TOOLKIT

AI Product Discovery Canvas

A structured framework for discovering, validating, and prioritizing AI product opportunities — with dedicated lenses for B2B, B2C, technical feasibility, and AI-specific requirements.

How to use this canvas

Work through the 5 zones in order. Zone 1 is mandatory for every product. Apply Zone 2 or Zone 3 based on your go-to-market model — or both if your product serves both. Always complete Zones 4 and 5 before committing to build. Use the Prioritization Output at the end to make a Go / Refine / Stop decision.

Zone	Name	Apply when
1	Core	Always — every product
2	B2B lens	Product sold to businesses
3	B2C lens	Product sold to individuals
4	Technical lens	Always — before committing to build
5	AI deep lens	When AI is a core component of the solution

1

CORE — UNIVERSAL

Problem, evidence, ICP, and hypothesis

Every product starts here. This zone forces you to articulate the problem with precision before jumping to a solution.

Key insight: A vague problem statement produces a vague product. If you cannot state the problem in one sentence without mentioning your solution, you are not ready to build.

Block	What it captures	Key question to answer
Problem statement	One sentence: who has what problem, in what context	Can I state this without mentioning my product?
Evidence	Data, signals, or research confirming the problem exists	Is this observed — or assumed?
ICP	Who specifically — segment, digital maturity, context	Would 10 people all fit this profile?

Current alternatives	What they do today — workarounds, competitors, nothing	Why haven't they solved this themselves?
Cost of problem	Time, money, or risk lost because it is unsolved	Can I quantify this in units they care about?
Hypothesis	If we build X, Y happens, measured by Z	Is this falsifiable — can I disprove it?

2

B2B LENS — PRODUCTS SOLD TO BUSINESSES

Buyer, pain, ROI, and sales dynamics

In B2B, the person who experiences the problem is almost never the person who buys the solution. Misidentifying the buyer is the most common reason B2B products fail to close deals even when users love them.

Key insight: The economic buyer cares about ROI and risk reduction. The end user cares about speed and ease. Your product must satisfy both — but your pitch must speak to the buyer first.

Block	What it captures	Key question to answer
Economic buyer	Who signs — not always who uses the product	What is their primary success metric?
End user	Who uses it daily — different needs from the buyer	What makes them love or resist this?
Org pain	Revenue, cost, compliance, or risk impact	Can I tie this to a P&L; line item?
Sales cycle	Decision length, number of approvals needed	How many people need to say yes?
ROI frame	Value in numbers the buyer cares about	What is the cost of inaction?
Integration risk	Systems to connect and who owns them	Who will push back on integration?

3

B2C LENS — PRODUCTS SOLD TO INDIVIDUALS

Emotion, habit, trust, and retention

B2C products live and die on emotion and habit. Users decide in seconds whether to try something and in days whether to keep using it.

Key insight: Trust is the first conversion barrier in B2C. Design your trust signal before your feature set.

Block	What it captures	Key question to answer
Emotional trigger	The feeling or moment driving the user to look	What emotion happens before they open the app?
Behavioral pattern	Existing habits to attach to or replace	What do they already do this fits alongside?

Trust barrier	Fear, unfamiliarity, cost, or privacy concerns	What makes a skeptic say "not for me"?
Viral loop	How one user naturally brings another	Does the product get more valuable when shared?
Retention hook	Habit, value delivery, social, or FOMO	Why open this tomorrow without a push?
Price psychology	Want or need — how users justify the spend	Monthly, one-time, or success-based?

4

TECHNICAL LENS — FEASIBILITY BEFORE COMMITMENT

Data, integrations, build vs. buy, and scale

Most product failures are not discovery failures — they are feasibility failures discovered too late. This lens forces honest conversations about what you can actually build.

Key insight: MVP shortcuts that seem harmless in month 1 become refactoring nightmares in month 6. Name them now.

Block	What it captures	Key question to answer
Data availability	What exists, where it lives, who owns it	Do we have the data — or must we create it?
Integration surface	APIs, databases, third-party tools required	Who controls the systems we need?
Build vs. buy	What to build in-house vs. purchase or compose	What is core to our differentiation?
Scalability flag	Where load or volume becomes a problem	At 10x users, where does this break first?
Tech debt risk	MVP shortcuts that hurt the team in 6–12 months	What are we deciding now that we'll undo later?
Feasibility score	1–5: 1 = very low effort, 5 = very high effort	Is effort proportionate to problem severity?

5

AI DEEP LENS — FOR AI-CORE PRODUCTS

Necessity, data, model type, risk, and feedback

The most important AI product question is the first one: is AI actually necessary? Many products called "AI" are better served by deterministic logic.

The AI necessity test: If a rule-based system could solve 90% of this problem at 10% of the cost and risk — you probably do not need AI. The remaining 10% is where AI genuinely earns its place.

Block	What it captures	Key question to answer

Is AI necessary?	Can this be solved without AI — and should it be?	Would rules solve 90% of this problem?
AI value prop	What AI uniquely enables vs. a deterministic system	What breaks without AI that works with it?
Data requirements	Volume, quality, labelling, and freshness needs	Do we have enough clean, representative data?
Model type	Classification, generation, retrieval, recommendation	Are we predicting, generating, or ranking?
AI risk	Hallucination, bias, explainability, regulatory exposure	What is the cost of a wrong AI output?
Feedback loop	How the model improves — signals that retrain it	How does the model get smarter over time?

Prioritization output

Score each signal 1–5. Problem severity and ICP clarity below 3 = refine first. Feasibility below 2 relative to severity = stop.

Signal	What to score	Score (1–5)	Weight
Problem severity	How painful is this for the ICP?	_ / 5	High
ICP clarity	How precisely can you describe who has this problem?	_ / 5	High
Solution feasibility	How buildable given current team and resources?	_ / 5	Medium
AI necessity	How much does the solution depend on AI specifically?	_ / 5	Medium
Market timing	Is the problem acute enough to drive action now?	_ / 5	Medium

Decision	When	Next step
GO	Avg 3.5+ across all signals, severity and ICP both 4+	Define MVP scope, assign team, set KPIs
REFINE	Avg 2.5–3.5, or any high-weight signal below 3	Return to discovery — more interviews or technical spikes
STOP	Problem severity below 2, or feasibility below 2 vs severity	Document learnings, free the team, revisit next cycle

AI Product Discovery Canvas

theglocalpm.com · Ali Mahmoud

1 CORE — UNIVERSAL (ALWAYS COMPLETE THIS ZONE)

PROBLEM STATEMENT
Who has what problem, in what context

EVIDENCE
Data or signals confirming the problem

ICP
Who specifically has this problem

CURRENT ALTERNATIVES
What they do today instead

COST OF PROBLEM
Time, money, or risk lost

HYPOTHESIS
If we build X, Y happens, measured by Z

2 B2B LENS — FOR PRODUCTS SOLD TO BUSINESSES

ECONOMIC BUYER
Who signs — not always who uses

END USER
Who uses daily — different needs

ORG PAIN
Revenue, cost, compliance, risk

SALES CYCLE
Decision length, approvals

ROI FRAME
Value in numbers buyer cares about

INTEGRATION RISK
Systems to connect, who owns them

3 B2C LENS — FOR PRODUCTS SOLD TO INDIVIDUALS

EMOTIONAL TRIGGER
Feeling driving them to look

BEHAVIORAL PATTERN
Habits to attach to or replace

TRUST BARRIER
Fear, cost, unfamiliarity

VIRAL LOOP
How one user brings another

RETENTION HOOK
Habit, value, social, or FOMO

PRICE PSYCHOLOGY
Want or need — how they justify spend

4 TECHNICAL LENS — ALWAYS COMPLETE BEFORE COMMITTING TO BUILD

DATA AVAILABILITY
What exists, where, who owns it

INTEGRATION SURFACE
APIs, DBs, third-party tools

BUILD VS. BUY
Core vs. commodity decisions

SCALABILITY FLAG
Where 10x load breaks this

TECH DEBT RISK
MVP shortcuts that hurt later

FEASIBILITY SCORE
1 (easy) to 5 (very hard)

5 AI DEEP LENS — WHEN AI IS A CORE COMPONENT

IS AI NECESSARY?
Could rules solve 90% of this?

AI VALUE PROP
What AI uniquely enables here

DATA REQUIREMENTS
Volume, quality, freshness

MODEL TYPE
Classify / generate / retrieve / rank

AI RISK
Hallucination, bias, explainability

FEEDBACK LOOP
How model improves over time

PRIORITIZATION OUTPUT — Score 1–5, then decide:

PROBLEM SEVERITY	ICP CLARITY	FEASIBILITY	AI NECESSITY	MARKET TIMING	DECISION
○ ○ ○ ○ ○	○ ○ ○ ○ ○	○ ○ ○ ○ ○	○ ○ ○ ○ ○	○ ○ ○ ○ ○	GO REFINE STOP