



Developer Handoff

Last Updated: 6 June 2026, 21:01

Prepared for TravAI Solutions
Project TravAI App

© We Are Affective Ltd 2026. All rights reserved.
Private and Confidential.

Contents

Introduction	3
Design Phase Deliverables	3
What Has Been Built	5
What Matters Most	6
What Matters Most	7
Known Constraints & Decisions	8
System Copy Requirements	10
System Copy Requirements	10
Acceptance Criteria	13
Emotional Acceptance Criteria	13
Brand and Voice Criteria	13
Interaction Acceptance Criteria	14
Performance Criteria Affecting Emotional Experience	15
Copy and Content Criteria	15
Pre-Launch Emotional Review	16

Introduction

This document exists to ensure that what was designed for TravAI actually gets built, and that the decisions made during the design phase are not revisited, diluted, or lost in translation. It captures the things that fall between the cracks: the emotional intent behind design decisions, the non-negotiables that cannot be compromised, the constraints that shaped the design, and the criteria by which the build will be judged. The research and strategy work revealed that TravAI's success depends entirely on transforming someone who feels burnt out and overwhelmed into someone who feels understood and catered for. That transformation cannot happen by accident during development.

This is not a functional specification, project plan, or design system. User stories, technical architecture, API documentation, and component libraries live elsewhere. This document focuses on the qualitative decisions that determine whether TravAI feels like a trusted guide or just another booking platform with AI marketing. It explains why certain approaches were chosen over others, what would break the experience if changed, and how to recognise when the build is staying true to the design intent versus when it is drifting away from it.

The development team, QA, product owners, and anyone responsible for making decisions during the build phase should read this document before writing their first line of code, and refer back to it whenever questions arise about priority, approach, or quality. When in doubt about whether a particular implementation choice serves the user experience or just the technical requirements, this document provides the answer. The goal is to build something that makes people want to book a holiday, not just something that allows them to.

Design Phase Deliverables

These documents form a complete system for understanding what TravAI needs to be and why. They emerged from structured workshops, exercises, and conversations designed to surface what the team already knew but had not yet articulated about the product's true purpose and potential. Each document captures a different dimension of the same insight: that TravAI succeeds or fails based on whether it makes someone feel genuinely understood rather than simply processed.

The research documents establish the emotional reality TravAI is designed to address. The brand and experience documents translate that understanding into specific voice, visual, and interaction principles. The technical documents ensure those principles can be implemented without compromise. Together, they provide both the why and the how for every decision that will be made during development.

Developers should read the Core Principles and Emotional Journey documents first, as these establish the fundamental criteria by which all implementation decisions should be judged. The Voice Guidelines and Visual Direction documents should be consulted whenever questions arise about tone, presentation, or user interface choices.

Document	Category	Primary Audience	Summary
Core Principles	Research & Strategy	All team members	Five non-negotiable principles that emerged from user research. Every feature, interaction, and design decision must serve these principles or risk breaking the core user experience.
Emotional Journey Map	Research & Strategy	Design, Development, QA	The complete emotional arc from burnt-out arrival to confident booking. Maps specific feelings, triggers, and turning points that the product must create to succeed.
Aspiration Gap Analysis	Research & Strategy	Product, Strategy	Documents the gap between where users arrive emotionally and where they need to be to book with confidence. Identifies the specific barriers that prevent this transformation.
Trust-Building Moments	Research & Strategy	Design, Development	Catalogues the precise moments when users decide whether to trust the AI or abandon the experience. Critical for implementation priority and QA testing.
Voice Guidelines	Brand & Experience	All content creators	Complete voice definition including rhythm rules, forbidden phrases, register guidelines, and examples. Essential for any text that appears in the product.
Visual Direction	Brand & Experience	Design, Development	Visual principles, density guidelines, colour approach, and interface behaviour. Includes specific directions that are definitively off the table.
Interaction Principles	Brand & Experience	Development, QA	How the AI should behave, respond, and guide users through conversations. Defines the difference between guidance and coercion.
Content Strategy	Brand & Experience	Content, Development	Framework for all written content including question design, explanation copy, and reasoning transparency requirements.
Technical Constraints	Technical	Development, Architecture	Non-functional requirements that support the emotional experience. Response time targets, conversation memory, and personalisation technical requirements.
Quality Criteria	Technical	QA, Development	Measurable signals that indicate whether the emotional experience is working. Success metrics tied to user feelings rather than just completion rates.

What Has Been Built

This is what has been built through the strategy and design phase, and what the development team is now responsible for bringing to life.

Who the brand is TravAI is a well-travelled friend who happens to work in the industry. Someone who listens before speaking, asks thoughtful questions, and remembers exactly what you said about needing to properly switch off this year. The voice is unhurried, quietly confident, and entirely focused on the person in front of it, never performing expertise, always using it in service of making someone feel genuinely understood.

Who the user is They arrive burnt out from their regular job and overwhelmed by the prospect of planning what should be a relaxing break. They fear wasting money on the wrong choice and spending their precious time doing research that feels like a second job. Success for them means booking a trip they feel confident about without having to manage every detail themselves, finally being able to look forward to their holiday instead of dreading the process of getting there.

The emotional arc, First contact: Curiosity mixed with hesitancy. They want to believe this will be different but have been disappointed before., **During the journey:** Progressive relief as they realise they are being listened to rather than processed. The AI demonstrates understanding before asking for vulnerability., **After sustained use:** Complete confidence in delegating travel planning. They feel catered for rather than served, and excited about their trip rather than exhausted by booking it.

Critical moments

1. **The opening conversation**, When the AI first asks about their needs instead of immediately presenting options. This is where trust begins or the experience fails entirely.
2. **First recommendation reveal**, When users see their own words reflected back as the reason for each suggestion. The AI must visibly connect input to output or lose credibility permanently.
3. **Budget vulnerability**, When financial information is requested. This must happen only after trust has been established through demonstrated understanding.
4. **Deviation response**, When a user goes off the suggested path. The AI must adapt rather than redirect, proving it guides without controlling.
5. **Booking confirmation**, When the user commits. They should feel relief and anticipation, not anxiety about whether they made the right choice.

Design principles, **Never assume the user's wants, listen for their insights**, Every recommendation must be visibly connected to something the user actually said, **Guide thinking without forcing paths**, Offer clear direction while always leaving space to deviate and explore, **Make the AI's reasoning visible at every step**, Show why each suggestion was made using the user's own words and priorities, **Reveal one decision at a time**, Present choices progressively to reduce cognitive load rather than reorganising it, **Build trust before asking for vulnerability**, Demonstrate understanding through early suggestions before requesting personal or financial information

What Matters Most

These are not preferences or nice-to-haves. They are load-bearing elements of the design that cannot be compromised, simplified, or deprioritised without fundamentally breaking the user experience. When technical constraints or timeline pressures create the temptation to simplify, these elements must be protected first. Everything else can be adjusted, but these cannot.

What Matters Most

Non-Negotiable	Why It Matters	Risk of Compromise
<p>The AI must initiate conversation before presenting any options</p>	<p>Trust begins when the product demonstrates it wants to understand rather than immediately sell. Users arrive sceptical from platforms that showed them everything upfront and left them overwhelmed. The opening interaction establishes whether this is guidance or another directory.</p>	<p>If users see a wall of options on first load, they immediately categorise TravAI as "booking.com with AI marketing" and bring all their existing resistance. The differentiation is lost in the first ten seconds and cannot be recovered.</p>
<p>Every recommendation must include visible reasoning that references the user's actual input</p>	<p>Users need to see the thread connecting what they said to what they received. "Because you mentioned needing somewhere the kids can roam freely" proves the AI was listening. Generic explanations like "based on your preferences" destroy credibility entirely.</p>	<p>Without specific reasoning, recommendations feel algorithmic rather than personal. Users conclude the AI isn't actually understanding them, trust collapses, and they either abandon or continue with reduced confidence that leads to booking anxiety.</p>
<p>Budget questions cannot appear before the AI has demonstrated understanding</p>	<p>Financial vulnerability requires established trust. Users must see evidence of the AI's ability to understand their travel needs before they feel safe sharing what they can afford. The sequence cannot be reversed for technical convenience.</p>	<p>If budget appears in early onboarding, users experience it as data extraction rather than personalisation. They either provide false information to protect themselves or abandon entirely. Either outcome breaks the personalisation engine the product depends on.</p>
<p>When users deviate from suggestions, the AI must adapt rather than redirect</p>	<p>The difference between guidance and control is revealed in how the system responds to user autonomy. Attempts to pull users back to original suggestions signal that the AI follows scripts rather than people, destroying the concierge illusion entirely.</p>	<p>If users feel herded back to predetermined paths when they explore alternatives, the experience becomes transactional rather than conversational. They lose confidence in the AI's flexibility and revert to manual browsing behaviour within the app.</p>

Non-Negotiable	Why It Matters	Risk of Compromise
Conversation state and reasoning must persist across all user sessions	Each interaction should build on previous conversations, with the AI remembering not just preferences but the specific language users employed. Starting over each time breaks the relationship that is central to the experience.	If users have to re-explain their needs in subsequent sessions, the AI feels like a search engine with memory problems rather than a trusted guide. The cumulative understanding that builds confidence never develops, and repeat usage drops dramatically.
No screen can present more than three primary options simultaneously	Progressive revelation reduces cognitive load for users who arrived already overwhelmed. More options might seem like better service but actually recreates the analysis paralysis they came to escape. The constraint forces quality of curation over quantity of choice.	If users face multiple complex decisions simultaneously, they experience the same overwhelm that drove them away from other platforms. They either freeze in analysis paralysis or make rushed choices they later regret, both of which undermine confidence in the product.

Known Constraints & Decisions

These decisions were made deliberately during the design phase, with full context from user research and strategic analysis. They represent considered choices about how TravAI should behave to create the emotional transformation from overwhelmed to understood. Revisiting them without this context risks undermining the coherence of the experience and returning the product to the category problems it was designed to solve.

Decision	Rationale	What Not To Do
The AI always asks before it shows	Users arrive burnt out from platforms that overwhelmed them with options immediately. Trust begins when the product demonstrates it wants to understand rather than immediately sell. Opening with conversation rather than directory establishes TravAI as fundamentally different.	Never show destination options, hotel listings, or flight results on the home screen. Never default to a browse/filter interface as the primary experience. Never present choices before gathering context about the user's specific situation and needs.

Decision	Rationale	What Not To Do
<p>Every recommendation includes specific reasoning that quotes the user</p>	<p>Generic explanations like "based on your preferences" destroy the personalisation illusion entirely. Users need to see their own words reflected back as justification for each suggestion to believe the AI actually understood them rather than just processed them.</p>	<p>Never use template reasoning text. Never say "based on your travel style" without defining what that style means using the user's actual language. Never present suggestions without visible connection to specific user input. Never hide the AI's decision-making process.</p>
<p>Deviations are accommodated, not corrected</p>	<p>When users go off suggested paths, the AI must adapt rather than redirect. This proves the system guides without controlling and maintains the concierge relationship rather than revealing a scripted experience underneath.</p>	<p>Never show "back to recommendations" prompts when users explore alternatives. Never repeat original suggestions when users have moved on. Never use nudges or UI patterns that encourage return to the predetermined path. Never treat user exploration as an error to correct.</p>
<p>Information density stays minimal</p>	<p>Users came here to escape overwhelm, not to reorganise it. Every screen should reduce cognitive load rather than present comprehensive options. The constraint forces quality curation over quantity of choice, which is the core product promise.</p>	<p>Never show more than three primary options on any screen. Never present full feature matrices or comparison tables. Never use dense layouts that require scrolling to understand all available choices. Never add options "for completeness" if they dilute focus.</p>
<p>Voice stays conversational, never corporate</p>	<p>The research revealed users want to feel like they're talking to a knowledgeable friend, not interfacing with a booking system. Corporate language immediately signals this is just another platform with AI marketing rather than a genuinely different experience.</p>	<p>Never use "powered by AI" or technology-forward language. Never say "seamless," "comprehensive," or other overused platform promises. Never use passive voice when the AI is doing the work. Never write copy that could appear in a corporate presentation or press release.</p>
<p>Conversation state persists indefinitely</p>	<p>Each interaction builds on previous conversations, with the AI remembering not just preferences but the specific language users employed. Breaking this continuity destroys the relationship that differentiates TravAI from search-based alternatives.</p>	<p>Never reset conversation context between sessions. Never ask users to re-explain previously discussed preferences. Never treat returning users as new users. Never lose the specific language users employed in favour of categorised preferences.</p>

Decision	Rationale	What Not To Do
No urgent or scarcity language ever appears	Users arrive stressed and need to feel unrushed. Any pressure tactics immediately break the calm, supportive experience that is core to the brand promise. The voice stays unhurried even when highlighting time-sensitive information.	Never use countdown timers in copy. Never say "book now," "limited availability," or "don't miss out." Never use exclamation marks as default punctuation. Never manufacture urgency to drive conversion. Never pressure users to move faster through any part of the experience.
Progress indicators show conversation stages, not transaction steps	The experience is positioned as a guided conversation leading to understanding, not a funnel leading to purchase. Progress indication should reinforce the consultation metaphor rather than revealing a predetermined sales process.	Never use traditional checkout progress bars. Never label stages with transactional language like "payment" or "booking." Never show steps that reveal the experience as a standard purchase flow. Never use completion percentage that treats conversation as data collection.
Visual hierarchy prioritises content over product features	The design should feel like looking at travel content with subtle AI assistance, not using an AI product that happens to show travel. Interface elements should never compete with destination imagery, trip details, or conversation content for attention.	Never make AI branding more prominent than trip content. Never use interface colours that compete with travel photography. Never foreground product features over travel outcomes. Never design screens that read as software rather than travel guidance.

System Copy Requirements

System Copy Requirements

System copy encompasses all the small-scale text that appears during user interactions, the words that guide, reassure, inform, and respond to users as they move through the product. Unlike marketing copy or feature descriptions, system copy often appears at moments of uncertainty, error, or transition, making it critical for maintaining the trust and emotional experience that TravAI depends on. Every piece of system text must reinforce the feeling that users are being guided by someone who understands them, not processed by a system that is managing them.

The voice guidelines apply with particular intensity to system copy because these moments often determine whether a user continues with confidence or abandons with frustration. A poorly written error message can undo the trust built through several successful interactions. Empty state copy that feels generic can make users question whether the personalisation they experienced earlier was real.

Category	Examples	Tone Guidance	Review Required
Error Messages	Network timeouts, booking failures, invalid inputs, server errors	Never blame the user or use technical language. Frame as temporary setbacks, not failures. "Something's not quite working on our end, give us a moment to sort this out" rather than "Error 500: Server timeout." Maintain the helpful friend voice even during problems.	Content lead + Product owner review for all error states before development begins. QA testing must verify tone as well as functionality.
Empty States	No search results, incomplete profiles, first-time user screens	Avoid generic "no results found" language. Frame as conversation continuation: "I haven't found the perfect match yet, tell me more about what matters to you" rather than "0 results." Use empty states as opportunities to deepen understanding.	Content lead review for all empty state scenarios. Must align with conversation flow and personalisation principles.
Loading States	AI processing responses, booking confirmations, search results	Reinforce the AI as actively working for the user: "Looking for places that match what you described" rather than "Searching..." Show the AI's thought process when possible: "Checking availability for family-friendly options in your budget range."	Copy review required for any loading state longer than 3 seconds. Must feel like the AI is present and working, not just processing.
Confirmation Messages	Booking success, preference saves, trip updates	Celebrate quietly without manufactured excitement. "Your trip is sorted, time to look forward to it" maintains the calm satisfaction voice. Reference specific user inputs: "Confirmed: 5 days in Croatia, just like you wanted."	Content lead review for all confirmation flows. Must reinforce relief and confidence, never anxiety about the choice made.
Onboarding Copy	Welcome messages, feature introductions, first conversation prompts	Establish the conversation immediately: "Let's figure out what you actually want" rather than "Welcome to TravAI." Focus on user benefit, not product capability. No feature tours or platform explanations.	Product owner + Content lead review for all onboarding sequences. Critical for first impression and voice establishment.

Category	Examples	Tone Guidance	Review Required
Notification Copy	Push notifications, email updates, in-app alerts	Maintain the friend voice in all channels: "Your trip to Barcelona is coming up, here's what to expect" rather than "Upcoming booking reminder." Never use urgency or interruption language.	Content lead review for all notification templates. Must feel helpful, never intrusive or sales-driven.
Success States	Trip recommendations delivered, bookings completed, preferences captured	Reference the conversation that led to this moment: "Based on what you told me about needing proper downtime, here's what I found." Frame success as understanding achieved, not transaction completed.	Content lead review for all success flows. Must demonstrate visible connection between user input and outcome.
Help & Support Text	FAQ responses, contact forms, troubleshooting guides	Write as if the user is talking to someone who has context from their previous conversations. No corporate boilerplate or generic support language. "Having trouble with something? I'm here to help" rather than "Submit a ticket."	Content lead review for all support touchpoints. Must maintain relationship continuity and personal tone.

Review Process

All system copy must be reviewed and approved before implementation begins, not during QA testing. The content lead reviews for voice consistency and emotional appropriateness. The product owner reviews for strategic alignment and user journey coherence. Any system copy that fails review must be rewritten before development continues.

System copy changes during development require the same review process as original copy. Quick fixes or developer-written placeholder text cannot ship without content review, as these moments often have disproportionate impact on user trust and brand perception.

Critical Review Triggers

Immediate content review is required for any copy that appears during error states, first-time user interactions, or moments of user vulnerability such as payment processing. Copy that communicates system limitations, technical constraints, or service unavailability requires elevated review to ensure problems are framed as temporary setbacks rather than product failures.

Any copy that references AI capabilities, explains algorithmic decisions, or acknowledges when the system doesn't understand user input requires review for technical accuracy and voice consistency.

Acceptance Criteria

Define what "done" means for the build phase. These are specific, verifiable criteria that will be tested during QA and validated before release. Each criterion includes the verification method and the person responsible for sign-off.

Emotional Acceptance Criteria

The AI creates trust through conversation, not features, [] First screen shows conversation prompt, not destination options, [] User can complete initial preference conversation without seeing any booking options, [] AI asks minimum 3 contextual follow-up questions before showing first recommendation, [] No "powered by AI" or technology language appears anywhere in the interface, **Verification method:** User testing with 5 first-time users, measuring completion rate of initial conversation, **Sign-off required:** Product Owner + Content Lead

Users feel understood through specific, not generic personalisation, [] Every recommendation includes reasoning that quotes user's actual words, [] No template phrases like "based on your preferences" appear in suggestion explanations, [] AI references specific user constraints (group size, dates, budget, needs) in all communications, [] Recommendation reasoning reads like human conversation, not algorithmic output, **Verification method:** Content audit of all suggestion templates + user feedback on "felt personalised", **Sign-off required:** Content Lead + UX Lead

Moments of vulnerability are preceded by trust demonstration, [] Budget questions appear only after user has seen personalised recommendations, [] Payment information requested only after user has confirmed trip interest, [] No financial fields in initial onboarding sequence, [] Personal information requests include specific justification for why they improve recommendations, **Verification method:** User journey audit + drop-off analysis at each vulnerability point, **Sign-off required:** Product Owner + UX Lead

Users experience relief, not additional overwhelm, [] Maximum 3 primary options presented on any single screen, [] No feature matrices, comparison tables, or dense information displays, [] Loading states show AI working for user, not generic "processing" messages, [] Every interaction reduces decisions user needs to make independently, **Verification method:** Cognitive load testing + screen density audit, **Sign-off required:** UX Lead

Brand and Voice Criteria

Voice maintains conversational warmth without corporate inflation, [] All copy passes read-aloud test with development team (sounds natural when spoken), [] No sentences longer than 20 words in user-facing interface copy, [] No passive voice when AI is taking action ("Here's what I found" not "Results have been generated"), [] All copy uses "you" focus rather than "we" or "TravAI"

focus, **Verification method:** Copy audit against voice guidelines + read-aloud review session, **Sign-off required:** Content Lead

System copy maintains friend voice during problems, [] Error messages never blame user or use technical language, [] Empty states continue conversation rather than showing generic "no results", [] Loading messages show AI actively working: "Looking for places that match what you described", [] All system messages maintain helpful tone without manufactured enthusiasm, **Verification method:** Error state testing + system copy review against voice guidelines, **Sign-off required:** Content Lead + Product Owner

Visual hierarchy prioritises content over product features, [] AI interface elements never compete with travel content for visual attention, [] Trip imagery and details dominate screen real estate over system UI, [] Product branding stays subtle, never overshadowing destination content, [] Interface feels like browsing travel content with AI assistance, not using AI software, **Verification method:** Visual audit + user perception testing ("does this feel like travel or tech?"), **Sign-off required:** UX Lead + Creative Director

Interaction Acceptance Criteria

AI demonstrates guidance without coercion, [] When users explore alternatives, AI adapts rather than redirecting to original suggestions, [] No "back to recommendations" prompts when users deviate from suggested path, [] AI asks follow-up questions about user deviations instead of ignoring them, [] Users can change any previously stated preference without starting conversation over, **Verification method:** User journey testing with intentional deviation scenarios, **Sign-off required:** UX Lead + Product Owner

Conversation state persists across all user sessions, [] Returning users see reference to previous conversations in greeting, [] All stated preferences carry forward to new sessions without re-entry, [] AI remembers specific language user employed, not just categorical preferences, [] Previous reasoning visible in recommendation history across sessions, **Verification method:** Multi-session user testing + conversation state audit, **Sign-off required:** Technical Lead + Product Owner

Progressive revelation reduces cognitive load effectively, [] Destination selection completes before flight options appear, [] Hotel choices shown only after destination confirmation, [] Activities and experiences revealed only after accommodation selection, [] No step shows comprehensive view of all decisions simultaneously, **Verification method:** User flow testing + cognitive load measurement at each step, **Sign-off required:** UX Lead

AI reasoning remains transparent and accessible, [] Users can ask "why did you suggest this?" and receive specific, personal answer, [] All AI decisions reference visible user input or previously stated preferences, [] Reasoning explanations use user's own language when possible, [] No algorithmic black boxes in user-facing explanations, **Verification method:** Transparency testing + reasoning explanation review, **Sign-off required:** Product Owner + Content Lead

Performance Criteria Affecting Emotional Experience

Response times support conversation flow, [] AI responses to user questions appear within 3 seconds maximum, [] No loading state exceeds 8 seconds without explanation of delay, [] Conversation flow never interrupted by technical performance issues, [] Initial recommendation generation completes within 15 seconds of final preference input, **Verification method**: Performance testing under realistic load conditions, **Sign-off required**: Technical Lead

Conversation memory performs reliably, [] Previous conversation context loads within 2 seconds on return visits, [] No user reports losing stated preferences between sessions, [] AI references to previous conversations appear automatically, not on request, [] Conversation history accessible within 3 clicks from any screen, **Verification method**: Memory persistence testing + user experience validation, **Sign-off required**: Technical Lead + Product Owner

Search and recommendation quality meets personalisation expectations, [] First recommendation set includes at least one option that clearly matches stated preferences, [] AI can explain reasoning for every suggestion using specific user input, [] No generic or obviously irrelevant recommendations in first results, [] Alternative suggestions improve relevance when user provides feedback, **Verification method**: Recommendation quality testing + user satisfaction scoring, **Sign-off required**: Product Owner + Technical Lead

Copy and Content Criteria

All system copy passes voice review before implementation, [] Error messages reviewed and approved by Content Lead before development, [] Loading states use approved copy that maintains conversation voice, [] Confirmation messages reference specific user choices, not generic success language, [] Help text maintains friend voice, never corporate support language, **Verification method**: Complete copy audit against voice guidelines, **Sign-off required**: Content Lead

Onboarding establishes relationship immediately, [] First interaction is AI question, not feature explanation or app tour, [] Welcome message focuses on user benefit, not product capability, [] No mention of AI features or technical capabilities in opening sequence, [] User can begin preference conversation within 2 interactions of first screen, **Verification method**: First-time user experience testing + onboarding flow review, **Sign-off required**: Product Owner + Content Lead

Recommendation explanations feel personal and specific, [] Every suggestion includes "why this suits you" reasoning in user's language, [] No template explanations or generic personalisation claims, [] Reasoning references specific conversation details, not inferred preferences, [] Users can identify their own words reflected in explanation copy, **Verification method**: Personalisation testing + reasoning copy review, **Sign-off required**: Content Lead + UX Lead

Pre-Launch Emotional Review

The Pre-Launch Emotional Review is the final checkpoint before TravAI ships to users. Unlike technical QA that validates functionality, this review validates feeling. It asks the critical question: will someone who experiences this product actually feel understood and catered for, or will they recognise it as another booking platform with AI marketing? The review exists because emotional experience cannot be retrospectively designed into a product. It must be preserved throughout the build phase or rebuilt entirely.

What the review covers

The review walks through the complete user journey against each Feel Factor layer, checking that emotional intent has survived the development process intact. Reviewers experience the product as a first-time user would, paying particular attention to moments of vulnerability, choice, and AI interaction. Each critical moment identified in the design phase is tested for its ability to create the intended emotional response.

The conversation initiation sequence receives concentrated attention. Reviewers confirm that users encounter questions before options, that the AI demonstrates listening before presenting solutions, and that trust-building happens visibly rather than through claims. The first recommendation reveal is tested for specificity and connection to user input. Budget and payment flows are validated for their positioning after trust establishment rather than data collection.

The review also examines system responses to user deviation. Reviewers deliberately go off suggested paths to confirm the AI adapts rather than redirects, that exploration is encouraged rather than corrected, and that the conversation continues rather than resetting. Error states, empty states, and edge cases are tested for voice consistency and relationship maintenance.

Visual execution undergoes scrutiny for emotional impact. The density audit confirms that screens breathe rather than overwhelm, that travel content dominates over interface elements, and that the overall impression reads as guidance rather than transaction. Loading states, confirmation messages, and micro-interactions are evaluated for their contribution to the calm, unhurried experience that differentiates TravAI.

Who participates

The Design Lead facilitates the review and holds final authority on emotional experience quality. The Content Lead validates voice consistency across all user-facing text and system copy. The Product Owner confirms that strategic intent has been preserved and that compromise decisions during development have not undermined core differentiators.

Development leads participate to understand how technical implementation choices affect emotional outcomes and to identify any constraints that might require design adjustment rather than forcing

suboptimal user experience. QA leads join to understand what they should be testing for beyond functional compliance.

Client stakeholders from TravAI attend as observers and validators, confirming that the built experience aligns with their understanding of the brand promise and user needs. Their role is to validate strategic direction rather than to suggest tactical changes at this late stage.

What the review produces

The review generates a formal Emotional Experience Audit document that categorises findings into three levels: Experience Affirmed, Adjustment Required, and Launch Blocked. Each finding includes specific evidence, impact assessment, and recommended resolution.

Experience Affirmed findings confirm that emotional intent has been successfully implemented and that specific aspects of the experience deliver the intended user feeling. These become reference points for post-launch optimisation and expansion decisions.

Adjustment Required findings identify areas where emotional intent is present but execution needs refinement before launch. These typically involve copy revisions, interaction timing adjustments, or visual hierarchy modifications that can be addressed within the current development sprint without architectural changes.

Launch Blocked findings indicate fundamental breaks in the emotional experience that cannot be resolved through minor adjustments. These are rare but absolute. Examples might include technical constraints that force budget questions into early onboarding, AI reasoning that remains generic despite personalisation claims, or system responses that feel corporate despite voice guidelines. Launch Blocked findings require either significant rework or strategic decision to accept compromised user experience.

The audit document includes specific user quotes from testing sessions, screenshots of concerning interactions, and measurable benchmarks for post-launch monitoring. Each finding links directly to the original design principles and research insights that justify the required standard.

Sign-off process

Launch approval requires unanimous agreement from the Design Lead, Content Lead, and Product Owner that no Launch Blocked findings remain unresolved. Adjustment Required findings may be accepted for post-launch resolution if they do not fundamentally undermine the core emotional transformation that TravAI depends on.

The threshold for approval is not perfection but coherence. The experience must feel unified in its commitment to making users feel understood rather than processed. Individual elements may require improvement, but the overall emotional arc from overwhelmed arrival to confident booking must be intact and demonstrable.

Any stakeholder can block launch based on Emotional Experience Audit findings, but blocks must reference specific design principles or user research insights rather than personal preference or feature wishes. The audit document provides the framework for these decisions, ensuring they remain tied to user benefit rather than internal opinion.

Post-launch review framework

The emotional experience requires ongoing validation because user expectations evolve and competitive context shifts. Post-launch reviews follow a structured cadence designed to catch drift before it becomes entrenched.

The 30-day review focuses on first-impression metrics and onboarding experience quality. Real user conversations are analysed for trust-building effectiveness, abandonment patterns are mapped against emotional journey stages, and customer feedback is audited for language that suggests users feel processed rather than understood.

The 90-day review examines sustained usage patterns and relationship development. Users who have completed multiple bookings are interviewed about their evolving relationship with the AI, conversation quality is assessed for personalisation depth, and repeat usage behaviour is analysed for signals that the guidance relationship is strengthening or degrading.

The 6-month review evaluates strategic positioning and competitive differentiation. The brand voice is audited for consistency and distinctiveness, user language about TravAI is compared to language about competitors, and the emotional positioning is tested for continued relevance and effectiveness in the market context.

Each review cycle can trigger immediate adjustments to preserve emotional experience quality or inform longer-term strategic evolution of the brand and product approach.