PM Career Navigation Platform

PMpath by Piyush Funde

PRODUCT REQUIREMENTS DOCUMENT

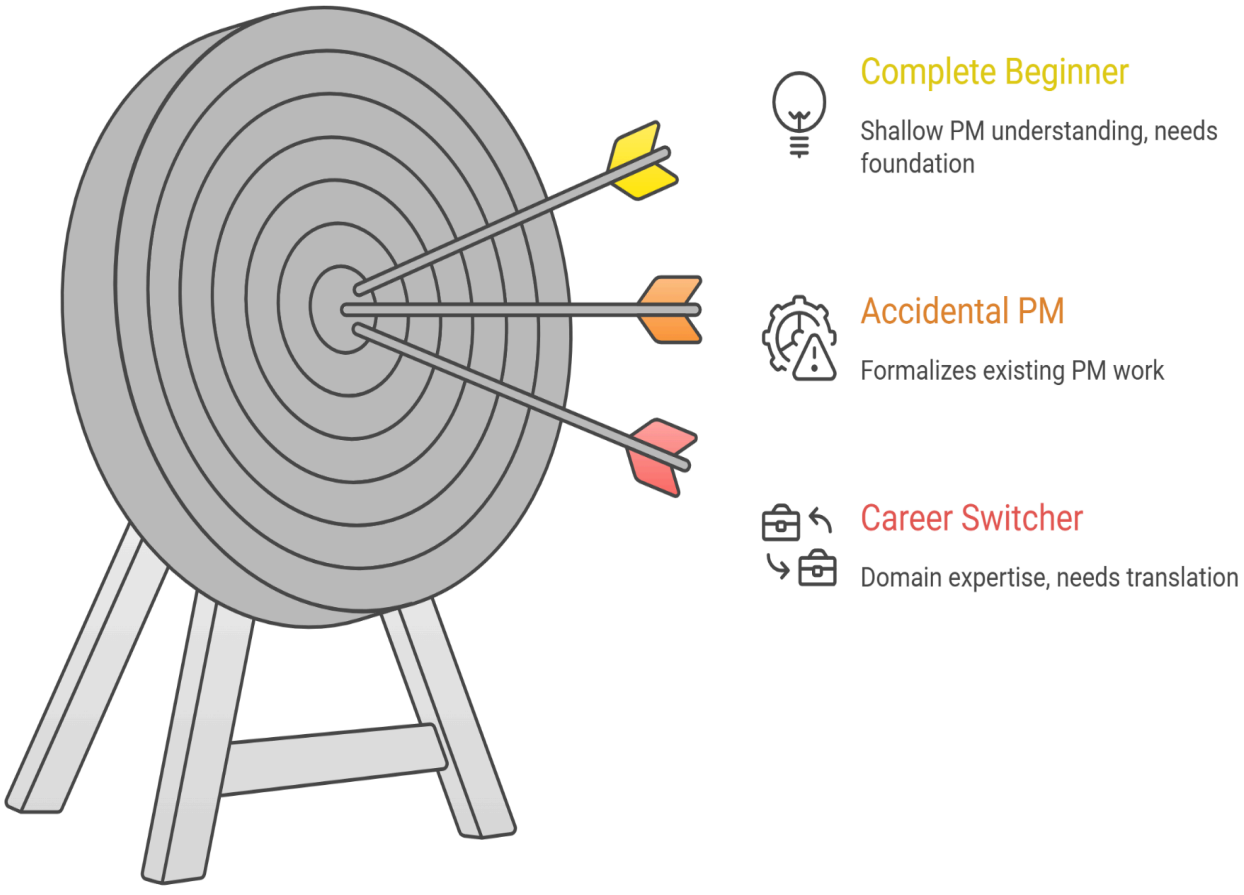

Link: pmpath.piyushfunde.com

1. Discovery Insights: Who Is the User and What Did We Learn?

1.1. Research Approach

Discovery was conducted in two phases. The first phase was secondary research across 40+ sources covering the Indian PM hiring landscape, competitor platforms, candidate journey data, salary benchmarks, and community sentiment. The second phase was 5 primary user interviews with PM aspirants.

1.2. Who the user is

Three distinct user types emerged from the research. They are meaningfully different in what they need and should not be treated as a single persona.

The first type is the career switcher with domain expertise but no PM title. This person has 5 to 10 years of experience in consulting, operations, finance, or a domain-heavy role. They have done product-adjacent work such as requirements gathering, sprint facilitation, stakeholder management, and impact measurement. They cannot claim it because the vocabulary does not match PM job descriptions and their title filters them out of ATS before any human reads their resume. Their core problem is translation and access, not knowledge.

The second type is the accidental PM or developer who needs to formalize. This person is already doing PM-like work, either as a developer, an APM at an early stage startup, or a technical lead who has drifted into product decisions. They do not recognise their existing work as PM experience because it does not look like what they read about in PM content. Their core problem is identity clarity and structured frameworks to articulate what they already do.

The third type is the complete beginner with shallow PM understanding. This person has decided to pursue PM based on one or two external signals such as a mentor suggestion, LinkedIn content, or hype around the role. They have not spoken to PMs, have not studied JDs in depth, and have a surface-level understanding of what the job involves. Their core problem is reality check and a structured foundation before anything else.

1.3. What the research revealed

Six patterns emerged consistently across both primary and secondary research.

The first and most important pattern is that the real problem is not lack of knowledge. It is a lack of calibrated self-knowledge. Every user in the interviews was making strategic decisions based on an incomplete or incorrect understanding of where they stood relative to what employers actually want. One user applied 50+ times and received no callbacks. Another applied 20 to 30 times, got one callback, failed the interview, and stopped applying entirely. A third has been in learning mode for a year without applying once. In all three cases the problem was not skill deficiency. It was the absence of an honest external signal telling them where they stood and what to do next.

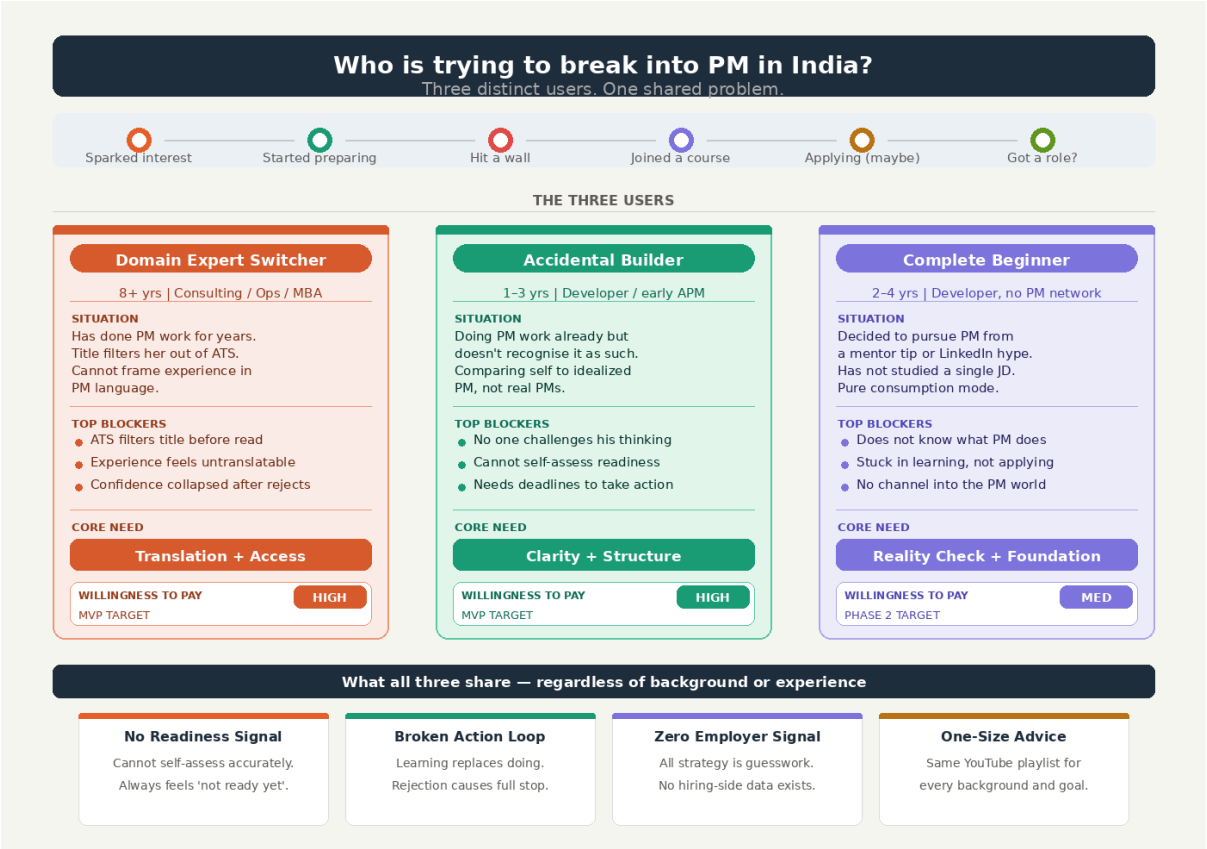

The second pattern is that confidence is broken before the skill gap becomes relevant. At least three of the five users interviewed had stopped taking meaningful action entirely, not because they lacked ability but because rejection or the absence of feedback had collapsed their belief that the path was open to them at all. This is not a motivation problem. It is a structural problem that no amount of content solves.

The third pattern is that preparation defaults to passive consumption without forcing functions. Every user described some version of watching videos, reading blogs, and doing courses as a substitute for applying. The course they had enrolled in was being used not primarily as a skills upgrade but as a socially acceptable reason to delay the risk of rejection. Any product that offers self-paced learning without built-in accountability will replicate this exact behavior.

The fourth pattern is that existing experience almost always translates to PM but users cannot see it themselves. Every user interviewed had done work that maps to PM competencies. None of them could frame it in PM language without help. The gap is articulation, not experience.

The fifth pattern is that the one-size-fits-all advice problem is not about scattered resources. It is that wildly different people receive identical guidance. A developer with 3 years of experience building 8 products and a Deloitte consultant with an MBA and 8 years of stakeholder management are given the same YouTube playlist and the same resume template. The starting points are so different that generic advice is actively harmful for both of them.

The sixth pattern is that nobody has employer-side information. All five users were making their most important strategic decisions based on zero direct signal from hiring managers. One user inferred his skill gap from the silence after a VP interview without ever being told what went wrong. Another got told her resume looked JD-heavy but had no idea how to fix it. The absence of employer-side feedback is not just frustrating. It makes the entire preparation loop directionally wrong.

2. Problem Prioritisation

The discovery process involved two rounds of filtering. The first round used **secondary research across 40+ sources** covering the Indian PM hiring landscape, competitor platforms, candidate journey data, and community sentiment to generate an initial long list of pain points. The second round validated and weighted these through **5 primary user interviews** with aspirants at different stages of the PM transition journey. Pain points that could not be validated by at least one of the **two sources were discarded**.

The remaining **9 pain points** were scored on two dimensions: frequency, meaning how many aspirants face this problem, and impact, meaning how severely it blocks progress toward a PM role. Each dimension was rated on a scale of 1 to 5, producing a combined score out of 25.

Three pain points scored 25 out of 25 and were classified as **P0 Critica**l.

The first is the **absence of a readiness signal**. Almost every aspirant in both primary and secondary research was stuck in an indefinite preparation loop not because they lacked content but because they had no way to know when they were ready to apply. Without an external benchmark, the default human behavior is to keep consuming more material because learning feels like progress without requiring the risk of rejection.

The second is **confidence-action loop breakdown**. Rejections, or even the absence of callbacks, triggered a complete stop in application activity for multiple users. This is not a motivation problem. It is a structural problem. There is no mechanism that keeps people moving after a setback, and no small win to anchor confidence before the next attempt.

The third is **ATS and title-based screening**. Career switchers, who make up a significant share of PM aspirants, are being eliminated before any human reviews their application. The problem is not their experience or their ability. It is that their title does not pattern-match to what automated filters are looking for, and nobody tells them this until they have already wasted 20 to 30 applications.

These three problems were selected as the primary focus because they share three properties. They are **universal**, meaning they affect aspirants regardless of background or experience level. They are **blockers**, not friction points, meaning they stop the journey entirely rather than just slow it down. And they are **upstream**, meaning solving them unblocks the rest of the preparation journey rather than optimising a step within it.

The remaining six pain points scored between 12 and 20 and are real problems worth solving, but they either affect a narrower segment, have partial workarounds already available in the market, or become less acute once the **three P0 problems are addressed**. They will be considered for **subsequent product phases**.

Here is the link to the sheet for a detailed overview of the prioritisation. click here

3. Proposed Solution: What Are We Building and Why This Way?

3.1. The one-line version

A personalised PM readiness platform that tells you exactly where you stand, builds a learning path specific to your background, and prepares your resume for the roles you are actually qualified for.

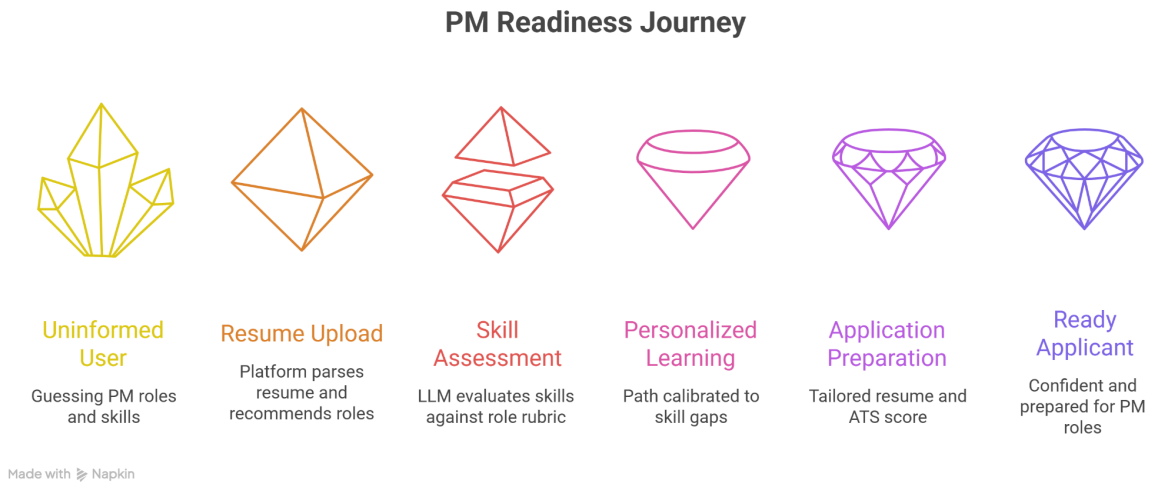

3.2. What the MVP does

The product is an 8-step gated journey. The gating is the product. It replaces the endless consumption loop with a path that ends in an application.

The user uploads their resume. The platform parses it and recommends which PM role fits their background using plain language explanations and real examples, so a consultant or developer can make an informed choice rather than guessing. They then answer 8 to 12 situational multiple choice questions or voice and text input for subjective questions, which the LLM evaluates against a role-specific rubric to produce a skill gap map and a starting readiness score. This is the first honest external signal most users have ever received about where they stand relative to what employers actually want.

A personalised learning path follows, with depth calibrated to the gap. Weak areas get extended modules, harder quiz thresholds, and a required project submission such as a case study, a PRD draft, or a prototype. Strong areas get a refresher and move faster. Every stage is gated by proof of completion. A streak counter maintains daily momentum throughout, and a rejection recovery prompt fires automatically whenever a user logs a rejection so the platform keeps them moving rather than letting them stop.

When the readiness score crosses 75 the platform nudges the user to start applying. They can paste any JD and receive a tailored resume built from their onboarding data, uploaded resume, completed learning path, and JD keywords, with a before and after ATS score showing exactly what changed and why. Users who have been filtered out by their title get a specific reframing suggestion for that application. At full completion, a shareable certificate with a score breakdown is issued at a public URL, giving users a credibility signal to attach to cold outreach and referral requests rather than just a generic ask for help.

3.3. What makes this different

Every existing platform sells content, community, or mentorship. None start by diagnosing where the specific user actually is. None tell a consultant that her experience already maps to B2B PM and show her how to frame it. None address ATS screening as a structural fix. None give the user a signal for when to stop preparing and start applying. The closest competitors, Upraised at 1.75 lakhs and Exponent at roughly 1000 rupees per month, do none of these things.

3.4. Three principles the MVP is built on

Calibration over content: tell users where they stand specifically, not just what to study next. Behavior over knowledge: gating, streaks, nudges, and rejection prompts are design decisions not features. Translation over training: most users already have relevant experience, the platform helps them frame it in PM language rather than asking them to start from scratch.

3.5. Metrics and Measurement

North Star Number of users who receive at least one PM interview callback after completing 50 percent or more of their learning path.

Leading Metrics (measurable in first 30 days)

Onboarding completion rate: percentage who finish all four onboarding steps including proficiency assessment. Target 65 percent plus.

Day-7 retention: percentage who return and complete at least one task within a week of onboarding. Target 45 percent plus.

Stage completion rate: average stages completed before going inactive. Target 4 or more.

Resume builder activation: percentage of users above 60 percent readiness who use the resume builder. Target 55 percent plus.

Lagging Metrics (confirm real outcomes)

Interview callback rate: percentage who report at least one callback after using the resume builder. Target 20 percent plus for users with 6 or more stages complete.

Offer rate: percentage who report a PM job offer within 6 months. Target 8 percent plus in the first cohort.

Certificate share rate: percentage of certificate earners who share their link. Target 40 percent plus.

NPS: collected after first certificate earned. Target 45 plus.

4. Implementation Plan: How We Are Building This

4.1. Tech stack and approach

The platform is built on Next.js 14 with App Router, Prisma for the database layer, Supabase for storage, and NextAuth for Google OAuth. All AI features run on Claude via the Anthropic API, with OpenAI as a fallback. Development happens in VS Code using Claude Code for agentic implementation. The architecture is designed around 7 agentic workflows: Resume Parser, PSI Reframer, Gap Analyzer, Conversation Agent with Proficiency Evaluator, Learning Path Generator, Stage Evaluator, and Resume Builder Agent.

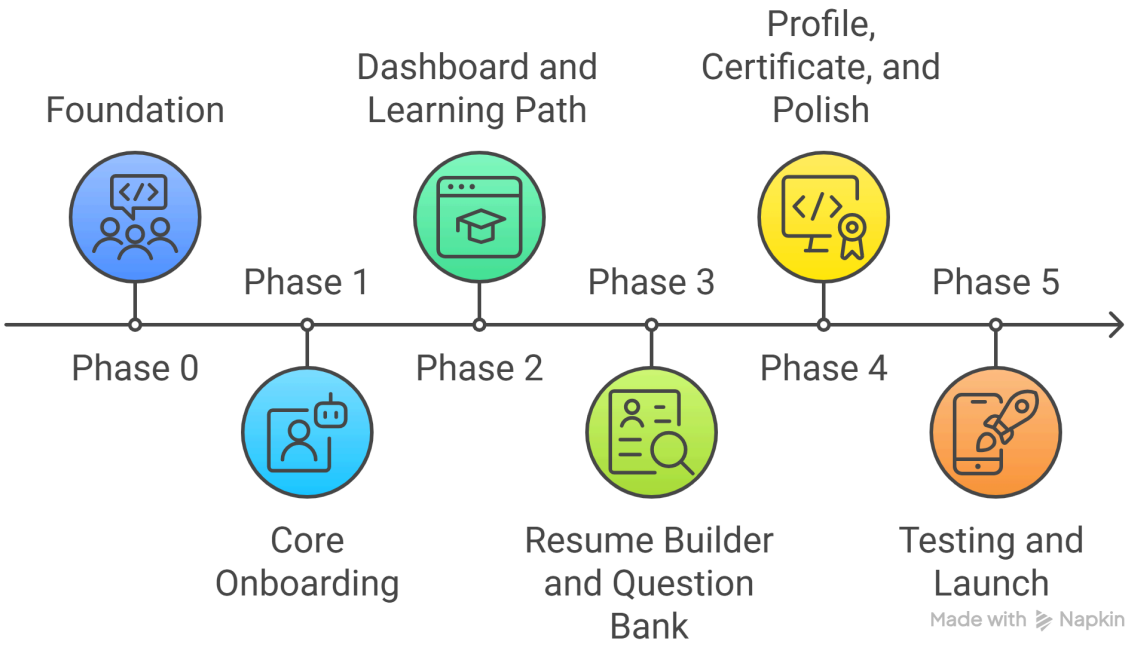

MVP Timeline

4.2. Phase 0: Foundation

Project setup, database schema, Google OAuth, seeding of the skill taxonomy, learning stages, and initial question bank. This phase also sets up the AI service abstraction layer that all agents sit on top of.

4.3. Phase 1: Core Onboarding

Resume upload and parsing, PSI reframing of work experience, gap analysis against the skills taxonomy, the conversational drill-down with voice and text input, proficiency assessment questions, and the full analysis dashboard with readiness score and skill gap visualization.

4.4. Phase 2: Dashboard and Learning Path

Main dashboard with readiness score widget and sub-scores, the full learning path with gated stages, resource completion tracking, quiz and assignment submission with AI evaluation, stage unlock logic, streak system, and activity logging.

4.5. Phase 3: Resume Builder and Question Bank

JD storage, resume generation pipeline, TipTap editor, before and after ATS scoring, PDF and DOCX download, resume version management, question bank listing, AI-evaluated practice, and attempt history.

4.6. Phase 4: Profile, Certificate, and Polish

Public shareable profile page, certificate milestones at Stage 1, 50 percent, and full completion, activity heatmap, verification badge, discovery module quiz, landing page, edge case handling, and performance optimization.

4.7. Phase 5: Testing and Launch

End to end flow testing, AI output quality review across 20 sample onboardings, mobile responsiveness, SEO for public profile pages, analytics via Vercel, error monitoring via Sentry, and first user cohort onboarding.

4.8. Critical path

Schema and auth must complete before resume upload. Skill taxonomy must be seeded before gap analysis can run. Readiness score depends on PSI entries and proficiency assessment. Resume builder depends on readiness score and PSI data. Public profile depends on PSI entries. Content creation for learning stages, question bank, and role descriptions runs in parallel to engineering throughout all phases.

5. Instruction Design: How the Platform Is Built and How It Works

5.1. Overview

The platform is built across 5 phases. Each phase delivers a self-contained set of features that build on the previous one. The development environment is VS Code with Claude Code handling implementation. All AI features run on the Anthropic API through a shared service abstraction layer so every agent in the system calls through the same interface. The database runs on Postgres via Prisma with Supabase handling file storage.

5.2. Phase 0: Foundation

The foundation phase establishes everything the product depends on before any user-facing feature is built. This includes the database schema, authentication, and the initial data that powers personalisation throughout the platform.

Google OAuth is the primary authentication method. The database schema covers all core entities: users, resumes, PSI entries, skill scores, learning stages, stage progress, question attempts, generated resumes, streaks, and applications. These tables are created via Prisma migrations and remain stable throughout the build.

Three data sets are seeded into the database before any feature is built. The first is the skill taxonomy: 8 skill categories covering product thinking, analytical skills, user research, technical acumen, communication, execution, business acumen, and leadership, with approximately 35 specific skills across them. The second is the role weight distribution, which defines how much each skill category contributes to the readiness score for each PM role type. Consumer PM, Growth PM, Technical PM, Platform PM, AI PM, and General PM each have a different weight profile stored in the database and configurable without code changes. The third is the learning path: 12 stages with sub-topics, curated resource lists, quick check prompts, and gate assignment rubrics including passing thresholds. An initial question bank of 50 entries is also seeded at this stage.

5.3. Phase 1: Core Onboarding

The onboarding phase is the most critical part of the product. The quality of every downstream feature, the skill gap map, the learning path depth, the readiness score, and the resume output, depends entirely on what is extracted here.

The onboarding flow has four sequential steps.

The **first step** is resume upload and parsing. The user uploads their resume in PDF or Word format. The Resume Parser agent reads the raw text and returns a structured profile containing the user's current title, company, years of experience, domain expertise, technical skills, and evidence of PM-adjacent work. The agent is designed to surface implicit PM signals that the user may not have labelled as such, including requirements gathering, sprint facilitation, stakeholder management, and user research. This is important because most career switchers have done PM work without recognising it.

The **second step** is PSI reframing. The PSI Reframer agent takes every work experience extracted from the resume and rewrites it in Problem, Solution, Impact format. A bullet that reads "built a configuration tool" becomes "identified that the operations team was spending four hours per day on manual setup, designed a self-serve configuration tool, reducing that time to twenty minutes." Every reframed entry is stored in the database with the original text, the reframed version, and the skill IDs from the taxonomy that the entry provides evidence for. These PSI entries become the source material for the resume builder later in the flow.

The **third step** is the gap-filling and proficiency assessment conversation. The user answers 8 to 12 questions through a conversational interface that supports both text input and voice input via the Web Speech API. The Conversation Agent runs in two modes sequentially. In gap-filling mode it asks questions designed to surface PM-relevant experience not captured in the resume, such as times the user identified a problem unprompted, influenced a decision without formal authority, or measured the impact of their own work. In proficiency assessment mode it asks role-specific situational questions calibrated to the user's chosen PM type. One example question for a Technical PM target is: "A senior engineer tells you your feature will take three months. The planning meeting said three weeks. What do you

do?" Each answer is evaluated by the Proficiency Evaluator sub-agent against a rubric that assesses specificity, structure, PM vocabulary, and domain relevance. Every skill category receives a proficiency tag of Strong, Developing, or Needs Work based on the combined evidence.

The **fourth step** is the gap analysis output. The Gap Analyzer agent maps all collected data against the skill taxonomy weighted for the user's target role and calculates the initial readiness score. The score uses a 50/30/20 weighting: 50 percent from stage completion, 30 percent from quiz performance, and 20 percent from project quality. At this stage, since no stages have been completed, the score reflects only the proficiency assessment data. The user sees a skill gap screen showing each category as Strong in green, Developing in amber, or Needs Work in red, each with a one-line explanation. Their starting readiness score is displayed here for the first time.

5.4. Phase 2: Learning Path

The learning path consists of 12 sequential stages covering all 8 skill categories. Stages are gated: a user cannot access Stage 5 without completing Stage 4. This gating is a deliberate behavioral design decision that prevents the platform from becoming a passive content library.

Each stage contains sub-topics with curated articles and videos, short comprehension checks called quick checks, and a gate assignment at the end. The gate assignment is a practical task: writing a PRD, designing an experiment, analysing a funnel, or building a prototype brief. The Stage Evaluator agent scores each submission against the rubric stored in the database, returns a score out of 100, and either unlocks the next stage or returns specific feedback pointing to which sub-topics the user needs to revisit before retrying.

The depth of each stage is personalised based on the proficiency tags from onboarding. A skill category tagged Strong receives one refresher article and a 3-question check with a lower passing threshold. A category tagged Developing receives 2 to 3 articles and a standard assignment. A category tagged Needs Work receives the full resource set, a harder quiz threshold, and a required full project submission.

The streak system tracks consecutive days with at least one qualifying action on the platform. Qualifying actions include marking a resource as read, attempting a quiz, submitting an assignment, and logging a job application. The streak count is visible on the dashboard at all times. Milestones at 7, 14, and 30 days trigger a brief celebration moment and a badge stored on the user's profile. If a streak breaks, a single recovery nudge is surfaced the following day with one specific suggested action rather than a generic prompt to return.

The readiness score updates every time a stage is completed. When the score crosses 70 the platform surfaces a dedicated screen informing the user they are ready to begin applying, showing 5 role suggestions matched to their profile. The learning path remains accessible in parallel with the application phase.

5.5. Phase 3: Resume Builder

The resume builder generates a tailored resume for each job the user wants to apply to. The user pastes a job description or a link to one. The Resume Builder agent constructs a personalised resume by drawing from four data sources in priority order: the PSI entries from onboarding, the original uploaded resume, the user's learning path completion record, and the keywords extracted from the job description.

The agent selects the most relevant PSI entries for the specific role, reorders and rewrites bullets to match the language of the job description, adjusts the summary statement to address what the employer is looking for, and flags any title-based ATS risk with a suggested reframe specific to that application.

Two ATS scores are displayed side by side above the output: the original resume's estimated compatibility with the job description and the tailored resume's score, with a one-line explanation of the most significant improvement. The ATS Scorer agent generates both scores by evaluating keyword match, role alignment, experience framing, and formatting. This score is an estimate, not a certified output, and is disclosed clearly to the user.

The output is displayed in a TipTap WYSIWYG editor with changed sections highlighted so the user can review and accept or reject individual modifications. The resume can be downloaded as PDF or DOCX. Up to 15 generated versions and 3 stored job descriptions are supported per user.

5.6. Phase 4: Public Profile and Certificate

Every user has a public profile page accessible via a shareable URL without requiring login. The page displays their name, target PM role, current readiness score, skill category breakdown, completed stages with dates, and a heatmap of platform activity over the previous 90 days.

Three certificate milestones are issued automatically based on progress. The first is issued after Stage 1 completion and carries an In Progress label. The second is issued at 50 percent stage completion and carries a Midway label. The third is issued at full completion and carries a Verified label. Each certificate exists at its own public URL and includes the user's name, target role, readiness score, skill breakdown, stages completed, and completion date. Three sharing options are available: copy link, share to LinkedIn with a pre-filled caption, and download as PDF.

The discovery module is a publicly accessible quiz at a separate route requiring no login. It presents 10 questions where each answer maps to a PM role type. After the final question it displays the top two matching role types with plain language descriptions of the day-to-day reality and an honest summary of the downsides. A call to action routes the user to the main onboarding flow.

5.7. Phase 5: Testing and Launch

Before deployment the full platform flow is tested end to end using at least 5 representative profiles covering a developer, a consultant, an MBA, a designer, and a recent graduate. AI output quality is reviewed manually across all agents with particular attention to whether PSI

rewriting produces accurate impact statements rather than fabricated metrics, whether the skill gap analysis feels specific to the individual rather than generic, and whether ATS scores move meaningfully between original and tailored resumes.

Mobile responsiveness is reviewed across all pages. Error monitoring is configured via Sentry. Analytics are set up via Vercel Analytics. The platform is deployed to production and the first user cohort is onboarded, with drop-off points monitored before any new features are considered.