BYOND BOUNDRYS

BYOND BOUNDRYS CONSULTING

SUPPORTING ANALYSIS

AI Opportunity Scorecard

A simple, evidence-based explanation of which AI ideas are worth doing first and why.

SAMPLE / PREVIEW DOCUMENT Illustrative sample showing the format a client receives after an AI Discovery engagement. Scores, quotes and figures below are fictional and for demonstration only.

PREPARED FOR	Apex Precision Components Pvt. Ltd. (fictional sample)
ENGAGEMENT	AI Discovery Engagement — full-day onsite Discovery
PREPARED BY	Byond Boundrys Consulting
OPPORTUNITIES SCORED	7 candidate workflows (5 primary + 2 parked ideas, within the approved 5–7 workflow scope)
DOCUMENT STATUS	Illustrative sample — replace with client findings after discovery

WHAT THIS DOCUMENT DOES This scorecard complements the AI Discovery Findings Report. It shows how five to seven candidate workflows are compared using one consistent method, so leadership can see why one workflow is recommended for the first pilot and can challenge any rating if the evidence does not support it.

How to use it

- Read the plain-English definitions of the five scoring criteria first.
- Then review the one-page summary ranking.
- For each opportunity, read the explanation behind the score not just the number and the evidence it is based on.
- A high score means “good candidate for the next phase,” not “guaranteed ROI.”
- Use the Decision Guide (Section 08) to turn the ranking into an actual choice.

01 / SCORING METHOD

What the numbers actually mean

Every opportunity is rated from 1 (weak) to 5 (strong) across five questions. The weighted total helps compare opportunities consistently rather than by gut feel or how exciting an idea sounds.

Question we ask	Plain-English meaning	Weight
Business value	If this works, will it materially improve time, quality, revenue, cost or customer experience?	30%
Feasibility	Can we realistically build and operate it with the current process, systems and constraints?	25%
Data readiness	Do we have enough clean, accessible and approved information to make it work?	20%
Implementation complexity	Can the first pilot be implemented with manageable integration, process-change and operating complexity?	15%
Risk	Can mistakes, privacy, approvals, exceptions and human-control requirements be managed safely?	10%

How to read a 1–5 rating

- 5 — Strong: clear evidence that this criterion is favourable for the opportunity. For implementation complexity, 5 means low or well-contained complexity.
- 4 — Good: mostly favourable; some preparation is required but no major blocker is visible.
- 3 — Mixed: potential exists, but important unknowns or dependencies need work.
- 2 — Weak: significant preparation, risk or uncertainty before this should be prioritised.
- 1 — Poor: not suitable for near-term investment based on current evidence.

Where the weights come from

Business value carries the highest weight (30%) because a technically easy project that does not matter to the business is not worth doing. Feasibility (25%) and data readiness (20%) come next because they indicate whether the project can actually be delivered with the client's current process, systems and information. Implementation complexity (15%) rewards opportunities whose first pilot can be kept bounded, with manageable integrations, operating change and dependencies. Risk (10%) reflects whether mistakes, privacy, approvals, exceptions and human-control requirements can be managed safely. The weighted total supports comparison, but a material control gap can still prevent an opportunity from proceeding.

02 / SUMMARY RANKING

Which opportunities rise to the top

These are illustrative sample scores. In a live engagement, the numbers are based on interviews, observed workflows, available data and client constraints not on how novel the idea sounds.

Rank	Opportunity	Score	Decision	Why
1	Customer Enquiry Assistant	4.5 / 5	Start first	Frequent workflow, clear knowledge sources, measurable service outcomes and manageable human review.
2	Internal Knowledge Assistant	4.1 / 5	Prepare next	Strong value, but content ownership and access permissions should be cleaned up first.
3	Purchase Order Extraction	4.0 / 5	Pilot after first use case	Good automation potential; needs document samples, field rules and ERP integration details.
4	Sales Forecasting	3.3 / 5	Validate data first	Potential value is high, but historical data quality and forecasting baseline need validation.
5	Inventory & Reorder Insights	3.0 / 5	Validate data first	Useful direction, but stock and lead-time data currently lives in three disconnected sheets.
6	Fully Autonomous Customer Agent	2.7 / 5	Do not start here	Too much decision and communication risk before the business has proven a controlled assisted workflow.
7	AI Resume Screening	2.3 / 5	Not recommended now	Low hiring volume today means limited near-term value versus the fairness and compliance risk of screening automatically.

LEADERSHIP TAKEAWAY The scoring is designed to stop “interesting AI ideas” from being mistaken for good first projects. The first pilot should be valuable, feasible, sufficiently data-ready, manageable in implementation complexity and controllable in risk Customer Enquiry Assistant is the only opportunity where every criterion scores 4 or above, making it the strongest balanced candidate for the first pilot.

03 / OPPORTUNITY REVIEW

Customer Enquiry Assistant

Overall score: 4.5 / 5 | Decision: Start first

Criterion	Rating	Reason
Business value	5	Repeated enquiry handling affects both customer response and staff time — roughly 640 enquiries a month.
Feasibility	4	A controlled workflow can be limited to selected enquiry types and the two highest-volume channels.
Data readiness	4	Useful knowledge exists, but it needs ownership and cleanup across ~40 shared-drive folders.
Implementation complexity	5	A narrow first pilot can be contained to two high-volume channels, selected enquiry types and human review, keeping implementation relatively bounded.
Risk	4	Human approval and source citations can control early-stage risk; pricing and complaints stay human.

DECISION LOGIC Recommended because it gives the clearest combination of business value, measurement and control. It is also the only opportunity where every criterion scores 4 or above.

EVIDENCE SOURCE Based on: enquiry-volume estimates from the CRM export, shadowing of two customer-service agents, and the anonymous 11-response team survey.

04 / OPPORTUNITY REVIEW

Internal Knowledge Assistant

Overall score: 4.1 / 5 | Decision: Prepare next

Criterion	Rating	Reason
Business value	4	Reduces internal search time and dependence on 4–5 experienced individuals.
Feasibility	4	Technically straightforward once document scope is controlled.
Data readiness	3	Documents are available but ownership, duplicates and permissions need cleanup first.
Implementation complexity	5	A focused document set and internal-only first version keep implementation comparatively contained once ownership and permissions are resolved.
Risk	5	Internal-only use, source citations and access controls make this manageable.

DECISION LOGIC Strong second use case because the same knowledge preparation can support both internal and customer-facing workflows — the two pilots share one underlying knowledge base.

EVIDENCE SOURCE Based on: shared-drive walkthrough with the IT Manager, and interviews describing 6–8 hours/week lost to internal knowledge search.

05 / OPPORTUNITY REVIEW

Purchase Order Extraction

Overall score: 4.0 / 5 | Decision: Pilot after first use case

Criterion	Rating	Reason
Business value	4	Reduces manual reading and re-keying of structured information — roughly 110 POs a week.
Feasibility	4	Works well when document types and target fields are defined.
Data readiness	4	Historical examples are usually available, but quality must be checked.
Implementation complexity	4	One document type, defined fields and one downstream process keep the first pilot bounded.
Risk	4	Confidence thresholds and human exception review can control errors.

DECISION LOGIC Worth pursuing, but the first pilot should remain focused on one document format and one downstream process (Tally ERP) before expanding to other document types.

EVIDENCE SOURCE Based on: review of 14 sample purchase-order documents and an interview with the Operations Lead on current error rates (≈ 2–3%).

06 / OPPORTUNITY REVIEW

Sales Forecasting

Overall score: 3.3 / 5 | Decision: Validate data first

Criterion	Rating	Reason
Business value	5	Better forecasting could materially improve planning ahead of automotive OEM production cycles.
Feasibility	3	Requires careful model design and process adoption from the sales team.
Data readiness	2	Historical consistency and feature availability are not yet confirmed in Tally.
Implementation complexity	2	Requires historical-data validation, forecasting design, a meaningful evaluation period and adoption within the sales process.
Risk	3	Forecasts can support decisions, but should not silently replace business judgement.

DECISION LOGIC Do not reject it; first validate data quality and baseline forecasting performance before committing to a build.

EVIDENCE SOURCE Based on: a 20-minute Finance interview and a light review of exportable order-history fields in Tally ERP.

06 / OPPORTUNITY REVIEW (CONTINUED)

Inventory & Reorder Insights

Overall score: 3.0 / 5 | Decision: Validate data first

Criterion	Rating	Reason
Business value	4	Better reorder timing could reduce both stockouts and excess raw-material holding.
Feasibility	3	Achievable, but requires combining data currently spread across three disconnected sheets.
Data readiness	2	Stock levels, lead times and consumption rates are not consistently recorded today.
Implementation complexity	2	Requires data consolidation across three disconnected sheets before a reliable pilot can be implemented.
Risk	3	Low risk in principle, since recommendations would remain advisory rather than automatic ordering.

DECISION LOGIC Directionally promising, but this is a data-quality project before it is an AI project. Revisit once inventory data is consolidated into one system of record.

EVIDENCE SOURCE Based on: a walkthrough of three separate inventory spreadsheets with the Operations Lead; no single source of truth was found.

07 / OPPORTUNITY REVIEW

Fully Autonomous Customer Agent

Overall score: 2.7 / 5 | Decision: Do not start here

Criterion	Rating	Reason
Business value	4	Could eventually reduce manual handling substantially.
Feasibility	2	End-to-end autonomy across real customer cases creates many edge cases.
Data readiness	3	Knowledge exists, but exception coverage is not mature enough.
Implementation complexity	2	End-to-end autonomy creates many edge cases, escalation paths and testing requirements, making the first implementation complex.
Risk	1	Customer commitments, incorrect answers and escalation errors create avoidable risk.

DECISION LOGIC Start with an assistant that drafts and routes work. Increase automation only after the evidence shows where it is safe — this is the natural “phase 2 or later” version of Opportunity #1.

EVIDENCE SOURCE Based on: leadership's stated ambition during the AI Fit Call, weighed against the risk profile observed in customer-facing interviews.

07 / OPPORTUNITY REVIEW (CONTINUED)

AI Resume Screening

Overall score: 2.3 / 5 | Decision: Not recommended now

Criterion	Rating	Reason
Business value	2	Hiring volume is low (≈ 10–15 roles/year), limiting near-term value.
Feasibility	3	Technically available off-the-shelf, but would need careful bias and fairness testing.
Data readiness	2	No structured historical hiring data exists to validate screening criteria against.
Implementation complexity	2	Off-the-shelf components exist, but fairness testing, governance and the low hiring volume make the implementation burden high relative to the value.
Risk	2	Automated screening carries fairness, compliance and reputational risk disproportionate to the benefit here.

DECISION LOGIC This idea was raised once during interviews but is not a good fit for this business's current hiring volume. Revisit only if hiring scales materially.

EVIDENCE SOURCE Based on: a single mention during the Managing Director interview; not independently validated with HR data.

08 / DECISION GUIDE

What to do with the scorecard

The number is only a comparison tool. Leadership still needs to confirm scope, owner, data and acceptable risk.

01	**Choose one first pilot** Do not start several AI projects at the same time unless the business already has the delivery capacity.
02	**Confirm the assumptions** Validate volumes, source data, integrations and process ownership behind the chosen score.
03	**Set a baseline** Measure the current workflow before development starts.
04	**Use the 90-Day Pilot Roadmap** Turn the selected opportunity into 30-day readiness, controlled pilot execution, evaluation evidence and an explicit proceed / revise / pause decision.
05	**Re-score after the pilot** New evidence may make a previously lower-ranked opportunity more attractive.

Common scoring mistakes to avoid

- Scoring business value based on how impressive the idea sounds in a meeting, rather than measured frequency and business impact of the underlying problem.
- Rating data readiness as high simply because a spreadsheet exists somewhere, without checking whether it is current, complete or owned.
- Letting one enthusiastic stakeholder's favourite idea skip the scoring process entirely.

APPENDIX / FULL SCORING WORKSHEET

All seven opportunities, side by side

A single reference table for workshops and leadership review — the same ratings shown on the individual pages above, gathered in one place.

Opportunity	Value (30%)	Feasibility (25%)	Data readiness (20%)	Complexity (15%)	Risk (10%)	Total
Customer Enquiry Assistant	5	4	4	5	4	4.5
Internal Knowledge Assistant	4	4	3	5	5	4.1
Purchase Order Extraction	4	4	4	4	4	4.0
Sales Forecasting	5	3	2	2	3	3.3
Inventory & Reorder Insights	4	3	2	2	3	3.0
Fully Autonomous Customer Agent	4	2	3	2	1	2.7
AI Resume Screening	2	3	2	2	2	2.3

HOW TO RE-USE THIS TEMPLATE This same five-question, weighted-scoring method can be applied to future AI opportunities: business value, feasibility, data readiness, implementation complexity and risk. Score consistently, show the evidence, and let the total guide — not dictate — the decision.