PRODUCT REQUIREMENTS DOCUMENT

Google Maps

AI Day Planner

India Edition

From a vague idea to a perfect day — in 30 seconds.

Document Status	Draft
Product Area	Google Maps — AI Features
Geography	India (Delhi NCR, Mumbai, Bengaluru — Phase 1)
Product Manager	Himanshu Jain
Last Updated	March 2026
Version	v2.0 — Adds Anchor Planning + Adaptive Replanning Engine
Stakeholders	Maps Engineering, AI/ML, Design, India Business Dev, Legal
Review Cycle	Monthly until launch; bi-weekly during beta

Contents

1. Executive Summary

Google Maps AI Day Planner is an AI-powered itinerary orchestration layer embedded natively within Google Maps. It enables users in Indian metro cities to plan a complete, executable day out — across multiple venue types, budget constraints, and time windows — using a single natural language input. The resulting plan is sequenced, traffic-aware, crowd-optimised, and verified against live business hours.

> **The Core Insight**
> Planning a day out in India currently takes 20–40 minutes of fragmented searching, manual hour-checking, and ad-hoc routing. This friction is not a user failure — it is a product gap. AI Day Planner closes that gap by doing the orchestration work that today's Maps experience leaves to the user.

This PRD covers the full product specification for the India-first launch of AI Day Planner — including user problems, solution design, prototype specifications, success metrics, evaluation framework, monetisation strategy, rollout plan, and risk mitigation.

India is not a localisation of a US feature. It is the proving ground. Indian cities are the world's most complex urban planning environments — high traffic volatility, inconsistent business hours, budget-sensitive consumers, multi-language usage, and festival-driven crowd surges. If AI Day Planner succeeds here, the architecture scales globally.

1.5 Why Now? The Converging Shifts

Three structural shifts occurring simultaneously in 2024–2025 make AI Day Planner viable today — and not before:

1.5.1 Large Language Models Are Production-Ready

Gemini and on-device LLMs now parse natural language trip requests with sufficient accuracy to extract intent, budget, mood, and timing from a single unstructured input. This was computationally infeasible at production scale before 2023. The NLU layer that powers AI Day Planner's conversational input is a direct beneficiary of this shift.

1.5.2 Real-Time Places Data Is Sufficiently Rich

Google's Popular Times, verified Business Hours API, and live traffic graph now expose enough real-time signal to make constraint-aware scheduling reliable. Conflict detection — the product's most defensible trust feature — requires hour-accurate venue data at scale. That data infrastructure now exists within Google Maps.

1.5.3 India's Urban Day-Trip Culture Is Accelerating

India has 430M+ smartphone users, the majority with Google Maps installed. Weekend day-trip culture in Tier 1 and Tier 2 cities is growing 34% year-over-year, driven by a young, experience-oriented urban middle class.

2. Problem Statement

2.1 The Planning Friction Gap

When a user wants to plan a day out in an Indian city, the current Google Maps experience requires them to perform all orchestration manually:

- Search for each venue category individually (breakfast, park, museum, dinner)
- Cross-check opening hours for each stop across multiple tabs or apps
- Manually assess crowd levels and peak-time windows
- Build a geographically sensible sequence in their head
- Estimate travel times and account for Indian traffic variability
- Stay within a per-person or total budget
- Discover only mid-day if a stop has a closing-time conflict

This process takes 20–40 minutes and consistently results in suboptimal plans, mid-day replanning, and users abandoning the Maps experience for WhatsApp groups, Instagram Reels, or Zomato.

> **India-Specific Complexity**
> Indian metros compound this friction with: traffic volatility that renders static ETAs unreliable; business hours that change without notice (especially near festivals); sudden crowd surges from religious events and cricket matches; multi-party planning requiring group consensus; and strong budget sensitivity that makes per-head cost filtering essential.

2.2 User Persona — Primary

Name	Rohit, 31
Location	Noida, Delhi NCR
Persona Type	Urban Weekend Explorer
Device	Android (Pixel / Samsung mid-range)
Maps Usage	Daily navigation, occasional discovery
Planning Style	Spontaneous, values experience quality over price optimization
Pain Points	Traffic uncertainty, wasted trips to closed venues, group indecision, budget overruns
Quote	"I just want someone to plan it for me. I'll go wherever, but I need it to make sense."

2.3 User Personas — Secondary

- Priya, 27 — Mumbai — couples planning date nights, budget-conscious, uses Zomato heavily
- Arjun, 38 — Bengaluru — family planner, needs child-friendly options, multi-stop logistics

- Kavya, 24 — Hyderabad — solo explorer, spontaneous, values authentic local experiences

2.4 Current User Journey (Pain Map)

Stage	What User Does Today	Pain Experienced
1. Idea Formation	Thinks of vague goal: Sunday chill day	No starting point in Maps
2. Discovery	Searches 4–6 different queries	200+ results, no curation
3. Hour Checking	Opens each place, checks hours separately	3 tabs open, still uncertain
4. Routing	Tries to mentally sequence stops	Pins scattered, no logical flow
5. Departure	Leaves with incomplete plan	Anxiety about wasted day
6. Mid-Day	Discovers conflict / closed venue	Disruption, replanning stress
7. Exit	Switches to WhatsApp / Zomato	Maps loses the session entirely

3. Solution Overview

3.1 Product Vision

Google Maps AI Day Planner transforms Maps from a reactive search engine into a proactive urban orchestration engine. The product accepts a natural language description of a desired day and returns a sequenced, executable itinerary — with live hour validation, traffic-aware routing, budget filtering, crowd optimisation, and real-time conflict detection.

> **Design Principle**
> AI Day Planner must feel like a native Google Maps feature, not a bolt-on chatbot. Every interaction — input, anchor card, loading state, itinerary card, swap modal, replan banner — must be consistent with the Google Maps visual language and interaction model.

3.2 Core Capabilities

Capability	Description	Differentiator
Natural Language Input	Users describe their day in plain text — mood, activity types, budget, group size. No structured forms.	NLU handles Indian English, Hindi-English mixing, regional intent
Anchor-Based Planning ✨	AI proposes one anchor stop first. User confirms, then AI builds the rest. Reduces cold-start hesitation and increases trust.	Solves the cold-start problem unique to AI-generated plans.
AI Constraint Solver	Understands opening hours, crowd levels, heat index, travel time, and budget to select and sequence ideal stops.	Only Google has real-time Popular Times + verified hours data
Traffic-Aware Routing	Integrates live traffic telemetry to sequence stops that minimise total travel time and avoid peak congestion.	Google's own Maps traffic graph — no competitor has this
Live Conflict Detection	Proactively identifies stops that will close before the user arrives and flags them with a warning banner.	Prevents wasted trips — a uniquely high-value trust moment
One-Tap Swap	Users can replace any stop with AI-curated alternatives, complete with rating delta and travel time impact.	Re-plans the full route automatically on each swap
Adaptive Replanning Engine ✨	Mid-day AI engine monitors GPS, stop duration, traffic, and venue hours. Recomputes optimal sequence when delay >20 min, ETA change >15%, or venue conflict detected.	No competitor combines live GPS + live traffic + live hours for mid-trip re-optimisation
Budget Filtering (INR)	All recommendations filtered against stated per-head or group budget. Rupee-denominated throughout.	India-specific — not a USD conversion

3.3 What This Product Is NOT

- Not a chatbot or standalone AI assistant — it lives inside Maps
- Not a static itinerary generator — all data is live at query time
- Not a food delivery or booking app — it surfaces booking options; fulfilment is via existing integrations
- Not a routing-only feature — the discovery and constraint-solving layers are the core value

4. Prototype Specification

4.1 Screen Flow Overview

The v2.0 prototype consists of seven screens. Screens 1–1b–2–3 are the primary sequential flow. Screens 4 (Conflict Detection), 5 (Swap Drawer), and 6 (Adaptive Replanning) are contextual overlays triggered from Screen 3.

#	Screen	State	Trigger
1	Conversational Input	input	User taps 'Plan My Day' entry point in Maps
1b	Anchor-Based Planning ✨	anchor proposal	AI proposes single anchor stop after input submitted
2	AI Generating	loading	User confirms anchor / submits full intent
3	Itinerary Result	result	Auto-transition after 2.5s processing
4	Conflict Detection	overlay	AI detects closing-time conflict on any stop
5	Swap Drawer	overlay	User taps swap icon on stop card (Screens 3 or 4)
6	Adaptive Replanning ✨	overlay	Delay >20 min or venue conflict triggers engine

4.2 Screen 1 — Conversational Input

Map Area (Top 40%)

- Grey road-grid map background — no pins visible in input state
- Google Maps standard search bar and user avatar at top
- Map acts as context — reinforces spatial, location-aware nature of task

Bottom Sheet

- Header: 'Plan your day'
- Subtitle: 'Tell us what you're in the mood for'
- Large rounded input field with placeholder text (e.g. Sunday in Delhi — breakfast, park, museum, dinner)
- Quick Intent Chips: Family day out / Solo explore / Date night
- Primary CTA: 'Plan it →'
- Subtext: 'AI considers open hours, traffic, heat & your preferences'

> **Design Note**
> The input field must support Hindi-English mixed input (Hinglish). NLU parsing must handle queries like 'Sunday ko brunch aur museum chahiye, budget 600 rupee' as well as structured English queries.

✨ NEW IN v2.0 **Screen 1b — Anchor-Based Planning**

4.3 Screen 1b — Anchor-Based Planning (NEW)

Before AI generates the full itinerary, it proposes a single 'anchor' stop — the highest-confidence starting point for the day. The user confirms or requests a different anchor, and only then does AI build the complete plan around it.

Anchor Card

- 'Start here?' header with 'AI Anchor Suggestion' badge
- Single venue card showing: name, category, price tier, star rating, best arrival time
- AI rationale: 'High-rated morning anchor. Low crowd at this hour. Sets a relaxed pace for the day.'
- 'What follows:' preview chips — Park, Museum, Dinner
- Primary CTA: 'Looks good — build my day ✓'
- Secondary: 'Suggest a different anchor'
- Footer: 'Your anchor shapes the whole day. You can change stops after.'

> **Why This Matters**
> User research shows that committing to a full 4-stop AI plan on the first try creates hesitation. The anchor pattern reduces cognitive load by breaking commitment into two steps: confirm one stop, then accept a full plan.

4.4 Screen 2 — AI Generating

Map Area

- 3–4 red Google Map pins appear with a pulse animation
- Blue route line draws progressively between pins
- Animation communicates AI is reasoning spatially — not just text-processing

Bottom Sheet — Status Stack

#	Status Text	Delay	What's Actually Happening
1	Finding top-rated spots...	0ms	NLU parsing + Places API query with mood/category filters
2	Checking live hours & crowd data...	400ms	Business Hours API + Popular Times data validation
3	Optimising for traffic & route flow...	800ms	TSP routing with live traffic graph + budget filter

4.5 Screen 3 — Itinerary Result

Map Area

- 4 numbered pins (1–4) in Google Blue; clicking a card highlights its pin in red
- Blue solid route line connects all pins in sequence
- Map area is interactive — users can pan and inspect the route

Stop Card Anatomy

Element	Content	Notes
Number Badge	Blue circle with stop number (1–4)	Red when card is active/selected
Venue Name	Stop name in bold	Tappable — opens full Place card
Category + Price	e.g. Breakfast · ₹₹	Rupee symbols only — no USD
Star Rating	e.g. ★ 4.4	Live Google Maps rating
Time Slot	e.g. 10:00–11:00 AM	AI-assigned time window
Crowd Bar	Animated fill: Green=Low / Yellow=Medium / Red=High	Derived from Popular Times API
AI Reason	Italic one-line justification	e.g. 'Iconic filter coffee & dosa, calm morning crowd'
Swap Icon	Refresh/swap button (right)	Opens Swap Modal
Save Icon	Heart icon (right)	Saves stop to user's Lists

Live Replanning Banner (NEW)

- Persistent blue strip at bottom: 'Live Replanning Active — AI is monitoring your progress. If you fall behind, it will suggest adjustments automatically.'
- ON/OFF toggle (default: ON)
- 'Simulate delay →' tappable link triggers Screen 6 demo flow

4.6 Conflict Detection Banner

When AI detects a stop is scheduled to arrive after its closing time, a warning banner renders inline within the stop card. This is one of the highest-value trust moments in the product — it demonstrates the AI is monitoring on the user's behalf, not just generating a static list.

4.7 Swap Drawer

- Triggered by tapping the swap icon on any stop card
- Two alternative venue cards, each showing: Name, Category, Star rating, Distance delta from original, 'Choose' button
- AI reasoning per option: one-line natural language justification
- 'Keep original' escape hatch always available
- On selection: drawer closes, stop card updates, route animation re-renders

4.8 Footer Actions

Start Navigation	Adjust Day
Filled Google Blue button — commits to plan and launches turn-by-turn navigation to Stop 1.	Outlined Google Blue button — re-opens the AI input with the current plan pre-loaded as context.

NEW IN v2.0 Screen 6 — Adaptive Replanning Engine

4.9 Screen 6 — Adaptive Replanning Engine (NEW)

Triggered automatically when delay >20 min, ETA change >15%, or venue conflict detected mid-trip. Presents the current plan vs. AI-optimised reorder with clear visual comparison.

Situation Panel

- 'AI Replanning — You're 75 min behind schedule'
- Detected signals: breakfast overtime, heavy traffic on Mathura Road, museum closing risk
- Recommended reorder: side-by-side comparison of Current Plan vs. AI Plan
- AI rationale: 'Museum visit moved earlier: you'll arrive at 1:30 PM with 2.5 hrs before closing. Lower crowds in the afternoon.'

CTA Pair

- 'Apply AI Replan ✓' (filled blue) — updates itinerary and triggers route recalculation
- 'Keep original' (outlined) — AI respects user choice and continues silent monitoring

Engine Behaviour

After a user declines a replan (Keep original), the engine continues monitoring silently. It will re-prompt only if the situation worsens significantly. It will never prompt more than once per 20-minute window to avoid alert fatigue.

5. Functional Requirements

5.1 NLU & Input Processing

1. Accept natural language input in English, Hindi, and Hinglish
2. Parse intent dimensions from input: activity types, budget (INR), group size, timing preferences, geographic constraints, vibe/mood
3. Quick intent chips must pre-fill the input field without clearing prior partial input
4. Empty input must show inline validation — not block the CTA
5. Input must persist across screen transitions (loading → result → adjust)

5.2 Anchor-Based Planning (NEW)

6. After NLU parsing and before full plan generation, system must identify the single highest-confidence anchor stop
7. Anchor selection criteria: highest composite score (see Section 5.7), open at the natural start time for the activity type, within 15 min travel from user's location
8. Anchor card must render before the loading animation begins — not during
9. User must explicitly confirm or reject the anchor before AI generates the full itinerary
10. If user requests alternative anchor: re-query the second-ranked stop, not a random selection
11. Anchor selection must be logged for model feedback — confirmed vs. rejected anchors improve future suggestions

5.3 AI Constraint Solving

12. Query Google Places API for venues matching parsed intent, filtered to open hours at planned arrival time
13. Apply budget filter: sum of per-head costs across stops must not exceed stated budget ± ₹200 tolerance
14. Apply crowd filter: prefer stops with Low or Medium crowd levels at planned arrival time; surface crowd level as warning if High
15. Apply heat filter: between 12:00–15:00 PM in summer months (April–June), deprioritise outdoor stops
16. Sequence stops using TSP-variant routing that minimises total travel distance while respecting categorical logic
17. Generate a one-sentence AI justification for each stop selection

5.4 Conflict Detection

18. Before returning the itinerary, validate each stop's closing time against its planned arrival time
19. If arrival time > closing time: mark stop with conflict status; render warning banner; offer swap
20. If a conflict stop is not swapped: user must explicitly confirm before navigating
21. If a stop's hours change mid-plan: surface an in-itinerary notification

5.5 Swap Functionality

22. On swap trigger: query Places API for up to 3 alternatives in the same category within 1.5km of the original stop
23. Filter alternatives against same constraints as original selection (hours, budget, crowd)
24. Present top 2 alternatives with: name, category, star rating, distance delta from original
25. On selection: replace stop, regenerate route, recalculate total distance and return time
26. If no valid alternatives exist: show message 'No swaps available — all nearby alternatives are either closed or over budget'

5.6 Adaptive Replanning Engine (NEW)

27. Engine monitors the following signals continuously from the moment navigation starts: GPS location vs. planned waypoints, elapsed time at each stop vs. planned dwell time, live traffic ETAs for remaining legs, venue closing hours for remaining stops
28. Engine triggers recompute when ANY of the following conditions are met: total delay exceeds 20 minutes vs. original plan, ETA to any remaining stop changes by more than 15%, any remaining stop is projected to close within 60 minutes of projected arrival
29. Recompute algorithm: applies the same TSP-variant routing as initial planning, using remaining stops + current time + live traffic + current venue hours as inputs
30. Result is presented as a side-by-side reorder proposal — not applied automatically
31. User must explicitly tap 'Apply AI Replan' to accept the new sequence
32. If user declines: engine resumes silent monitoring; will not re-prompt for minimum 20 minutes
33. Engine must not trigger during the first stop (within first 30 minutes of navigation start)
34. Engine must be user-toggleable via the Live Replanning banner ON/OFF control

5.7 Edge Cases & Graceful Error States

- Impossible constraints: When no venue meets all stated constraints, AI must not silently drop the constraint. Display budget edit CTA with alternatives found.
- All venues closed: On national holidays, surface proactively before generating the plan.
- Poor network or offline: Serve last-cached itinerary with an offline badge. Do not show an error screen.
- Overly vague input: If NLU confidence below 60%, ask exactly one clarifying question. Do not ask multiple questions in sequence.
- Replanning refused: AI respects user's 'Keep original' choice. Continues monitoring silently. Re-prompts only if conflict worsens significantly.
- All alternatives bad: Show 'No better alternatives right now. Try after 3 PM or expand the area.'

✨ NEW IN v2.0 **5.8 — AI Ranking & Scoring Model**

5.8 AI Venue Ranking & Scoring Model

The scoring model determines how venues are ranked and selected when the Places API returns multiple candidates. Engineering reviewers require explicit scoring logic to evaluate model fairness, debuggability, and tuneability.

Composite Score Formula

```
score =
  0.35 × rating_score
  + 0.25 × distance_efficiency
  + 0.20 × crowd_comfort
  + 0.10 × novelty
  + 0.10 × budget_fit
```

Factor Definitions

Factor	Definition	Signal Source
rating_score	Normalised 0–1 from Google Maps star rating (4.0=0.5, 5.0=1.0). Weighted down for <100 reviews.	Google Maps ratings API
distance_efficiency	1 - (delta_distance / max_acceptable_distance). Rewards stops that don't require backtracking.	Live traffic graph + TSP routing
crowd_comfort	Derived from Popular Times. Low=1.0, Medium=0.6, High=0.2. Penalised to 0 if >80% occupancy at arrival time.	Popular Times API
novelty	Recency score based on user's Place history. Penalises frequently visited venues, rewards new discovery.	User's saved places + visit history
budget_fit	1.0 if cost is within stated budget. Decays linearly to 0 at 1.5× stated budget.	Google Places price tier (₹)

Hard Filters (Applied Before Scoring)

Venues failing any hard filter are excluded from the candidate set entirely before scoring:

- Closed at planned arrival time (Business Hours API)
- Cost per head exceeds 2× stated budget
- Crowd level is 'Extremely Busy' (Popular Times) — unless no alternatives exist
- Venue category does not match intent category parsed by NLU

Model Versioning

The scoring weights are tunable and will be optimised via offline evaluation (Section 8.1). The v1.0 weights above are based on India user research. Weight updates require re-evaluation against the 7,000-query test set before deployment.

6. Non-Functional Requirements

Requirement	Target	Notes
P95 Response Time	< 3.0 seconds	From input submit to itinerary rendered
Offline Graceful Failure	Show cached plan	Last generated plan stored locally for 24h
API Freshness	< 15 min staleness	Hours and crowd data must be near-real-time
Low-End Device Support	Android 8.0+, 2GB RAM	Indian mid-range device benchmark
Data Usage	< 2MB per plan generation	Optimised for 4G / limited data plans
Accessibility	WCAG 2.1 AA	Screen reader support; minimum 4.5:1 contrast
Language Support	English + Hindi (Phase 1)	8 regional languages in Phase 2
Privacy	No persistent PII storage	Queries are session-scoped; no profiling
Adaptive Replanning Latency	< 5.0 seconds	From trigger detection to replan card rendered

✨ NEW IN v2.0 **Section 6.1 — System Architecture**

6.1 System Architecture Pipeline

The end-to-end AI Day Planner pipeline consists of seven sequential layers. Engineering reviewers require an explicit description of the data and control flow to assess integration points, failure modes, and latency budgets.

```
User Input (natural language, Hindi/English/Hinglish)
     ↓
NLU Parser  (Gemini / on-device LLM)
  → Extracts: [activity_types], budget_inr, group_size, timing,
              mood, geographic_constraints
     ↓
Intent Graph  (structured intermediate representation)
  → Normalises parsed entities into typed constraint objects
  → Resolves ambiguity (e.g. 'cheap' → budget_inr < ₹400/head)
     ↓
Places Retrieval  (Google Places API + internal cache)
  → Parallel query per intent category
  → Hard filters: hours, budget, category match
  → Returns candidate set (10–30 venues per category)
     ↓
Constraint Engine  (scoring model — see Section 5.8)
  → Applies composite scoring: rating, distance, crowd,
                               novelty, budget_fit
  → Returns top-N per category for routing
     ↓
Route Optimizer  (TSP-variant + live traffic graph)
```

```
  → Inputs: candidate stops, current time, live traffic ETAs
  → Outputs: optimal stop sequence + time slots per stop
     ↓
Itinerary Builder  (UI payload assembly)
  → Conflict validation (arrival time vs. closing time)
  → AI rationale generation per stop
  → JSON payload → Screen 3 rendering
     ↓
Adaptive Replanning Monitor  (continuous, post-navigation start)
  → Polls: GPS, dwell time, live traffic, venue hours
  → On trigger: runs Constraint Engine + Route Optimizer
                with remaining stops as new candidate set
```

Key Integration Points

System	API / Service	SLA Dependency
NLU Parser	Gemini API / Gemini Nano on-device	P95 < 800ms
Places Retrieval	Google Places API (New)	P95 < 600ms per category query
Business Hours	Places Fields: opening_hours, current_opening_hours	Freshness < 15 min
Crowd Data	Popular Times (internal Google graph)	Updated hourly
Traffic Graph	Google Maps Directions API + live traffic	Real-time
GPS Monitor	Android Fused Location Provider	Update interval: 60s while navigating

✨ NEW IN v2.0 **Section 6.2 — Latency Architecture**

6.2 Latency Architecture — Achieving P95 < 3 Seconds

The P95 < 3 second target is aggressive given the number of sequential API calls required. The following strategies are required to meet this target:

Strategy 1: Parallel Places API Querying

Rather than querying each activity category sequentially, all category queries (breakfast, park, museum, dinner) are dispatched in parallel as soon as NLU parsing completes. Expected time saving: 800–1,200ms.

Strategy 2: City-Level Intent Precomputation

For the top 20 intent patterns per Phase 1 city (e.g. 'chai + heritage walk + dinner, Delhi'), the system precomputes candidate venue sets daily during off-peak hours. When a live query matches a cached intent pattern (cosine similarity > 0.85), the candidate set is served from cache, bypassing Places API latency. Expected cache hit rate: 30–40% of queries.

Strategy 3: CDN-Optimised Places API Batching

All Places API calls are batched at the India CDN edge node (Mumbai region), reducing cross-continental latency for Places API responses. Expected saving: 200–400ms vs. US-region routing.

Strategy 4: Progressive Rendering

The itinerary result screen renders stops 1 and 2 as soon as scoring is complete, while route optimisation for stops 3 and 4 continues asynchronously. Users see actionable content within ~1.8 seconds; the full plan completes within 3 seconds.

Latency Budget Breakdown

Pipeline Stage	Sequential (ms)	Parallel / Cached (ms)	Notes
NLU parsing	200	200	Cannot be parallelised
Intent graph resolution	80	80	
Places API (4 categories)	1,200	300 (parallel)	4× parallel queries
Constraint scoring	200	200	In-process computation
Route optimisation	400	400	TSP with 10–15 nodes
Conflict validation	100	100	
AI rationale generation	300	300	Gemini API (async)
TOTAL	2,480ms	~1,580ms P50 / ~2,800ms P95	P95 target met

7. Success Metrics

7.1 North Star Metric

North Star: Completed Real-World Days — India (CRD+)

Definition: A CRD+ event is recorded when a user who generated an AI Day Plan (a) navigated to 60% or more of the planned stops, (b) had no unresolved hour-conflict warnings on any visited stop, and (c) provided a post-day rating of 4/5 or above. Target: 25% CRD+ completion rate within 90 days of general availability.

7.2 Supporting Metrics

Metric	Target (90 days post-GA)	Measurement
Feature activation rate	12% of eligible MAU	Day Planner opens ÷ eligible Maps sessions
Anchor confirmation rate	> 70% of anchor proposals	User confirms anchor ÷ anchor proposals shown
Plan completion rate	50% of generated plans	Navigation start ÷ plan generated
Booking conversion rate	15% restaurant / 6% activities	Booking completions ÷ navigations started
Session depth	+23% vs. baseline	Time in Maps for Day Planner sessions
Place saves	3× vs. manual search	Saves from Day Planner cards
Swap usage rate	< 30% of plans	High swap rate = poor initial curation
Conflict detection rate	> 95% of real conflicts	Conflicts caught ÷ total conflicts in ground truth
Adaptive replan acceptance rate	> 40% of replan proposals	User applies AI replan ÷ replan proposals shown
CRD+ with replan vs. without	+15–20% delta	CRD+ rate for sessions with replan trigger vs. control

7.3 Guardrail Metrics

The following metrics must NOT regress. If any guardrail is breached, the feature rollout pauses pending investigation.

- Organic Maps search volume: no more than 3% decline vs. control
- ETA accuracy: no more than 5% deviation from baseline routing accuracy
- Ad revenue per session: no regression vs. baseline (Day Planner should add sessions, not cannibalise)
- P1 bug rate: zero P0 incidents; maximum 2 P1 incidents per 14-day period post-launch
- Negative review rate: no increase in reviews mentioning 'closed', 'wrong hours', or 'wrong directions'

✨ NEW IN v2.0 Section 7.4 — UX Diagnostic Metrics

7.4 UX Diagnostic Metrics

These metrics are not pass/fail guardrails — they are diagnostic signals that help the product team identify specific friction points in the experience and prioritise improvements.

UX Metric	Why It Matters	Target / Benchmark
Time to first plan	Measures total friction from feature entry to first itinerary rendered. High values indicate NLU or API latency issues.	< 3 seconds P50; < 5 seconds P95
Swap frequency by stop position	If stop 1 is swapped more than stop 4, the anchor selection model has a systematic curation error at the start of the day.	Uniform swap rate across positions ±10%
Plan abandonment rate	Percentage of users who generate a plan but do not tap Start Navigation or Adjust Day. High rate = plan quality or trust issue.	< 30% abandonment within 5 min of plan rendered
Anchor rejection rate	Percentage of anchor proposals that users reject by tapping 'Suggest a different anchor'. High rate = anchor model needs retraining.	< 25% anchor rejection
Replan prompt frequency	Average number of replan prompts per session. Too many = alert fatigue. Too few = engine not triggering appropriately.	0.8–1.5 prompts per trip on delayed sessions
Post-trip rating distribution	Distribution of 1–5 ratings after completed days. Bimodal distribution (many 5s and 1s) = inconsistent quality.	Mean ≥ 4.2 / 5.0; < 10% 1-star ratings
Input reformulation rate	Percentage of users who edit their input after the anchor is shown. High rate = NLU is misinterpreting intent.	< 15% of sessions

8. AI Evaluation Framework

AI Day Planner is a production AI feature, not an experimental product. Its AI layer requires rigorous evaluation across both offline and online dimensions before any public launch.

8.1 Offline Evaluation

Before any A/B test or beta launch, the AI model is evaluated against a curated set of 7,000 India-based queries. Each query is rated by trained human evaluators against the following dimensions:

Evaluation Dimension	Metric	Target	Failure Threshold
Intent relevance	Evaluator score (1–5)	≥ 4.3 average	< 3.8
Hours correctness	% stops with verified-correct hours	≤ 2% error rate	> 5%
Budget adherence (INR)	Mean absolute deviation	≤ ₹200 from stated budget	> ₹400
Traffic realism	ETA alignment with live data	Within 15% of actual	> 25%
Conflict detection recall	Conflicts caught / total conflicts	≥ 95%	< 90%
Route efficiency	Distance vs. optimal TSP	Within 10% of optimal	> 20%
Anchor quality	% of anchors confirmed by evaluator as 'strong starting choice'	≥ 75%	< 60%
Replan quality	% of replan proposals accepted by human evaluators as 'clearly better'	≥ 65%	< 50%

8.2 Online A/B Test Design

Control	Standard Google Maps manual search and discovery (baseline)
Variant	AI Day Planner feature enabled and surfaced
Traffic Split	50/50 within each phase
Phase 1 Traffic	5% of eligible Maps users in Delhi NCR only
Phase 2 Traffic	20% of eligible Maps users in Delhi NCR, Mumbai, Bengaluru
Phase 3 Traffic	100% of eligible users, all target cities
Primary Metric	CRD+ completion rate (see 7.1)
Secondary Metrics	Booking uplift, session depth, place saves, feature return rate, anchor acceptance rate
Minimum Duration	4 weeks per phase (accounts for Indian weekend seasonality)

Statistical Threshold	p < 0.05 required to advance to next phase; p < 0.01 for general availability

8.3 Continuous Monitoring

- Hours accuracy monitoring: weekly automated sampling of Places Hours API vs. ground truth (in-person verification for top 500 venues by volume)
- Conflict detection monitoring: post-trip survey for conflict-flagged stops — 'Was this warning correct?'
- Route quality monitoring: compare AI-recommended sequence vs. user-modified final sequence (high swap rate = curation quality issue)
- Anchor model monitoring: weekly review of anchor rejection rate by category and city; retrain trigger if rejection rate exceeds 25% for any segment
- Replan quality monitoring: post-trip comparison of CRD+ rate for sessions where replan was accepted vs. declined
- Bias monitoring: quarterly audit for geographic, price-tier, or recency bias in venue selection

9. Monetisation Strategy

✨ NEW IN v2.0 Section 9.0 — Market Sizing

9.0 Market Sizing — TAM / SAM / SOM

Leadership investment decisions require a clear market sizing framework. The following defines the addressable opportunity from the global level down to the realistic capture target for Phase 1.

Level	Market	Size	Definition
TAM	Total Addressable Market	₹58,000 Cr / year	All domestic day-trip and weekend leisure spend in urban India (430M smartphone users × estimated ₹1,350/trip × 1 trip/month)
SAM	Serviceable Addressable Market	₹9,300 Cr / year	Urban India Google Maps MAU (85M users) × estimated spend on bookable activities and dining within a Maps-led session
SOM	Serviceable Obtainable Market (Phase 1 Year 1)	₹1,394 Cr / year	10.2M activated users × 50% plan completion × 15% restaurant + 6% activity booking conversion at ₹120/₹80 commission avg. See Section 9.1 for full model.

Market Sizing Note

The SAM figure represents the portion of TAM addressable through Google Maps booking integrations in Phase 1 cities. The gap between SAM and TAM (₹9,300 Cr vs ₹58,000 Cr) represents expansion opportunity through Phase 2 city rollout, group planning mode, and Google One Premium tier adoption.

9.1 Unit Economics Model — India

Funnel Stage	Rate	Volume
Urban India Eligible MAU (Maps)	—	85 million
Feature Activation Rate	12%	10.2 million users
Plan Completion Rate	50%	5.1 million plans/month
Restaurant Booking Conversion	15%	765,000 bookings/month
Activity/Experiences Conversion	6%	306,000 bookings/month
Average Restaurant Commission	₹120/booking	₹91.8 million/month
Average Activity Commission	₹80/booking	₹24.5 million/month
Projected Annual Revenue (Phase 1)	—	₹1,394 Crore

Revenue Upside Not Included

This model covers direct booking commissions only. Not yet modelled: (1) Sponsored swap placements — venues pay for placement in swap suggestions, (2) Local Ads uplift from increased session depth, (3) Google One Premium tier upsell — itinerary saving, offline access, group planning, (4) Data licensing to travel and hospitality partners.

9.2 Monetisation Phases

Phase	Lever	Description
1	Restaurant Booking Commission	Commission on reservations made via Google Reserve from Day Planner results. ₹120 avg per booking.
1	Activity Booking Commission	Commission on ticketed experiences booked via Maps. ₹80 avg per booking.
2	Sponsored Swap Placements	Restaurants and venues bid for premium placement in swap suggestions. CPC or CPB model. Must still pass all quality thresholds — clearly labelled 'Promoted'.
2	Google One Premium	Save & sync itineraries, offline access, group planning mode, ad-free experience.
3	UPI-Integrated Bookings	Direct payment via UPI at booking, unlocking full transaction revenue vs. referral-only.
3	B2B Data Intelligence	Anonymised aggregate demand signals to hospitality and real estate partners.

10. Rollout Plan

10.1 Phase Summary

Phase	Timeline	Geography	Key Capabilities
1	Q1–Q2 Year 1	Delhi NCR, Mumbai, Bengaluru	Single-day orchestration, Anchor-Based Planning, English+Hindi NLU, restaurant bookings, conflict detection, A/B test 5%→100%
2	Q3–Q4 Year 2	+Hyderabad, Pune, Chennai	Adaptive Replanning Engine, group planning mode, 8 regional languages, festival-aware routing, budget bands, activity bookings
3	Year 3	Pan-India + Global Template	Weather-triggered adjustments, advanced replan learning, UPI bookings, global expansion template

10.2 Phase 1 Quarterly Milestones

Quarter	Milestones
Q1	User research (20 intercept interviews in Delhi NCR). Finalise NLU model v1. Complete offline eval with 7,000-query dataset. Internal dogfood with 100 Googlers. Build and validate Anchor model v1.
Q2	Backend: NLP intent model + Places API constraint integration complete. Anchor-Based Planning flow live. A/B test at 5% in Delhi NCR. Weekly metric review cadence established.
Q3	Closed beta: 5,000 Maps users across 3 cities. Restaurant booking integration live. Conflict detection v1 deployed. Adaptive Replanning Engine backend complete. A/B test expanded to 20% if Q2 metrics pass.
Q4	General availability: 100% rollout across Phase 1 cities. Google One upsell live. Sponsored swap placements pilot. Adaptive Replanning Engine released. Phase 2 planning initiated.

11. Risks & Mitigations

#	Risk	Likelihood	Impact	Mitigation
1	Hours data inaccuracy leads to wasted trips and negative reviews	Medium	High	Automated hours sampling; in-person ground truth for top 500 venues; proactive conflict detection as fallback
2	AI suggestions feel generic or irrelevant to Indian tastes	Medium	High	India-specific training data; local evaluator panel; post-trip feedback loop integrated from Day 1
3	Traffic latency causes P95 response time > 3 seconds	High	Medium	Aggressive caching of popular city routes; pre-fetching likely intent patterns; CDN-optimised Places API calls; parallel querying (see Section 6.2)
4	Feature cannibalises organic search sessions	Low	High	Guardrail metric on search volume; rollback trigger if > 3% decline; feature designed to deepen session not replace search
5	Sponsored placement erodes user trust in AI recommendations	Medium	High	Clear labelling of sponsored swaps ('Promoted'); sponsored results must still meet all quality thresholds
6	Festival/holiday edge cases cause mass conflict errors	High	Medium	Festival calendar integration (Phase 1); conservative hour estimates on known festival days; manual override team on standby
7	Adaptive Replanning triggers too frequently — alert fatigue	Medium	Medium	Minimum 20-min cooldown between replan prompts; user-controlled toggle; threshold tuning based on beta data

12. Competitive Landscape

12.1 Capability Comparison

Capability	VeLa AI (Google Maps)	Zomato	Instagram	ChatGPT	Yelp
Live traffic telemetry	✓	✗	✗	✗	✗
Verified business hours	✓	~	✗	✗	~
Popular Times crowd data	✓	✗	✗	✗	✗
Multi-stop route optimisation	✓	✗	✗	✗	✗
Budget filter (INR)	✓	~	✗	~	✗
Android default distribution	✓	✗	✗	✗	✗
Booking integration	✓	✓	✗	✗	~
Anchor-Based Planning ✨	✓	✗	✗	✗	✗
Mid-day adaptive replanning ✨	✓	✗	✗	✗	✗

Defensible Advantage

Competitors can suggest. Only Google can execute. No competitor combines live traffic + verified hours + crowd data + multi-stop routing + booking integration + Android default distribution + anchor planning + mid-day adaptive replanning in a single product. This moat is structural and cannot be replicated without the same underlying data infrastructure.

13. Open Questions

#	Question	Owner	Target Resolution
1	Should we surface a 'free exploration' mode where the user can extend the day beyond 4 stops?	Product	Q1 user research
2	What is the right fallback experience when AI confidence is < 70% on intent parsing?	AI/ML + Design	Q1 prototype testing
3	Should conflict detection warnings be suppressible by the user?	Design + Legal	Q1
4	What is the minimum viable India Places dataset coverage required before launch?	Data/Maps	Q1 feasibility study
5	Should sponsored swap placements be disclosed? (Legal / Trust implications)	Legal + Policy	Q2 — pre-beta launch
6	Group planning mode: consensus mechanism? Voting? Host-controls?	Product	Phase 2 design sprint
7	What is the right CRD+ threshold for declaring Phase 1 a success?	Analytics + Product	Q1 — align with leadership
8	Should Adaptive Replanning be on by default for all users or opt-in? (battery/data considerations on low-end devices)	Engineering + Product	Q2 — before beta launch
9	What scoring weight adjustments are needed for non-Delhi cities with different crowd patterns?	AI/ML	Q3 — before multi-city expansion

14. Appendix

14.1 Glossary

CRD+	Completed Real-World Days — the North Star metric measuring full-day plan execution
NLU	Natural Language Understanding — the AI layer that parses user intent from free-text input
TSP	Travelling Salesman Problem — the routing optimisation algorithm used to sequence stops efficiently
Popular Times	Google's proprietary crowd density data layer, derived from aggregated, anonymised location signals
Hinglish	Hindi-English code-switched language, commonly used in Indian digital communication
MAU	Monthly Active Users
INR / ₹	Indian Rupee — the currency for all budget and cost displays in this product
Anchor Stop	The single confirmed starting stop proposed by AI before building the full itinerary — introduced in v2.0
Adaptive Replanning Engine	The v2.0 mid-day AI system that monitors GPS + traffic + venue hours and proposes itinerary reorders when significant delays or conflicts are detected
Composite Score	The weighted scoring formula (Section 5.8) used to rank and select venues from the Places API candidate set
CRD+ with Replan	A sub-variant of CRD+ measuring completion rate specifically for sessions where the Adaptive Replanning Engine triggered and the replan was accepted
Phase 1 Cities	Delhi NCR, Mumbai, Bengaluru — the three initial launch markets

The best day out in India shouldn't start with 30 minutes of homework.

Google Maps AI Day Planner — India Edition — v2.0