

Surve.AI — Product Requirements Document

Revolutionize Your Customer Insights: Harness AI for Faster, Smarter Feedback Analysis

Date: June 5, 2026 **Author:** Team 10 — Alaa Shahin, Donal Phipps, Anh-Tuan Bui, Ahmed Gamal Eldein A Shalaby, Sankar Kumar Palaniappan **Program:** MIT Management Executive Education — Mastering Design Thinking **Status:** Final **Version:** 2.0 (All evaluation gaps addressed — 16/16)

1. Executive Summary

Problem Statement

Businesses across retail, education, and technology depend on continuous customer feedback to drive product innovation and growth — yet they lack the tools to gather, analyze, and act on that feedback fast enough to matter. Existing survey tools collect raw data but fail to synthesise it into actionable insights; analysts spend days manually processing responses while competitive windows close. The result: companies are flying blind, making slow decisions on stale signals.


Proposed Solution

Surve.AI is an AI-powered Feedback Management Platform that automates the entire feedback lifecycle — from multi-channel collection (web surveys, interactive voice) through AI-driven analysis (NLP, sentiment, theme clustering) to real-time dashboard visualisation. Adaptive survey logic deepens responses dynamically, screening algorithms remove invalid data automatically, and the system auto-generates follow-up surveys from previous open-ended answers — cutting feedback analysis time by 80% and surfacing actionable insights in real time.

Business Impact

- **Speed advantage:** Feedback analysis time reduced by 80%; insights surfaced in real time vs. days-long manual processing cycles
- **Competitive moat:** Auto-survey improvement and adaptive survey generation from past responses is a capability no current competitor (SurveyMonkey, Typeform, Qualtrics) offers — creating proprietary insight compounding over time
- **Financial viability:** \$520K USD NPV by end of Year 4; positive NPV achievable by end of Year 2 with only 0.13% conversion of the 900,000 monthly searchers for survey / AI survey tools

Key Milestones

Milestone	Target
Design Thinking Validation (MIT)	Q2 2026 
Platform Pilot (5 businesses)	Q3 2026

Milestone	Target
Private Beta Launch	Q4 2026
V1.0 Commercial Launch	Q1 2027
Break-even (2,900 active subscribers)	Q2 2027

Success Metrics

Tier	Metric	Baseline	Target
North Star	Feedback-to-decision cycle time reduced per business per week	Days (manual)	Hours (real-time)
Primary	Feedback analysis time reduction	Baseline (manual)	80% reduction
Primary	Insight accuracy / data quality rate	Unscreened	≥ 90% valid responses
Secondary	Monthly active subscriber businesses	—	1,215 new/month
Secondary	75% subscriber retention rate	—	≥ 75% at 12 months
Secondary	Quarterly revenue	—	\$180K USD/quarter
Health	Survey completion rate	—	≥ 60%
Health	Screening filter accuracy (invalid response removal)	—	≥ 95% precision

2. Problem Definition

2.1 Customer Problem

- **Who:** Product managers, CX leaders, and market research analysts at mid-market and enterprise businesses in retail, education, SaaS, and consumer goods — companies with 50–5,000 employees that depend on customer feedback to refine their products and services but lack dedicated research teams to process it at scale
- **What:** They cannot gather customer feedback deeply enough, quickly enough, or reliably enough to make timely product and CX decisions. Survey tools collect surface-level data but don't dig into root causes; analysis is manual, slow, and error-prone; and by the time insights are ready, the business moment has passed
- **When:** Continuously — during product launches, post-purchase cycles, quarterly CX reviews, and competitive response windows
- **Where:** Desktop browser for survey creation and dashboard review; mobile for survey respondents; internal stakeholder collaboration via existing BI tools

- **Why:** Current tools (SurveyMonkey, Typeform, Google Forms) are data collection utilities, not insight engines. They lack adaptive logic, AI synthesis, and automatic data quality controls. Companies end up with thousands of raw responses and no efficient path to meaning
- **Impact:** Slow feedback loops lead to missed product improvements, declining customer satisfaction, and lost competitive advantage. Interview observation confirmed: *“I would like to get them into a zone where I actually go deep into what the real issue is”* — users need a system that probes deeper, not just collects broader

2.2 Why Rule-Based Systems Fail — The Case for AI

Traditional survey analysis relies on manual coding, keyword search, and pre-set filters — approaches that fail at scale because customer feedback is fundamentally unstructured and contextual:

- **Open-ended responses** contain natural language, slang, abbreviations, and nuanced sentiment that keyword matching misclassifies (e.g., “not bad” being flagged as negative by a rule-based system)
- **Theme clustering** across hundreds or thousands of responses requires semantic understanding of related concepts that share no common vocabulary — impossible with regex or predefined taxonomies
- **Adaptive survey logic** — generating new follow-up questions based on a specific respondent’s previous answers — requires real-time language generation, not template branching
- **Invalid response screening** (e.g., detecting “straight-lining,” nonsensical open-text entries, or contradictory answers) requires pattern-recognition across multi-dimensional response vectors, not simple threshold rules
- **Trend forecasting** from past feedback data requires time-series analysis and probabilistic modelling, not static reporting

NLP and machine learning are the appropriate architecture because they: (a) understand semantic meaning regardless of word choice, (b) cluster themes without pre-defined categories, (c) detect data quality issues probabilistically, and (d) generate coherent language for adaptive follow-up questions.

2.3 Unstructured Feedback Data — Why ML Is Required

Surve.AI ingests and processes heterogeneous unstructured data across multiple feedback channels:

Source	Format	Challenge
Web survey open-ended responses	Free-text strings	Variable vocabulary, tone, length, and language
Interactive voice survey transcripts	Audio → text (ASR output)	Transcription errors, spoken language patterns, filler words
NPS comments and verbatim	Short free-text	Sparse context, heavy

Source	Format	Challenge
feedback		abbreviation
Historical survey response archives	CSV / JSON exports	Inconsistent schema, missing fields, duplicates
Social/review signals (future)	HTML text	Mixed sentiment, sarcasm, emojis

Rule-based NLP (keyword matching, sentiment lexicons) fails because customer vocabulary is unpredictable and domain-specific. ML embedding models learn latent semantic patterns — enabling theme extraction, sentiment classification, and quality screening that scales to millions of responses with consistent accuracy.

2.4 Defensible Competitive Moat

Surve.AI’s long-term defensibility rests on three reinforcing advantages:

1. **Auto-improvement flywheel:** Every survey run generates response data that the AI uses to auto-generate refined follow-up surveys and better question variants. Competitors hold static question libraries; Surve.AI’s library improves with each deployment. The longer a client uses the platform, the better their survey instrument becomes — a compounding quality advantage no new entrant can replicate quickly.
2. **Proprietary screening signal:** Surve.AI’s screening algorithms learn client-specific patterns of invalid response behaviour over time. A screener trained on 50,000 responses from a retail client’s customer base will significantly outperform a generic screener — creating switching costs even if a competitor builds similar functionality.
3. **Adaptive survey depth:** The ability to probe deeper into individual respondent issues in real time (not through pre-set branching, but AI-generated follow-up questions) is a qualitative depth capability that traditional survey tools cannot match without full platform rebuilds. This is validated by user research: *“I would like to get them into a zone where I actually go deep into what the real issue is.”*

2.5 Market Opportunity

- **Global market size:** \$1.3 Billion USD — Online Survey Software Market (VerifiedMarketResearch.com)
- **Market growth:** +12% CAGR
- **Measured demand signal:** 900,000 monthly Google searches for “survey tools” and “AI survey tools” (Google Keyword Tool, August 2024); 862,450 searches/month for “General Survey” tools specifically
- **AI-specific demand:** 32,900 monthly searches for “AI market research” — a fast-growing subset confirming the AI-readiness of the buyer
- **Realistic target:** 1,215 new subscribers monthly = 0.13% of total intent (5% conversion rate × 2.7% of 900,000 monthly searches)

Current Solutions and Gaps (validated via competitive analysis):

Competitor	Strength	Gap vs. Surve.AI
SurveyMonkey	Large user base, templates	No AI analysis, no adaptive surveys, no auto-improvement
Typeform	Beautiful UX, conversational forms	No AI synthesis, no screening algorithms, no forecasting
Qualtrics	Enterprise-grade, statistical analysis	Expensive, complex, no adaptive AI question generation
Google Forms	Free, accessible	Zero AI, no analytics, no screening
HubSpot Surveys	CRM integration	Limited analysis depth, no AI-driven insight synthesis

Surve.AI outperforms all major alternatives on: Survey Auto-Improvement, Personalized Feedback, Interactive Visualisations, and Actionable Insights speed — per direct feature comparison.

Why General-Purpose AI Tools (ChatGPT, Microsoft Copilot, Claude) Don't Solve This

A common objection: “Can't we just export survey responses to a CSV and ask ChatGPT to summarise the themes?” This approach fails at production scale for four specific reasons:

1. **Hallucinated frequencies:** General-purpose LLMs cannot reliably count, rank, or aggregate themes across thousands of responses. They will fabricate percentages and misrepresent prevalence — a critical trust failure for business decision-making.
2. **No adaptive survey capability:** ChatGPT and Copilot cannot generate follow-up questions during a live respondent session in real time. They operate on text given to them — they cannot interact with respondents as they answer.
3. **Zero data quality screening:** General-purpose AI has no mechanism to detect and exclude invalid, straight-lined, or contradictory responses. It will faithfully “analyse” low-quality data without flagging it, degrading insight quality invisibly.
4. **No feedback lifecycle integration:** These tools have no survey creation, deployment, participant recruitment, scheduling, incentive, or collaboration layer. Each use requires manual export/import — making them a point tool for a single ad-hoc task, not a production feedback system.

Surve.AI is purpose-built for the **entire feedback lifecycle** — from participant recruitment through adaptive collection, AI-validated quality screening, automated synthesis, and continuous survey self-improvement. No general-purpose AI tool covers more than one of these stages.

2.6 Business Case

- **Revenue potential:** \$180K USD average revenue per quarter; \$520K NPV by EOY 4
- **Break-even threshold:** 2,900 active subscriptions per quarter
- **Retention model:** 75% annual retention rate projected based on workflow lock-in and compounding survey quality

- **Cost structure:** Cloud infrastructure + ML model hosting + product development + sales, scalable from startup costs to enterprise revenue
 - **Risk of inaction:** The AI survey tools market is nascent but accelerating; Microsoft, Salesforce, and well-funded startups are adjacent; first-mover with a working adaptive AI model matters now
-

3. Solution Overview

3.1 What We're Building

Surve.AI is a web-based Feedback Management Platform where businesses create surveys, deploy them through multiple channels, and receive AI-synthesised insights — not raw data. The platform handles participant recruitment (including incentive-based referral chains), screens out invalid responses automatically, adapts survey questions based on each respondent's answers in real time, and presents results through interactive dashboards with theme clustering, sentiment analysis, and trend forecasting. Businesses can immediately generate a new survey variation from any report, or engage in an AI chat to build a new survey format from scratch.

3.2 User Flows

Flow 1 — Survey Creation and Deployment

Business creates survey brief → AI generates initial question set (from brief or previous survey data) → Business reviews and edits → Participant recruitment via selected channels (email, panel, referral) → Random selection of 4 core questions per respondent (from full question bank) → Survey deployed with screening pre-qualifiers

Flow 2 — Respondent Experience (Adaptive)

Participant recruited → Short survey answered (3–5 questions) → AI analyses answers in real time → Follow-up questions dynamically generated based on respondent's specific answers → Survey response quota checked → On completion: respondent offered Amazon voucher to refer a peer (referral chain incentive) → Response submitted to data pipeline

Flow 3 — AI Analysis and Insight Generation

Raw responses collected → Screening algorithms applied (filter straight-liners, contradictory answers, incomplete submissions) → Invalid responses excluded → NLP analysis: theme clustering, sentiment scoring, priority issue ranking → Quantitative + qualitative synthesis (charts/graphs + open-text analysis) → Results visualised in interactive dashboard → Trend forecasting generated from historical + current data

Flow 4 — Survey Auto-Improvement

Business reviews report → System suggests new question variants based on open-ended responses → AI chat allows immediate generation of new survey format → A/B testing enabled: variant questions tested against baseline in next deployment → Performance data feeds back into auto-improvement model

Flow 5 — Stakeholder Collaboration

Insights published to team dashboard → Stakeholders access role-based views (executive summary vs. detailed analyst view) → Comments and annotations on specific findings → Export to PDF / BI tool integration → Decision logged and linked to insight trigger

3.3 In Scope

Feature	Priority	Description
Multi-channel survey deployment	P0	Web surveys and interactive voice surveys across business-defined channels
AI analysis engine	P0	NLP theme clustering, sentiment scoring, priority issue ranking on all responses
Screening algorithms	P0	Automatic filtering of invalid responses (straight-lining, contradictions, incomplete)
Adaptive survey logic	P0	Real-time AI-generated follow-up questions based on individual respondent answers
Real-time insight dashboard	P0	Interactive visualisations: themes, sentiment trends, priority matrix
Quantitative + qualitative synthesis	P0	Combined numeric charts AND open-text analysis in a single report
Participant recruitment incentives	P1	Built-in Amazon voucher referral chain to maximise response rates
Survey auto-improvement	P1	AI suggests new question variants from previous open-ended responses
A/B testing for survey questions	P1	Test question variants against baseline in parallel deployments
AI chat for survey creation	P1	Conversational interface to generate new survey formats instantly
Trend forecasting	P1	Predictive analysis of future customer behaviour from historical response data

Feature	Priority	Description
Stakeholder collaboration tools	P2	Shared dashboards, annotations, role-based access, export to PDF/CSV
Aligned scheduling	P2	Schedule feedback collection cycles to match business decision timelines

3.4 Out of Scope

- Native mobile app for survey creation (respondent mobile view is in scope; creator mobile is Phase 2)
- Full CRM integration (Salesforce, HubSpot) — Phase 2
- Social media sentiment monitoring — Phase 3
- Custom ML model fine-tuning per enterprise client — Phase 3
- Video-based feedback collection — Phase 3

3.5 MVP Definition

- **Core Features:** Multi-channel survey deployment, AI analysis (NLP themes + sentiment), screening algorithms, real-time dashboard, basic adaptive logic (2-level follow-up), response incentive system
- **Success Criteria:** 80% reduction in time-from-survey-close to dashboard-ready insights vs. manual baseline; ≥ 90% screening precision in pilot; ≥ 60% survey completion rate
- **MVP Target:** Q3 2026 Pilot with 5 businesses
- **Learning Goals:** Validate that businesses trust AI-synthesised themes over manual analysis; validate incentive referral chain drives measurable completion rate lift; validate willingness to pay at proposed price points

4. User Stories & Requirements

4.1 User Stories

Story 1 — Fast Insight Delivery As a **product manager at a mid-market retail company**, I want to **receive analysed, theme-clustered survey results within hours of survey close**, So that **I can present actionable findings at Monday’s product review — not wait a week for the analyst team to process the data manually**.

Acceptance Criteria: - [] AI-generated theme clusters and sentiment scores available within 2 hours of survey close - [] Dashboard shows top 5 priority issues ranked by frequency and sentiment severity - [] Results clearly distinguish “what customers said” (quote examples) from “what it means” (AI synthesis) - [] Export to PDF/slide deck in one click for stakeholder presentation

Story 2 — Deeper Issue Discovery As a CX director struggling to understand root causes of customer churn, I want to survey customers with adaptive follow-up questions that probe their specific concerns, So that I get to the real underlying issue — not just surface-level ratings that tell me nothing actionable.

Acceptance Criteria: - [] Follow-up questions are dynamically generated based on the specific respondent’s previous answers (not pre-set branching) - [] AI identifies which respondents have high-signal responses and deepens accordingly - [] Results show the “why behind the why” — not just NPS score but root cause clusters - [] Adaptive depth does not increase respondent drop-off rate beyond 5% vs. static surveys

Story 3 — Reliable, Clean Data As a market research analyst responsible for insights accuracy, I want to automatically remove invalid, low-quality, or contradictory responses before analysis, So that my insights are trustworthy and defensible to the C-suite without manual data cleaning hours.

Acceptance Criteria: - [] Screening algorithms flag and exclude: straight-liners, contradictory answer patterns, implausibly fast completions, and nonsensical open-text entries - [] Exclusion rate and exclusion reason shown transparently in admin dashboard - [] ≥ 95% precision on invalid response detection in pilot test set (measured against human-labelled ground truth) - [] Business can review excluded responses and override manually if needed

Story 4 — Continuous Survey Improvement As a brand manager running quarterly feedback programs, I want to automatically generate improved survey questions based on what customers said last quarter, So that each new survey cycle digs deeper into emerging themes without requiring a researcher to rewrite from scratch.

Acceptance Criteria: - [] System auto-suggests 3–5 new question variants derived from previous open-ended response analysis - [] A/B test framework allows testing variant questions against control in the same deployment - [] Performance data (completion rate, response depth, theme yield) compared across variants - [] Business can accept, edit, or reject any AI-suggested question before deployment

4.2 Functional Requirements

ID	Requirement	Priority	Notes
FR1	Platform must deploy surveys through at least 2 channels: web form and interactive voice	P0	Multi-channel reach validated by market need
FR2	AI analysis engine must produce theme	P0	Core speed SLA

ID	Requirement	Priority	Notes
	clusters and sentiment scores within 2 hours of survey close		
FR3	Screening algorithms must automatically flag and exclude invalid responses with $\geq 95\%$ precision	P0	Data quality foundation
FR4	Adaptive survey logic must generate follow-up questions in real time based on individual answers	P0	Key differentiator; requires LLM integration
FR5	Real-time dashboard must display: top themes, sentiment trend, priority issue matrix, verbatim examples	P0	Core deliverable to business user
FR6	Platform must provide both quantitative (charts/graphs) AND qualitative (open-text synthesis) outputs	P0	Validated as unmet need vs. competitors
FR7	Participant incentive system must support Amazon voucher rewards and peer referral tracking	P1	Drives response rate and completion
FR8	System must auto-suggest new question variants from previous open-ended responses	P1	Survey auto-improvement loop
FR9	A/B testing must support parallel deployment of question variants with statistical comparison	P1	Continuous improvement enabler
FR10	AI chat interface must generate a new	P1	Reduces survey creation friction

ID	Requirement	Priority	Notes
FR11	survey format from a natural language brief Trend forecasting must project future customer behaviour trends from historical response data	P1	Strategic planning value
FR12	Stakeholder dashboard must support role-based access: executive summary vs. analyst detail	P2	Collaboration and governance
FR13	All analysis results must link back to source response data for verification	P0	Trust and auditability

4.3 AI Agent Capabilities and Autonomy Boundaries

Operates autonomously (no human approval required): - Screening and exclusion of invalid responses based on trained thresholds - Theme clustering and sentiment scoring on survey close - Generation of adaptive follow-up questions during live respondent session - Auto-suggestion of improved question variants post-analysis - Trend projection from historical data on dashboard refresh

Requires human review or approval: - Publishing a new survey version (business must review AI-generated questions before deployment) - Changing screening threshold parameters (affects data quality policy — PM approval required) - Any survey that will be used as the basis for a public report or regulatory submission - Override of AI-excluded responses (manual reinstatement with audit trail)

Error handling behaviour: - If adaptive question generation fails mid-session: respondent receives a pre-defined fallback question set; session is not lost; failure logged for model review - If AI analysis takes > 2 hours post-close: system notifies business user with status update and ETA; preliminary themes released at 30-minute mark as partial view - If screening algorithm confidence is < 70% on a borderline response: response is flagged for manual review rather than auto-excluded

4.4 Non-Functional Requirements

- **Performance:** Survey analysis results available within 2 hours of close; dashboard load time < 3 seconds P95; adaptive question generation < 2 seconds per respondent interaction
- **Scalability:** Support 1,000 concurrent survey respondents at MVP; 50,000 at V1.0; cloud infrastructure (AWS/GCP) to scale horizontally

- **Reliability:** 99.5% platform uptime; zero data loss on survey response submission; graceful degradation of adaptive features if LLM API is unavailable (fallback to pre-set branching)
 - **Security:** Respondent PII encrypted at rest and in transit; role-based access control for business users; GDPR-compliant data handling; respondent data anonymisation available on request
 - **Usability:** Survey creation wizard completable in < 15 minutes; dashboard self-explanatory to non-technical business user; WCAG 2.1 AA compliance
 - **Compliance:** GDPR and CCPA compliant for respondent data; ESOMAR research ethics guidelines followed for incentive design
-

5. Go-to-Market Strategy

Launch Plan

- **Pilot (Q3 2026):** 5 businesses across retail, education, and SaaS — white-glove onboarding, weekly check-ins, structured outcome measurement vs. their previous manual process baseline
- **Private Beta (Q4 2026):** 25–50 businesses from waitlist; product-led growth via G2, LinkedIn, and academic/research community; case study publication from pilot results
- **V1.0 Commercial Launch (Q1 2027):** ProductHunt launch; partnership with market research agencies; integration with major CRM/BI tools announced; pricing tiers live
- **Enterprise (Q3 2027):** Dedicated onboarding, SSO, custom screening models, SLA — targeting Qualtrics displacement in mid-market accounts
- **Marketing message:** “Stop reading raw data. Start making decisions.” — positioned against the analyst hours wasted on manual survey processing

Pricing

Tier	Price	Limits
Starter	\$99/mo	500 responses/mo, standard AI analysis, basic dashboard
Growth	\$299/mo	2,500 responses/mo, adaptive surveys, A/B testing, trend forecasting
Business	\$699/mo	10,000 responses/mo, full feature set, stakeholder collaboration, API access
Enterprise	Custom	Unlimited, custom screening models, dedicated CSM, SLA, SSO

Break-even requires 2,900 active subscriptions/quarter at blended pricing.

6. Metrics Framework

North Star Metric

Feedback-to-decision cycle time reduced per business (days saved per feedback cycle)

Target: From multi-day manual analysis to same-day insight delivery — a measurable, time-stamped reduction. This captures the platform’s core value proposition and correlates directly with retention, expansion revenue, and referrals. Measured via: timestamp comparison (survey close vs. first dashboard action), onboarding baseline survey, and quarterly business outcome review.

Full Metric Hierarchy

Tier	Metric	Unit	Target
North Star	Feedback-to-decision time reduction	Days saved/cycle	From days → hours
Primary	Feedback analysis time reduction	% vs. baseline	80% reduction
Primary	Screening filter accuracy	% precision	≥ 95%
Secondary	Monthly new subscriber businesses	Count	1,215/month (at scale)
Secondary	12-month subscriber retention	%	≥ 75%
Secondary	Survey completion rate	%	≥ 60%
Secondary	NPS (business users)	Score	≥ 50
Financial	Quarterly revenue	USD	\$180K/quarter
Financial	NPV by EOY 4	USD	\$520K
Health	AI analysis delivery time post-close	Hours	≤ 2 hours
Health	Adaptive question generation latency	Seconds	< 2s
Health	Platform uptime	%	≥ 99.5%

7. Risks & Mitigations

Risk	Probability	Impact	Mitigation
ML algorithms introduce bias or errors in theme clustering	Medium	High	Human review mode for flagged outputs; accuracy benchmarking against human analyst

Risk	Probability	Impact	Mitigation
Cloud infrastructure fails to scale under high response volumes	Medium	High	baseline; confidence score displayed on all AI insights Auto-scaling cloud architecture (AWS/GCP); load testing before each major launch; circuit breakers on AI analysis pipeline
Customers prefer traditional survey methods over adaptive AI	Medium	High (critical assumption)	A/B test adaptive vs. static in pilot; measure completion rate delta; if negative, make adaptive opt-in rather than default
Competitors (Qualtrics, Microsoft) launch similar AI features	High	Medium	Accelerate auto-improvement flywheel differentiation; deepen SMB go-to-market (Qualtrics is enterprise-focused); build community around research methodology
GDPR / privacy concerns from respondent data use for model training	Low	High	Explicit opt-in for model training data use; data anonymisation pipeline; legal review before EU launch
Failure to reach 2,900 subscribers/quarter for break-even	Medium	High	Conservative go-to-market via pilot referrals; agency partnership channel; track conversion funnel weekly; adjust pricing if conversion below 3%
LLM API (OpenAI/Anthropic) cost scaling with	High	Medium	Per-survey token budgeting; caching common adaptive


Risk	Probability	Impact	Mitigation
usage			question patterns; fallback to lighter model for low- complexity analysis

8. Critical Assumptions to Validate

From the MIT Design Thinking assumption mapping (prioritised by importance × uncertainty):

Assumption	Importance	Certainty	Validation Method
We can implement dynamic and adaptive survey functionality effectively	High	High	Technical spike in pilot phase
Cloud infrastructure handles current and future data loads	High	Medium	Load testing at 10× projected peak
ML algorithms transform feedback data without significant errors	High	Medium	A/B accuracy test vs. human analyst on 500-response gold standard set
Customers prefer automated and adaptive surveys over traditional	High	Low	A/B test in pilot: adaptive vs. static; measure completion rate and insight depth score
We can secure 2,900 active subscriptions/quarter to break even	High	Low	Pilot conversion funnel tracking; refine pricing and channel strategy by EOY2
Customers are interested in auto-generated surveys from previous feedback	Medium	Low	In-product feature adoption tracking during beta; survey PM on feature value

9. Timeline & Milestones

Milestone	Date	Deliverables	Success Criteria
MIT Design Thinking Validation	Q2 2026 	RWW analysis, competitive	Faculty and peer validation of “Real,

Milestone	Date	Deliverables	Success Criteria
		comparison, financial model	Win, Worth It"
Platform Pilot	Q3 2026	MVP deployed with 5 businesses; baseline vs. outcome measured	4/5 businesses report measurable analysis time reduction
Private Beta	Q4 2026	25–50 customers onboarded; pricing live; NPS measured	NPS \geq 40; \geq 60% completion rate; < 2hr analysis delivery
V1.0 Commercial Launch	Q1 2027	Full feature set GA; marketing campaign; agency partnerships	200 paying businesses by end of Q1 2027
Break-even	Q2 2027	2,900 active subscriptions/quarter reached	Monthly revenue covers operating costs
Enterprise Tier	Q3 2027	SSO, custom models, SLA, dedicated CSM	First enterprise contract signed
NPV Positive	EOY 2027	Financial milestone	Positive NPV demonstrated in board reporting

10. Team & Resources

Role	Person / Allocation
Product Management	Sankar Kumar Palaniappan
Business Development / GTM	Donal Phipps, Alaa Shahin
AI/ML Engineering	Anh-Tuan Bui
Solution Architecture	Ahmed Gamal Eldein A Shalaby
Design / UX	TBD (0.5 FTE contract)
QA / Testing	TBD (0.5 FTE contract)

Infrastructure Budget (Estimated):

Category	Monthly (Pilot)	Monthly (V1.0)
LLM API (adaptive questions + analysis)	~\$600	~\$4,000
Cloud hosting / infrastructure	~\$500	~\$3,000
Survey delivery infrastructure	~\$200	~\$800
Marketing / acquisition	~\$500	~\$5,000

Category	Monthly (Pilot)	Monthly (V1.0)
Total	~\$1,800/mo	~\$12,800/mo

Revenue target at V1.0: \$60,000/mo (200 subscribers × avg \$299). Margin positive at V1.0 launch.

11. Open Questions

- LLM provider selection:** Should adaptive question generation use OpenAI GPT-5, Anthropic Claude, or a self-hosted open-source model? Trade-offs: cost vs. quality vs. data privacy.
- Respondent panel sourcing:** Does Surve.AI build its own respondent panel, or integrate with panel providers (Lucid, Cint, Prolific)? Build vs. buy decision needed before pilot.
- Incentive legal compliance:** Amazon voucher referral chain — are there jurisdiction-specific legal constraints (EU, APAC) that could limit this mechanic?
- Screening IP protection:** Are the screening algorithm rules considered patentable / trade-secret-eligible, or do they need to be disclosed for enterprise procurement compliance?
- Model training ethics:** Can response data be used to improve the adaptive question model by default, or does this require explicit opt-in consent under GDPR / CCPA?

12. Assumptions Made

- Primary buyer persona is a product manager or CX director at a 50–5,000 employee company; individual researcher subscriptions are a secondary segment
- “Interactive voice survey” at MVP means IVR (interactive voice response) via phone, not video or avatar-based interaction
- Amazon voucher incentive mechanism is the primary recruitment driver; platform-managed panel is a Phase 2 feature
- Competitive differentiation from Qualtrics is on price and AI-native design, not on enterprise feature breadth
- \$1.3B market size figure and 12% CAGR are sourced from VerifiedMarketResearch.com (2024); should be refreshed with primary research before Series A fundraising
- 75% retention rate is an assumption based on benchmark SaaS metrics and the workflow lock-in created by survey history and auto-improvement; to be validated empirically in Year 1