

Loopback

Product Requirements Document · v3.0

Owner: Shiwal Singh | July 2026 | Public Showcase Version

Domain: fuffad.com (temporary) | Status: Active, Pre-Pilot

A note on why this is public: this is a working founder document, shared to show how I think through ambiguity, not a finished or audited business case. Every claim about users, pain points, or market size is labeled Hypothesis or Validated throughout. Where v3.0 differs from an earlier draft, most notably the market sizing in Section 3, I've kept the correction visible rather than quietly fixing it, because catching and correcting my own overreach is part of the process I want this document to show.

What changed from v2.0

Redundant framework layers were cut where a second framework added no new decision, for example Kano labels that duplicated MoSCoW. Every claim about users, pains, or market is now labeled as a hypothesis or as validated. Market sizing in Section 3 was rebuilt from public SaaS industry data and is explicitly labeled as a founder estimate rather than a researched figure. The document was trimmed from 39 pages toward a working, defensible version.

1. Executive Summary

Framework: Working Backwards (Amazon)

The Press Release

Loopback, for product teams who are tired of flying blind. Loopback is a new B2B SaaS product built in India for the US market. Its MVP is a zero-setup AI tool that reads a company's Freshdesk support tickets every week and delivers a three-point product intelligence digest to Slack and email automatically: no dashboards, no taxonomy, no analysts. Just the top pain point, top feature request, and one churn signal, every Monday morning before sprint planning. Loopback starts at \$49 per month with self-serve onboarding under 20 minutes.

The Product Bet

The problem (Hypothesis): B2B SaaS SMBs generally lack a system to extract product intelligence from support tickets. The data exists; no affordable, zero-setup tool synthesizes it at SMB price points.

The solution: Loopback connects to Freshdesk via API, runs AI synthesis weekly using Claude, and delivers a structured three-point digest via Slack and email.

The market (Hypothesis, see Section 3 caveat): B2B SaaS companies with 5 to 50 employees and 50 to 500 tickets per month, with no dedicated Voice of Customer tool. No competitor currently serves this exact segment at self-serve SMB pricing.

The moat: Support-team-first setup, opinionated output rather than a dashboard, \$49 to \$99 per month self-serve pricing, and connectors spanning both ticketing platforms and async communication channels.

The founder edge: 4.5 years inside B2B SaaS support and product operations at HighLevel, Maropost, and Vartopia, direct exposure to this exact problem across three different support stacks.

What Success Looks Like at 90 Days

Pilot Users	Paying Customers	MRR	North Star
3 real workspaces	1 on Stripe	\$49+	80% digest delivery rate

2. Problem Statement

Framework: 5 Whys

Surface problem (Hypothesis): B2B SaaS PMs at small companies often make roadmap decisions without structured input from support ticket data.

Why	Answer
Why don't PMs use ticket data for roadmap decisions?	Extracting themes from hundreds of unstructured tickets is too time-consuming to do consistently.
Why is it too time-consuming?	No tool automates it at SMB price points.
Why is there no affordable tool?	Existing tools (Enterpret, Productboard, Dovetail) are built for enterprise teams with \$36K to \$100K annual budgets.
Why don't SMBs use enterprise tools?	Cost, implementation complexity, and sales cycle don't fit SMB budgets or timelines.
Why has no product filled this gap?	Per-ticket AI synthesis only became cost-viable (under \$0.05/run) with recent LLM API pricing.

Root cause: the cost of extracting signal from ticket noise has historically exceeded the cost of ignoring it. LLM pricing has changed that equation. This is the timing bet the product is built on, labeled here as a bet, not a proven fact.

Failure Modes Loopback Addresses

Failure Mode	Problem	Loopback's Fix
Data Accessibility	Ticket data locked in Freshdesk/Zendesk, no export to PM workflow	Freshdesk API + CSV ingest, automatic
Synthesis Cost	Reading 200+ tickets/week takes hours of analyst time	Claude API synthesis in under 60 seconds, under \$0.05/run
Output Format	Dashboards require interpretation; PMs don't log in	Three-point digest to Slack/email, no login required
Pricing Access	VoC tools start at \$10K+/year with mandatory sales demos	Self-serve at \$49 to \$99/month, live in 20 minutes

Current Workarounds

Hypothesis, based on founder's own support-role observation, not yet interview-validated.

Workaround	Why It Fails
PM manually reads tickets	Unscalable above roughly 50 tickets/week; typically abandoned within 2 sprints
Support team writes summary	Inconsistent, recency-biased, often unread
Jira labels on escalated tickets	Captures only a fraction of signal; no qualitative synthesis
Enterprise tools (Enterpret, Productboard)	\$36K to \$100K/year, inaccessible to SMB
ChatGPT on exported CSV	No memory or automation; typically a one-off, abandoned habit

3. Market Opportunity

Framework: TAM / SAM / SOM

This section is a founder estimate, not a researched market sizing study. It was rebuilt using public SaaS industry statistics rather than the original v2.0 figures, which overstated the addressable market. Treat all numbers below as directional, not diligence-grade.

Public estimates of the global SaaS company count vary by source, from roughly 30,800 to 42,000 companies globally as of 2026, because "SaaS company" is defined inconsistently across sources. The U.S. alone accounts for roughly 17,000 of these. There is no single authoritative count of B2B SaaS companies with 5 to 50 employees; this segment has to be estimated top-down with wide error bars.

Level	Definition	Estimate	Reasoning (unvalidated)
TAM	All SaaS companies globally	~35,000 companies	Midpoint of published 2026 estimates (30,800 to 42,000). Assumes ~75% are B2B-oriented (~26,000 companies) with average willingness-to-pay of \$600 to \$1,200/year for a VoC tool if one existed, giving a rough TAM of \$16M to \$31M/year, not the \$4.2B originally claimed in v2.0.
SAM	B2B SaaS SMBs (5 to 200 employees) in US/Canada/Europe with 50 to 500 tickets/month and no VoC tool	Unknown, likely low single-digit thousands of qualifying companies	No public data source segments SaaS companies by ticket volume or VoC tool ownership. Must be validated through pilot outreach, not assumed.
SOM (12 mo)	US/Canada-first SMBs reachable via PLG, Freshdesk/Zendesk communities, content marketing	~\$500K ARR (unchanged from v2.0)	~1,000 paying customers at \$49/month average. Now understood as an ambitious percentage of a much smaller TAM than originally estimated, not a conservative slice of a large one.

What this means for the plan: the corrected TAM is dramatically smaller than v2.0 claimed. This doesn't kill the thesis, a \$16 to 31M/year addressable market can still support a meaningful solo-founder SaaS business at the SOM target, but the original "massive market" framing was not defensible and should not be used in investor conversations without a proper bottom-up study, for example actual Freshdesk/Zendesk customer counts by company size. Flagged as Open Question Q-09.

Why Now, the Timing Argument

This market gap likely existed for years before AI made synthesis affordable. Manually reading 500 tickets could cost roughly 8 hours of analyst time (about \$400 at \$50/hr); Claude API now does comparable synthesis in under a minute for about \$0.05 (Hypothesis). This cost inflection is what makes a self-serve SMB VoC product viable now, where it wasn't before.

4. Target Users and ICP

Framework: Persona Framework, all content below is Hypothesis until validated through pilot interviews.

Dimension	Qualifier	Disqualifier
Company type	B2B SaaS	B2C, marketplace, services, agency
Employees	5 to 50	1 to 4 (no PM role yet) or 200+ (enterprise budget)
Support volume	50 to 500 tickets/month	<50 (not enough signal) or 5,000+ (needs enterprise)
PM presence	At least 1 PM or Founder-as-PM	No dedicated product role
VoC tool	None in use	Already paying for Dovetail/Productboard/Interpret
Geography	US and Canada (Phase 1)	APAC or LatAm (Phase 3+)

Persona 1, the Founder-PM (MVP focus)

Founder or Head of Product at a 5 to 20 person B2B SaaS company (Hypothesis). Wears multiple hats; knows support tickets contain signal but rarely has time to read them consistently; has tried CSV-export-to-ChatGPT and abandoned it; evaluates tools by time-to-first-value. Definition of success: receives the Monday digest before sprint planning, and makes at least one roadmap decision per month citing a Loopback insight.

Persona 2, the Support Lead

Head of Support or CS Lead at a 20 to 50 person B2B SaaS company (Hypothesis). Currently writes manual summaries that go unread; would automate this if setup is trivial; sets up Loopback on the PM's behalf; will abandon if setup requires taxonomy work.

Persona 3, the PM at Scale (explicitly not MVP focus)

Product Manager at a 50 to 200 person B2B SaaS company. Has a slow, manual VoC process; larger budget tolerance (\$99 to \$299/month); every V1 decision optimizes for Persona 1 instead.

5. Jobs To Be Done

Framework: JTBD (Christensen), Hypothesis pending pilot validation.

ID	Job	Statement
FJ-01	Sprint Planning	When preparing for Monday sprint planning, I want the top 3 customer pain points from this week's tickets, so I can prioritize without spending hours reading tickets.
FJ-02	Churn Prevention	When a new spike in a ticket type appears, I want to be alerted, so I can investigate before it becomes a churn event.
FJ-03	Roadmap Justification	When presenting the roadmap to leadership, I want real customer evidence, so I'm not relying on anecdote.
FJ-04	Support Team Alignment	When my support lead asks what we're building next, I want to show their data influenced the roadmap.
EJ-01	Confidence (emotional)	I want to feel confident I haven't missed anything important in this week's support data.
SJ-01	Credibility (social)	I want to demonstrate my roadmap is grounded in systematic evidence, not gut feel.

Note: EJ-01 (Confidence) is treated as the primary retention driver in this document's design decisions. This is a founder hypothesis about what will drive renewal, not yet observed in real usage data.

6. Value Proposition

Framework: Value Proposition Canvas

One-sentence value proposition: Loopback gives B2B SaaS PMs the confidence that they are not missing anything important in their support data, delivered automatically every Monday morning, for the cost of a lunch.

Customer Pains (Hypothesis)	Pain Relievers
No time to read 200+ tickets/week	AI synthesis in ~60 seconds
VoC tools are enterprise-priced and over-engineered	10 to 100x cheaper than incumbents
Manual summaries are inconsistent	One consistent, automated source
Roadmap decisions made with incomplete data	Evidence-backed digest by default

7. Competitive Landscape

Framework: Positioning Matrix

Product	Target Segment	Pricing	Self-Serve SMB?
Enterpret	Enterprise	\$36,000+/yr	No
Productboard	Mid-market to Enterprise	\$118/mo floor; \$70 to 100K/yr enterprise	Limited
Dovetail	Enterprise/UX teams	Custom	No
UserVoice	SMB to Mid-market	\$899/mo Growth plan	Limited
Canny	Early-stage SMB	Free to \$400/mo	Yes (voting portal, not synthesis)
Loopback	5 to 50 person B2B SaaS	\$49 to \$99/mo self-serve	Yes

Positioning conclusion: every competitor with serious AI synthesis capability sits in the expensive and complex quadrant. The only affordable self-serve option (Canny) is a voting portal, not an intelligence tool. The affordable, zero-setup, AI-synthesis quadrant appears open. This has not been independently verified beyond the products listed above and should be re-checked periodically, since this space moves fast.

Moat, condensed to the two most defensible layers

1. Support-team-first setup, PM-only consumption, inverting the typical enterprise onboarding model where a PM or analyst must configure the tool.
2. Cross-platform ingestion spanning both ticketing and async channels, a genuine data moat if built, since no listed competitor currently spans both ticketing and Slack/Teams/WhatsApp at SMB pricing. This is a V2+ commitment, not yet built, and should be tracked as a roadmap dependency, not a current differentiator.

8. Product Strategy, Strategic Pillars

Pillar	Principle
P1	Time to Value is the Product. First value moment must happen within 20 minutes of signup.
P2	Opinionated Output Beats Dashboards. Every output is a decision, not raw data requiring interpretation.
P3	Support Lead Setup, PM Consumption. PM never needs to log in to get value.
P4	No Hallucinations, Ever. Every quote must be traceable to a real ticket. First 5 digests per pilot manually reviewed.
P5	SMB Pricing is a Strategic Commitment. Do not move upmarket until SMB is saturated.

9. MVP Scope

Framework: MoSCoW

MVP definition: a Freshdesk-connected weekly digest (Slack + email) surfacing the top pain point, top feature request, and one churn signal, self-serve setup under 20 minutes, zero taxonomy.

Constraint	Decision	Reason
Ticket sources	Freshdesk API + CSV fallback	Freshdesk is the native domain; CSV removes onboarding blockers for pilots
Output channels	Slack webhook + Resend email	Delivers value without requiring login; dashboard deferred to V1.1
AI model	Claude API (claude-sonnet-4-6)	Best-in-class synthesis, structured JSON output, under \$0.05/run at 200 tickets
Onboarding	Self-serve, under 20 min, no sales touch	Core product principle
Pricing	Free (3 pilots) then \$49 to \$99/mo via Stripe	Free validates product; paid validates willingness to pay

Build Phases

Phase	Weeks	Theme	Definition of Done
0	Week 1	AI validation	Claude correctly identifies top 3 themes from 20 real tickets, manually verified
1	Weeks 2 to 3	Core pipeline	Freshdesk CSV ingest + AI synthesis produces a readable digest
2	Week 4	Delivery	Digest delivered on schedule to Slack + email
3	Weeks 5 to 6	Onboarding	New user connects Freshdesk, receives first digest in under 20 minutes
4	Week 7	Launch	Landing page live, waitlist open, Stripe configured, first real pilot digest sent

10. Feature Requirements

Framework: MoSCoW + RICE. Kano labels dropped from v2.0, in the Must-Have tier, Kano's Basic/Performance/Delighter classification didn't change any prioritization decision that RICE hadn't already made. Kept only where it adds real signal, distinguishing Should/Could-tier delighters worth pursuing.

Must Have

Feature	RICE	User Story
F-01 Freshdesk CSV Ingest	40.5	Upload CSV export, no API setup required
F-02 Freshdesk API Connection	22.5	Connect via API key, auto-sync every 24h
F-03 AI Synthesis Engine	15.0	Identify top pain point, feature request, churn signal weekly
F-04 Weekly Slack Digest	42.75	Delivered Monday 8 to 9 AM user timezone
F-05 Weekly Email Digest	34.2	Sent via Resend, mobile-readable, unsubscribe link
F-06 Signup and Onboarding	22.5	Five-screen flow, first digest triggers immediately on completion
F-07 Stripe Billing	14.4	Self-serve upgrade, no sales touch

Should Have

Feature	RICE	Kano
F-08 Digest History Dashboard	7.2	Performance
F-09 Ticket Volume Trend Chart	6.3	Delighter
F-10 Salesforce Case Ingest	4.67	Performance

Could Have

Feature	RICE	Kano
F-11 Jira Issue from Insight	7.0	Delighter
F-12 Zendesk + Intercom Ingest	3.5	Performance

11. Success Metrics

Framework: OKR + North Star Metric

North Star: Weekly Digest Delivery Rate, the percent of connected workspaces that receive and open their weekly digest. Target: 80% by Week 8. Chosen because delivery alone isn't value; a digest that's delivered but not opened means the product isn't yet trusted enough to check.

Key Result	Target	Measurement
KR1: Activate 3 real pilot workspaces	Week 6	Supabase onboarding event log
KR2: 80% Weekly Digest Delivery Rate	Week 8	Resend open rate + Slack click data
KR3: 1 paying customer on Stripe	Week 10	Stripe MRR dashboard
KR4: 1 pilot PM cites a Loopback insight in sprint planning	Month 2	User interview / Slack confirmation
KR5: Week-4 retention above 50%	Week 4	% of activated workspaces still receiving digest

12. Technical Architecture

Layer	Technology	Free Tier?
Frontend	Next.js on Vercel	Yes
Backend	FastAPI on Render	Yes
Database + Auth	Supabase (PostgreSQL)	Yes (500MB)
AI Synthesis	Claude API (claude-sonnet-4-6)	Pay per token (\$0.003/1K)
Ticket Ingestion	Freshdesk REST API + CSV parser (MVP)	Freshdesk free tier
Email Delivery	Resend	Yes (3K/mo)
Slack Delivery	Slack Incoming Webhooks	Free
Billing	Stripe Checkout	No (2.9% + 30 cents)

Data Flow

1. Ingest: Freshdesk API (24h incremental sync) or CSV upload, stored in a Supabase tickets table. 2. Pre-process: filter to the current 7-day window, dedupe, remove closed/spam tickets. 3. AI Synthesis: Claude API returns structured JSON only, with pain point, feature request, and churn signal, each with theme, ticket count, severity, and exactly 2 supporting quotes. 4. Digest Builder: JSON validated against schema, assembled into a three-section digest, stored in a digests table. 5. Delivery: Slack webhook + Resend email, logged, retry once on failure. 6. Storage: last 8 weeks retained for a future dashboard; raw ticket data purged after 90 days.

AI Prompt Design, Non-Negotiable Rules

Output is always structured JSON, never open-ended prose. The system prompt explicitly instructs the model to only reference content present in the provided tickets and not extrapolate. The first 5 digests per pilot workspace are manually reviewed before delivery. Token cost target is under \$0.05 per weekly run at 200 tickets.

13. Risks and Assumptions

Framework: Risk Matrix, trimmed to the risks with the highest likelihood times impact.

Risk	Score	Mitigation
AI synthesis produces generic or hallucinated insights	High	Structured JSON schema; manual review of first 5 digests per workspace
Pilot users won't share real ticket data	High	CSV upload removes API dependency; any exportable platform qualifies
Users activate but don't return after first digest	High	Immediate first digest at onboarding, not next Monday; Day 3 follow-up
US buyers won't share customer data with an unknown solo founder	High	SOC 2 readiness signaling on landing page; DPA template ready pre-launch
No paying customers after 3 pilots	High	Charge at least \$1/month from Day 1 to establish payment behavior
Claude API cost exceeds viable unit economics	Medium	Modeled at ~\$0.05/200 tickets; monitor per-workspace cost weekly

Assumption Log, top 5 by risk-if-wrong

ID	Assumption	Validation Method	Deadline
A-01	SMB PMs will connect a real support tool from an unknown solo founder	3 pilots complete onboarding with real tickets	Week 6
A-02	A three-point weekly digest is the right output format	User interview after first digest	After pilot digest 1
A-04	Support leads will set up Loopback without hand-holding	Observe onboarding completion with zero founder support	Week 3
A-05	\$49 to \$99/month is acceptable for the value delivered	Charge at least one pilot user before Month 2	Month 2
A-06	20 tickets is enough to generate meaningful insights	Test AI synthesis on 20 real tickets, validate manually	Week 1

14. Out of Scope (MVP)

Feature	Category	Reason
In-app dashboard	V1.1	Slack + email covers the core job first
Zendesk / Intercom / Salesforce integration	V2	Validate on Freshdesk before adding sources
Multi-workspace accounts	V1.1	Single workspace simplifies auth/billing for MVP
Jira issue creation	V2	High value, not core to the initial job
Public feedback portal / roadmap tool / NPS survey	Never	Different product category entirely
Mobile app	V3	Web-first; digest already covers mobile via Slack/email
Custom AI taxonomy	V3	Complexity undermines the SMB self-serve positioning

15. Roadmap

Phase	Timing	Learning Question
MVP	Weeks 1 to 7	Does AI synthesis of support tickets produce insights PMs find useful and act on?
V1.1	Weeks 8 to 12	Do users come back after the first digest, and why?
V2	Months 3 to 5	Can Loopback ingest from multiple sources and still deliver a coherent digest?
V2.1	Months 5 to 6	Can Loopback ingest from async communication channels, not just ticketing?
V3	Months 6 to 12	Can Loopback evolve from a weekly digest to a continuous intelligence layer?

16. Decision Log

Kept to decisions with real trade-offs, not routine choices.

ID	Decision	Trade-off
DR01	Freshdesk first, not Zendesk/Intercom	Founder has hands-on Freshdesk experience; other platforms follow in V2
DR02	Three-point digest, not a dashboard	Opinionated output is the positioning; configurability would dilute it
DR03	Slack + email only in MVP, no in-app UI	Saves 2 to 3 weeks of frontend build; core job fully served without it
DR04	\$49 to \$99/month, not freemium-unlimited	Freemium with no conversion mechanism won't validate willingness to pay
DR07	HighLevel native integration deferred indefinitely	Proceeding without vendor due diligence and TOS review would create legal/platform risk; replaced with a platform-agnostic strategy (Slack/Teams/WhatsApp) that avoids single-vendor dependency

17. Open Questions

ID	Question	Deadline
Q-01	Should the digest be daily, weekly, or configurable in V1?	After pilot 1 feedback
Q-02	Charge pilot users from Day 1, or free for 30 days?	Before launch
Q-03	What's the minimum viable ticket volume for a useful digest?	Week 1
Q-05	Do we need SOC 2 signaling before US buyers share ticket data?	Month 2
Q-06	What is the final brand name before US GTM? (fuffad.com is temporary)	Pre-launch (Week 7)
Q-09	What is the real, bottom-up SAM for this product? Section 3's market sizing is a founder estimate built from public SaaS industry stats, not a validated study, and needs actual data before appearing in any investor-facing material.	Before any fundraising conversation