human review for sensitive or uncertain answers.

Finding 2

Knowledge is difficult to find

WHAT WE OBSERVED

Product details, policies and process information are spread across roughly 40 shared-drive folders, spreadsheets and people. No single source is treated as “current.”

WHY IT MATTERS

Employees waste an estimated 6–8 hours a week searching or interrupting colleagues. New staff reported taking 4–6 weeks to become fully productive on enquiry handling.

WHAT WE RECOMMEND

Create one governed knowledge layer with clear content owners, freshness dates and source citations.

Finding 3

Document handling is repetitive

WHAT WE OBSERVED

People repeatedly read purchase orders and copy fields such as customer, item, quantity, date and reference number into other systems — roughly 110 purchase orders a week.

WHY IT MATTERS

Manual entry consumes an estimated 15–20 staff-hours a week and creates avoidable transcription mistakes; the operations lead estimated a 2–3% error rate on manually keyed fields.

WHAT WE RECOMMEND

Use document extraction to pre-fill structured fields, then validate confidence and exceptions before posting data downstream.

Finding 4

No shared success measures for AI

WHAT WE OBSERVED

The business has ideas for AI, but no common baseline for response time, handling time, quality, adoption or exception rate.

WHY IT MATTERS

Without a baseline, it is difficult to prove whether a pilot actually helped, which slows down the decision to expand, revise or pause.

WHAT WE RECOMMEND

Measure the current workflow before the build and define 3–5 pilot metrics that leadership agrees to use.

Finding 5

AI ideas outnumber the business's delivery capacity

WHAT WE OBSERVED

During the visit, staff and leadership raised 9 separate AI ideas — from customer-enquiry assistance to fully autonomous pricing decisions. Seven candidate workflows were selected for formal evaluation in the companion AI Opportunity Scorecard, while the remaining ideas were parked before scoring.

WHY IT MATTERS

Attempting several AI projects at once, with a small internal team, risks under-resourcing all of them and producing no completed pilot within the year.

WHAT WE RECOMMEND

Use the companion AI Opportunity Scorecard to select one starting point, and revisit the remaining ideas once the first pilot has produced evidence.

06 / CURRENT WORKFLOW

How work happens today

The onsite Discovery day should make the current process visible before proposing technology. This sample shows the customer-enquiry workflow as observed during the agreed scope.

01	**Enquiry arrives** A customer sends a question by email, WhatsApp or phone. No single inbox owns all three channels.
02	**Staff member interprets it** The person decides what the customer is asking and whether they can answer it, averaging 4–6 minutes per enquiry just to interpret intent.
03	**Information is searched manually** The employee checks old messages, shared-drive files, spreadsheets, or asks a colleague — the single largest source of delay observed.
04	**Response is written from scratch** The answer is typed, reviewed informally, and sent — with no record of which source information was used.
05	**Follow-up is recorded inconsistently** Some details reach the CRM or another tracker; others remain only in the conversation thread, making later reporting unreliable.

WHERE TIME IS LOST The largest delays are not in “writing text.” They are in understanding the request, finding trusted information, deciding what is safe to say, and recording the result consistently. Across the shadowed sessions, searching for information took roughly 3× longer than drafting the actual reply.

What should stay human

- Pricing exceptions, commitments, contractual language and sensitive customer decisions.
- Any answer where the source information is missing, outdated or contradictory.
- Escalations, complaints and unusual requests that require judgement.

07 / RECOMMENDED WORKFLOW 1

Customer Enquiry Assistant — recommended first pilot

This is the highest-priority recommendation in the illustrative sample because it has frequent volume, clear business knowledge and measurable outcomes.

What changes for the team

01	**Enquiry enters one assisted workflow** Email and WhatsApp — the two highest-volume channels — are captured for the pilot.
02	**AI classifies the request** The system identifies the enquiry type (pricing, spec, order status, complaint, other) and required business information.
03	**Approved knowledge is retrieved** The assistant searches only the governed knowledge base and records the sources used.
04	**A draft response is prepared** The system creates a draft and flags missing information or low confidence rather than guessing.
05	**A person reviews when required** A team member approves, edits or escalates based on defined rules — pricing and complaints always route to a person.
06	**Outcome is logged** Category, response, source and resolution status are captured in the CRM for measurement and later improvement.

Why this is valuable

- Faster handling of repeated questions without removing human control — target: under 1 hour first response for routine enquiries, down from 6.4 hours.
- More consistent answers because drafts come from approved sources, reducing the 3-different-answers pattern observed in the sample review.
- Less dependence on 4–5 experienced employees for routine knowledge.
- A measurable pilot: response time, handling time, approval rate, correction rate and customer outcome can all be tracked from day one.

08 / RECOMMENDED WORKFLOW 2

Internal Knowledge Assistant

This recommendation addresses the same knowledge problem internally. It is useful after content owners and source documents have been cleaned up for the first pilot.

How it works

01	**Employee asks a business question** For example, “Which specification applies to this product?” or “What is our current return policy?” — the kind of question that currently interrupts 2–3 senior staff several times a day.
02	**Assistant searches approved information** Only the governed knowledge base is searched — starting with a limited, high-value document set, not every shared drive on day one.
03	**Answer is returned with its source** The employee sees the answer alongside the document or policy it came from, so it can be checked at a glance.
04	**Weak or unsupported answers are flagged** If the evidence is missing or unclear, the assistant says so instead of guessing — and the question is logged for a content owner to resolve.

Good first content

- Product specifications and approved product notes.
- Frequently used process SOPs.
- Current customer-service policies.
- Approved FAQs and standard operating guidance.

Controls to include

- Access should respect the employee's existing permissions.
- Each answer should identify the source used.
- Documents need owners and review dates.
- Unanswered questions should become a content-improvement queue.

09 / RECOMMENDED WORKFLOW 3

Purchase Order & Document Automation

Use AI to reduce repetitive document reading — not to remove validation.

01	**Document arrives** A purchase order or selected document enters the controlled intake channel — roughly 110 documents a week in the sample.
02	**Required fields are extracted** The system reads agreed fields such as customer, PO number, item, quantity, date and delivery information.
03	**Business rules check the result** Required fields, known customers, allowed formats and duplicate references are validated automatically.
04	**Exceptions go to a person** Low-confidence or unusual documents are routed for correction — expected to be 10–15% of volume initially.
05	**Approved data moves downstream** Validated fields are sent to Tally ERP or prepared for import, removing the manual re-keying step.

IMPORTANT Do not begin by trying to automate every document type. Pick one high-volume format (purchase orders), measure accuracy against the current 2–3% manual error rate, and expand only after the exception process works.

10 / TECHNOLOGY DIRECTION

What the solution needs — in business language

Technology should support the workflow above. The exact product choices can change; the required capabilities are more important than brand names at this stage.

Capability	Why it is needed	Examples of implementation choices
Secure model access	Generate, classify and summarise while keeping business controls around usage.	Enterprise LLM API through an approved cloud/provider.
Knowledge retrieval	Find the right internal information before generating an answer.	RAG/search layer over approved documents with source citations.
Workflow/orchestration	Move work through classify → retrieve → draft → review → log.	Application backend or workflow service with explicit business rules.
Business integrations	Read from and write to the systems already used by staff.	Email, Zoho CRM, Tally ERP, WhatsApp Business API as applicable.
Evaluation & monitoring	Measure quality, failures, cost and usage over time.	Test set, human review logs, quality dashboards, error/latency monitoring.
Access & audit	Control who can see what and keep traceability.	SSO/role-based access, audit logs, document permissions, retention rules.

Suggested architecture

01	**Business channel** Email, WhatsApp Business API or another controlled entry point.
02	**Workflow service** Applies routing, business rules and approval logic.
03	**Knowledge layer** Retrieves relevant approved content and records sources.
04	**AI model** Performs classification, extraction or drafting for the specific step.
05	**Human review & systems** People approve exceptions; approved outcomes are written to Zoho CRM/Tally ERP where appropriate.

06	**Monitoring** Quality, usage, corrections and errors are captured for continuous improvement.

11 / DATA & INTEGRATION READINESS

What must be prepared before development

Most AI project delays come from unclear source data and system ownership, not from the model itself.

Readiness item	What “ready” means	Owner to confirm
Knowledge sources	The pilot uses a named set of current documents. Old/duplicate versions are excluded.	Business content owner
Customer-enquiry examples	Representative past enquiries and approved good answers are available for testing.	Customer-service / sales lead
System access	Required APIs, exports or integration routes to Zoho CRM and Tally are confirmed.	IT / system owner
Permissions	The pilot knows which users can access which information.	IT + business owner
Data handling	Retention, sensitive data, vendor use and logging rules are agreed.	Management / security / legal as applicable

GO / NO-GO RULE Do not start the pilot until the workflow owner, approved source data, success measures and human-review rules are named.

12 / CONTROLS & RISKS

How to keep the pilot safe and useful

The goal is not zero risk. The goal is to make the important risks visible and controllable.

Risk	What could happen	Control we recommend
Incorrect answer	The assistant gives a confident answer that is not supported by approved information.	Source-grounded retrieval, citations, confidence rules and human approval for defined categories.
Outdated knowledge	A correct old document produces the wrong current answer.	Named document owners, review dates and version control.
Sensitive information exposure	A user sees data they should not have access to.	Role-based access, source permissions, logging and minimum-data design.
Over-automation	The team trusts automation for cases that require judgement.	Explicit exception categories and mandatory human decision points.
Poor adoption	Staff bypass the system because it adds friction or is unreliable.	Design with end users, measure edits/rejections, improve before rollout.
Vendor / model dependency	Over-reliance on one provider affects continuity or cost predictability.	Prefer swappable components and monitor usage cost against agreed budget.

13 / SUCCESS MEASURES

How leadership will know whether the pilot worked

Measure the current process before the build. Then compare the same measures during the controlled pilot.

Measure	Baseline (illustrative)	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 of staff effort per enquiry	30–40% reduction in staff minutes per enquiry
Draft acceptance rate	Not applicable before pilot	≥ 70% of AI drafts accepted with minor or no edits by week 6
AI draft correction rate	N/A — no AI-assisted drafting today	≤15–20% of routine drafts require substantial correction by week 6
Escalation rate	Not formally tracked today	Routine questions require fewer interruptions; genuine exceptions still escalate
User adoption	N/A — no assisted workflow exists today	≥ 80% of in-scope enquiries handled through the assisted workflow by week 8

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

14 / INVESTMENT & COMMERCIAL OVERVIEW

What this is likely to cost, in phases

Figures below are illustrative planning ranges for a business of this size. They are refined using the scope, dependencies, readiness findings and indicative effort recorded in the companion 90-Day Pilot Roadmap. They are not a quotation.

WHAT'S ALREADY INCLUDED The ₹55,000 + GST founding price for the first 10 paid AI Discovery Engagements includes this AI Discovery Findings Report, the AI Opportunity Scorecard and the 90-Day Pilot Roadmap. Everything in the table below — pilot development, later expansion/automation and ongoing hosting — is a separate, optional next investment if leadership decides to proceed.

Phase	What it covers	Illustrative range	Typical duration
Phase 1 — Pilot	Customer Enquiry Assistant for email + WhatsApp, one team, human-reviewed	₹8–14 lakh	6–8 weeks build + pilot
Phase 2 — Expand	Internal Knowledge Assistant + wider enquiry coverage	₹6–10 lakh	4–6 weeks
Phase 3 — Automate	Purchase order & document automation with ERP integration	₹7–12 lakh	5–7 weeks
Ongoing	Hosting, model usage, monitoring and minor improvements	₹1.5–2.5 lakh / month	Ongoing

How to think about payback

- If the Phase 1 pilot recovers even 30% of the estimated 55–70 staff-hours a week currently lost to manual search and drafting, that is roughly 18–21 staff-hours a week returned to higher-value work.
- Faster, more consistent responses to the top 10 accounts (≈ 60% of revenue) reduce the risk of losing business to slow or inconsistent service — a harder number to quantify, but the one leadership raised most often during interviews.
- We recommend treating the first controlled pilot as a proof point, not a committed spend for later phases. The Day 30 readiness decision in the 90-Day Pilot Roadmap is the first investment gate; the Day 90 decision determines whether to expand, revise or pause.

COMMERCIAL NOTE These ranges assume the client's existing systems (Zoho CRM, Tally and WhatsApp Business) are used as-is, with no major replatforming. A formal, binding quotation follows confirmed scope, readiness, integrations and data requirements. Readiness support, technical scoping, development and implementation are separately proposed and priced.

15 / 90-DAY PILOT ROADMAP HANDOFF

How the recommendation moves from readiness to a controlled pilot decision

The detailed staged plan is provided separately in the 90-Day Pilot Roadmap. The first 30 days cover readiness and preparation; the remaining period covers controlled pilot execution, measurement and evaluation.

01	**Confirm one recommended pilot workflow** Leadership confirms the recommended first workflow, accountable owner and scope boundaries.
02	**Days 1–30: readiness and preparation** Capture baselines, approve source data, define controls, confirm integrations and acceptance criteria; make the Day 30 proceed / revise / pause decision.
03	**Days 31–45: build and offline evaluation** Build the smallest usable pilot and test it against the approved evaluation set before live exposure.
04	**Days 46–60: controlled live pilot** Run with a limited user group, required human review and monitoring of response time, handling effort, corrections and escalations.
05	**Days 61–75: stabilise and measure** Fix recurring failures, retest difficult cases and compare pilot outcomes with the baseline.
06	**Days 76–90: evidence-based decision** Refresh the business case and decide whether to expand, revise or pause, with reasons, owners and the next decision date.

What we recommend not doing yet

- Do not launch a company-wide autonomous agent before one controlled workflow proves value.
- Do not connect every internal document source before content ownership and permissions are clear.
- Do not replace human approval in pricing, contracts, complaints or other high-impact decisions during the first phase.

16 / FINAL RECOMMENDATION

The decision we recommend leadership make

A good AI Discovery Findings Report should end with a clear decision, not a list of technologies.

RECOMMENDED DECISION Proceed to the readiness phase for a controlled Customer Enquiry Assistant pilot, subject to confirmation of the business owner, approved knowledge sources, baseline metrics and integration access. Use the companion 90-Day Pilot Roadmap to make the Day 30 proceed / revise / pause decision before committing to the controlled build.

Why this first

- The problem is frequent (≈ 640 enquiries/month) and visible to the business.
- The workflow can be limited to a manageable set of enquiry types and two channels (email, WhatsApp).
- Quality can be checked against approved answers and sources.
- Human review can remain in place while the team learns where automation is reliable.
- The result can be measured within 6–8 weeks, creating evidence for the next investment decision.

Decision checklist for the client

Decision	Status / owner
Confirm the first pilot workflow	_____
Name the executive sponsor	_____
Name the day-to-day workflow owner	_____
Approve the source data/content set	_____
Approve success measures	_____
Confirm next pathway: internal readiness / Pilot Readiness Sprint / controlled pilot proposal	_____

APPENDIX A / GLOSSARY

Plain-English definitions

Technical terms used elsewhere in this report, explained in business language.

Term	What it means in plain English
AI agent	Software that can take multi-step actions on its own (e.g. classify, look up, draft) rather than just answering one question.
RAG (retrieval-augmented generation)	A method where the AI first searches approved documents, then writes an answer based only on what it found — so answers can be traced to a source.
LLM (large language model)	The underlying AI technology that reads and generates natural-language text, such as drafting a customer reply.
Confidence / confidence threshold	A measure of how sure the AI is about an answer; below the agreed threshold, the case is routed to a person instead of being sent automatically.
Human-in-the-loop	A design principle where a person reviews or approves AI output before it takes effect — used throughout the recommended pilot.
Pilot	A small, time-boxed, controlled test of a new workflow with a limited group of users before any wider rollout.
Baseline	A measurement of how the process performs today, taken before any changes, so improvement can be proven rather than assumed.
Escalation	A case routed to a human because it is sensitive, ambiguous, or outside what the AI is allowed to decide.
Integration	A connection that lets the AI workflow read from or write to an existing system, such as the CRM or ERP.
Exception rate	The share of cases that could not be handled automatically and were routed to a person.

APPENDIX B / WHO WE SPOKE WITH

Interviews and documents reviewed

Provided for transparency and so the client can confirm nothing important was missed.

People interviewed (illustrative sample)

Role	Area	Format
Managing Director	Leadership / strategy	45-min interview
Head of Customer Service	Customer enquiries	60-min interview + shadowing
Customer Service Executive (x2)	Day-to-day enquiry handling	Shadowing, 45 min each
Regional Sales Manager	Sales & key accounts	30-min interview
Operations Lead	Purchase order processing	45-min interview + document walkthrough
IT Manager	Systems, CRM, integrations	45-min systems walkthrough
Quality Manager	Documentation & SOPs	30-min interview

Documents & systems reviewed

- 14 sample documents: purchase orders, product spec sheets, current FAQ document, 2 policy documents.
- Zoho CRM structure and current usage patterns.
- Tally ERP data fields relevant to order processing.
- Shared Google Drive folder structure (≈ 40 folders reviewed at a high level).
- Anonymous 11-response team survey on where time is lost.

END OF SAMPLE REPORT In a real engagement, all sample observations, workflow details, priorities, figures and recommendations are replaced with evidence collected during the client's own AI Discovery Engagement.