Implementation Roadmap

Last Updated: 4 November 2026, 20:15

Prepared for Albemarle Excellence

Project Albemarle Excellence App

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

Contents

Introduction 3
Phase Overview 3
Phase 1: Emotional Foundation 4
Phase 2: Journey Architecture 9
Deliverables 11
Success Criteria 13
Key Risk 14
Phase 3: Identity Layer 14
Deliverables 11
Success Criteria 13
What Makes This Phase Difficult 19
Phase 4: Polish & Pre-Launch 19
Deliverables 11
Success Criteria 13
The Pre-Launch Emotional Review 24
Post-Launch 25
Dependencies & Risks 29

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

Introduction

This roadmap exists to answer one question before it needs to be asked in a meeting: in what order does this product get built, and why that order rather than any other? The sequence of delivery for a product like this one is not a project management preference. It is a design decision, and it carries the same weight as any decision made during the design phase itself. A warm surface built on a structure that does not yet work produces something that looks like Albemarle and behaves like a different company. That is precisely what the workshops were describing when they mapped the eleven stages of the client journey and found that every peak sits inside a studio and every valley sits in the digital space around it. This document is about closing that gap in the right order, so that each phase of the build earns the next.

Three principles govern the sequencing throughout. The first is that emotional dependencies come before feature additions. The product cannot create the feeling of being known before the systems that allow recognition to exist are solid. Delight built on a floor that does not hold is worse than no delight at all, because it creates an expectation the product will then fail to meet consistently. So the structural conditions for warmth, the authorship rules, the recognition logic, the cross-studio knowledge transfer, are built before the experiences that depend on them. The second principle is that structural decisions come before surface ones. Voice, motion, and the fine-grained polish of edge-case screens matter enormously, but they matter at the end, when the architecture beneath them is stable and trusted. Polishing a surface that may still change is effort spent twice. The third principle is that the highest-risk assumptions are tested earliest, specifically in live conditions rather than in a staging environment. The December Saturday scenario from the Anti-Principles session is not a theoretical edge case. It is the test of whether the authorship principle holds when the pressure to compromise it is highest. That test needs to happen in Phase 2, with real stylists and real schedules, before any of the identity-layer features that depend on stylist trust are built.

This document sits alongside the sprint plans and task assignments that live in the team's project management tooling, but it is not a replacement for them and does not attempt to be. What it defines is the order of things and the reasoning behind that order: the deliverables for each phase, the success criteria that must be met before the next phase begins, and the scope boundaries that protect each phase from the well-intentioned additions that tend to arrive once a team has momentum. When a decision comes up during build about whether something belongs in Phase 2 or Phase 3, or whether a feature can be simplified without losing what matters, this document is where that conversation should start. The sprint plan tells the team what to do this week. This document tells them why the weeks are arranged the way they are.

Phase Overview

The product is built in four phases. Each one creates the conditions the next requires, and the order is not interchangeable. What follows is the rationale for the sequence before the specifics of each

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

phase are defined.

The logic is simple enough to state in a sentence: warmth without structure is performance, and structure without warmth is a booking system. Phase 1 builds the structural conditions for recognition. Phase 2 makes recognition travel. Phase 3 makes it mean something to the person being recognised. Phase 4 ensures every surface and every edge case reflects what the core of the product has already earned. Remove any phase from this sequence, or compress it into another, and the phase after it is building on something that has not yet been tested in live conditions. That is the shape of every digital product that ships warm and reveals itself to be cold under pressure. This one is sequenced specifically to prevent that outcome.

Phase	Name	Duration	Focus
1	Emotional Foundation	10 weeks	Recognition before request. The product knows who it is talking to, speaks in a human voice, and never puts words in a person's mouth without their involvement
2	Journey Architecture	8 weeks	Knowledge travels. What the chair learns reaches the host. What Mayfair knows is waiting in Edinburgh. The post-visit note arrives in the stylist's own voice, or honestly from the studio
3	Identity Layer	8 weeks	The relationship has depth. The Albemarle story exists and is worth returning to. The client can see, shape and trust what the product holds about them
4	Surface and Pre-Launch	6 weeks	Every screen, every edge case, every piece of motion and copy reflects the standard the core build has established. Nothing ships that the warmth of the first three phases cannot support

Phase 1: Emotional Foundation

Phase 1 is first because nothing that comes after it works without it. The features built in later phases, the post-visit note, the Albemarle story, the life moment recording, the cross-studio recognition, all of them depend on a product that already knows who it is talking to and speaks to that person honestly. Build those features on a foundation that greets a client of eight years as a stranger, or that sends messages under Daniel's name that Daniel did not write, and the warmth is surface-level from the start. It looks like Albemarle and behaves like a different company, which is exactly the description the workshops kept returning to and the problem this entire project exists to solve.

The three principles governing the roadmap sequence are all present in this phase. Emotional dependencies come first: a returning client cannot feel expected until the product is structurally capable of recognising them, so recognition is the first thing built. Structural decisions come before surface ones: the voice rules, the authorship logic, and the host relationship architecture are

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

established here, before anything is layered on top of them. And the highest-risk assumptions are tested earliest: the authorship principle, the claim that a message signed by a person is written by that person, is not a rule that can be validated in a staging environment. It needs to be tested against real stylists, real schedules, and the real pressure of a busy week before the post-visit note workflow is built to depend on it in Phase 2.

Phase 1 does not produce a complete product. It produces a product that earns the right to be built further. By the end of it, a returning client opens the app and sees their host's face and their usual appointment before they have tapped anything. A new client meets a person before they meet a form. Every automated message comes plainly from Albemarle, never from someone who did not write it. And the data the product holds about a client is visible to that client, in language they would recognise, with control over what stays and what goes. These are the preconditions for everything else. They are not features. They are the floor.

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

Deliverables

Deliverable	Priority	Notes
Home screen: returning client view	Must have	Opens on the host's name and photograph, the suggested appointment based on the client's usual stylist, service and time, and a single primary action. The full service menu sits at least one deliberate step back. This screen is the product's first act and its most important one
Suggested appointment engine	Must have	Derives a returning client's suggestion from their visit history. Surfaces after the third visit, as agreed in the design principles. Before that threshold, offers a soft guided path rather than a suggestion that could be wrong. Must not surface a suggestion involving a stylist who has left or is no longer available at that studio
Host relationship screen	Must have	The client's named host, with photograph and a direct message thread. Visible from the home screen in one tap. Includes an honest statement of when the host typically replies. Available from every primary screen in the app, not only the home screen
New client onboarding: host introduction first	Must have	The first screen a new client sees contains the host's name, face and a short note in the host's own voice. Followed by one open question about what they are hoping for. Account creation is requested only at the point of booking confirmation, not before
Service descriptions in plain language	Must have	Every service on the menu has a plain-language sentence explaining what it involves, how long it takes, and an honest price range with a note on what affects the final figure. No service ships with only an internal name and a "from" price. The cancellation policy appears alongside the service, written as a host would explain it
Stylist profiles with colleague descriptions	Must have	Each stylist profile includes a short description written by a colleague rather than a marketing team, naming what they are especially good at and, where relevant, the hair types and textures they work with most. The grid of faces with nothing else behind them is replaced by this
Authorship logic and audit infrastructure	Must have	The backend logic that separates human-authored messages from system-generated ones, records author and signature separately for every outgoing message, and flags any mismatch. This is not a Phase 2 feature. The rule that a message signed by a named person requires that person's involvement must be enforced from the first message the product sends

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

Deliverable	Priority	Notes
Automated message voice: Albemarle attribution	Must have	Booking confirmations, reminders, receipts, and all system-generated communications are written in a warm, plain institutional voice attributed to Albemarle. None of them borrow a stylist's or host's name or imply a person composed them. This applies to every automated string in the product, including error states and empty states
Client-visible data view, first version	Must have	A view showing what the product holds about the client, in plain language they would recognise, reachable within two taps of the home screen. Editable. Name and pronoun changes propagate to all studio-facing systems within 24 hours. Photographs private by default. This version covers the core fields; the full life moment and preference layers are Phase 3
Pricing transparency throughout booking flow	Must have	Price ranges appear alongside service descriptions and again at booking confirmation. The language confirms that the stylist will agree the final price before beginning. No client reaches the confirmation screen without having seen a realistic indication of cost
Banned phrase audit and copy baseline	Must have	Every client-facing string in the product, including component labels, empty states, error messages, loading states, and notification templates, reviewed against the Voice and Tone Guidelines and the banned phrase list before Phase 1 ships. No placeholder copy, no component library defaults, no system language in any state a client can read
No-availability handling, baseline	Must have	When a stylist is unavailable, the product names when they are next free, offers to log the client's interest in a cancellation, and provides a direct path to the host message thread. The dead-end empty calendar is replaced in this phase. The personal colleague suggestion language comes in Phase 2
Onboarding price and policy presentation	Should have	Pricing and the cancellation policy are presented in the new client onboarding flow, in plain language, before the client reaches the confirmation step. Written as a host would explain them, not as small print
Cross-studio data infrastructure, groundwork	Should have	The database architecture and API layer that will allow client history, preferences and stylist notes to travel across studios is built in Phase 1, even if the host-facing surfaces that use it are Phase 2 deliverables. Building the pipes before the taps is the correct sequence
Lapsed client handling, baseline	Nice to have	The product identifies clients who have not visited in a defined period and surfaces this to the host as a quiet prompt, not an automated message to the client. What, if anything, the host does with that information is their decision. No "we've missed you" logic is built or implied

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

The following must be confirmed before Phase 1 is considered complete and Phase 2 begins. These are not aspirational targets. They are the floor that Phase 2 builds on.

- A returning client with three or more visits opens the app and sees their host's name and face alongside a suggested appointment without navigating any menu, on both iOS and Android, with a load time under two seconds on a standard mobile data connection.
- A new client reaches a service suggestion and an appointment confirmation without being asked to create an account. Account creation is requested only at the confirmation step.
- The authorship audit log records the author and the signature separately for every outgoing message sent during the final two weeks of Phase 1 testing. Zero mismatches between author and signature across that period.
- No automated message in the product, including booking confirmations, reminders, and error states, carries a named stylist's or host's signature. Confirmed by a full string audit against the authorship specification before launch.
- Every service on the menu displays a plain-language description, an honest price range, and a note on what affects the final figure. Confirmed by walkthrough of the full service menu against the copy specification.
- The client-visible data view is reachable within two taps of the home screen on both platforms, displays every personal field the product holds in plain language, and allows any field to be edited or removed without a manual step. Confirmed by walkthrough with five test users who have not read the documentation.
- Name and pronoun changes entered in the client-visible data view are reflected in all studio-facing systems within 24 hours. Confirmed by an automated test across at least three studio records before launch.
- No client-facing string contains any word or phrase from the banned phrase list established in the Voice and Tone Guidelines. Confirmed by a completed copy audit signed off by the nominated content reviewer before the phase ships.
- The no-availability state for any stylist offers a next-availability date, a path to register interest in a cancellation, and a direct link to the host message thread. Confirmed by walkthrough across at least four different unavailability scenarios on both platforms.
- The home screen, host relationship screen, new client onboarding flow, and service menu are reviewed against the ten-o'clock test established in the Design Principles Builder session. A tired person on the sofa at ten at night should be able to complete a booking without confusion, without reading any copy twice, and without needing to contact support. Confirmed by user testing with at least six participants before launch.

What Phase 1 Does Not Include

Defining the scope boundary of Phase 1 is as important as defining its deliverables. Well-intentioned additions during a build phase tend to arrive with good reasons and real costs. What follows is the agreed boundary, and anything outside it belongs in a later phase regardless of how natural it might feel to include it now.

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

- The post-visit note workflow is not a Phase 1 deliverable. The authorship logic that makes it trustworthy is built in Phase 1, but the draft-and-approve flow, the chair note prompts, and the morning-after stylist interface belong in Phase 2. Building the note workflow before the authorship infrastructure has been tested in live conditions would reverse the sequencing principle that governs the whole roadmap.
- The Albemarle story is not a Phase 1 deliverable. The client-visible data view built in this phase is its structural precursor, but the story presentation, the visit-by-visit record with stylist notes and optional photographs, is a Phase 3 feature that depends on Phase 2's post-visit note workflow producing entries worth showing.
- Cross-studio host surfaces are not Phase 1 deliverables. The data infrastructure is groundwork built in this phase. The host-facing screens that make a Mayfair client's history visible to an Edinburgh stylist before they arrive are Phase 2.
- Life moment recording is not a Phase 1 deliverable. The client-visible data view built here does not include the life moment field or the host-facing prompt to add one. That feature carries the most significant emotional risk in the product and belongs in Phase 3, after the trust established in Phases 1 and 2 has made it safe to introduce.
- The Albemarle story sharing feature, the invitation-only long-standing client feature, and the seasonal suggestion logic are all Phase 3 features and are not scoped here.
- Motion and micro-interaction polish, beyond what is required for the product to feel calm and unhurried on first use, belongs in Phase 4. Phase 1 ships something that works with clarity and warmth. It does not ship something that is fully finished at every edge.
- Detailed analytics beyond the success criteria measurements listed above are not in scope for Phase 1. The measurement infrastructure needed for Phase 2 and Phase 3 is specified in those phases. Setting up comprehensive analytics before the product is running in live conditions tends to produce data that reflects the build rather than the experience.

Phase 2: Journey Architecture

Phase 2 begins from a specific point of confidence. By the end of Phase 1, a returning client opens the app and is recognised before they say a word. A new client meets a person before they meet a form. The authorship rules are live, the voice is consistent, and the structural conditions for trust exist. What Phase 1 cannot yet do is make any of that travel. The warmth stays in the room where it was first created, which is exactly the description of Albemarle before this project began. Phase 2 is about closing that gap at every seam: studio to app, app to studio, Mayfair to Edinburgh, one host to the next, chair notes to the morning after.

The post-visit note is the central deliverable of this phase and deserves to be named plainly as such before the full deliverables table is set out. It is not simply a feature. It is the moment the product proves that the studio's warmth was real rather than situational. A note that arrives the same evening, written in the stylist's own voice, reflecting back what the client actually said in the chair, extends the experience past the pavement in a way nothing else in the product can replicate. It is also the feature most at risk during this phase, because the pressure to fill a gap with something

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

generated and sent invisibly will be felt for the first time here, in live conditions, on a busy Saturday. The December Saturday scenario exists in this document because that is when the authorship principle is genuinely tested, and Phase 2 ends when the team can demonstrate it held.

The other deliverables in this phase share the same underlying logic: knowledge that currently sits in one place needs to reach the people and screens that need it, before the client arrives or sends a message rather than after. Cross-studio recognition, rhythm-based host prompts, the no-availability colleague suggestion, the cancellation reply handling, all of them are expressions of the same structural commitment. The care belongs to Albemarle, not to one room.

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

Deliverables

Deliverable	Priority	Notes
Post-visit note: draft and approve workflow	Must have	After each appointment, the product generates a draft for the stylist built from their chair notes. The draft is clearly marked as a draft and cannot be sent without the stylist reading and amending it. The interface makes editing quick and skipping structurally impossible. The authorship audit log records the stylist's approval before any note is sent under their name
Post-visit note: honest holding message	Must have	When a stylist cannot write or approve a note on the day, the product sends a short plain message from the studio, not from the stylist. It names that the stylist's notes will follow and gives an honest timeframe. No name is borrowed. The message is warm, brief, and truthful
Chair note capture interface	Must have	A lightweight stylist-facing screen for recording observations during or immediately after an appointment. Fast to use, with structured prompts for the core fields, colour formula, length, texture, styling approach, and a free-text field for anything specific to this client and this visit. Notes entered here become the source material for the draft-and-approve flow
Cross-studio client history: host view	Must have	Before any appointment at any studio, the receiving host has access to the client's full visit record, including their usual stylist's notes, their colour formula, their preferences, and any stated details about their hair. This requires the data infrastructure groundwork built in Phase 1 to be surfaced in a host-facing interface for the first time
Cross-studio recognition: client experience	Must have	A client visiting a different studio is never asked for information the product already holds. Their history is waiting before they arrive. The receiving stylist knows the formula. The experience of being known must hold across all eleven studios, confirmed by testing the Edinburgh scenario with real client records before this phase ships
Rhythm-based host prompt	Must have	When a client is approaching the interval at which they typically visit, the product surfaces a quiet prompt to their host. Not an automated message to the client. A note visible only to the host, suggesting this client may be due, with one-tap access to their record and a suggested message draft for the host to edit and send in their own voice if they choose to

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

Deliverable	Priority	Notes
No-availability: colleague suggestion	Must have	The baseline no-availability handling from Phase 1 is extended. When a stylist is unavailable, the product suggests a named colleague with a short description written by the unavailable stylist themselves, or a close colleague if the stylist is unreachable. The language is personal rather than categorical. "Daniel is away that week. Priya is who he would suggest for a precision cut" rather than "other stylists available"
Cancellation reply handling	Must have	When a client cancels, the product sends a warm reply from the studio that leaves the door open without any rebook pressure. A simple path to rebook is present but not foregrounded. No guilt language, no promotional offer, no urgency. The reply is reviewed against the voice guidelines before it ships, and the banned phrase list is confirmed clear
Host change handover, baseline	Must have	When a client's host changes, through departure, leave, or reassignment, the client's full record transfers before any contact is made by the new host. The process is not manual. The new host has access to the client's history, their notes, their rhythm, and any relevant context before they send a first message. The client experiences no gap in recognition
Post-visit note: Albemarle story first entry	Must have	Each approved post-visit note creates the first or next entry in the client's Albemarle story, which is built in full in Phase 3. Phase 2's responsibility is ensuring the entry exists, is written to the client rather than about them, and is accurate. The story presentation layer is out of scope here; the content layer is not
Lapsed client: host-visible flag	Should have	The product identifies clients who have not visited in a defined period and surfaces this quietly to their host as a flag in the client record view. No automated message is sent to the client. What the host does with the information, if anything, is their decision. The product does not prompt the host to contact the client, only to be aware
Authorship breach alert	Should have	An automated alert fires immediately if any outgoing message is flagged by the authorship audit as carrying a named person's signature without a confirmed approval record. This is not a weekly report. It fires in real time and routes to the product lead and the named operations contact. One breach is treated as a protocol failure requiring a specific fix, not a rounding error in a dashboard
Pilot studio selection and rollout plan	Should have	Phase 2 features involving real stylist workflows, specifically the chair note capture interface and the draft-and-approve flow, should be piloted at two or three studios before full rollout. The pilot selection should include at least one high-volume studio, so the December Saturday conditions are approximated before full launch

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

Deliverable	Priority	Notes
Phase 2 measurement instrumentation	Should have	The tracking infrastructure needed to report against the success criteria below is built and confirmed working before any Phase 2 feature reaches production. The metrics that matter here are operational, repeat information requests, stylist note completion rates, client reply rates to notes, and they need to be live from the first day of the pilot

Success Criteria

The following must be confirmed before Phase 2 is considered complete and Phase 3 begins.

A client visiting any Albemarle studio for the first time under their account is never asked for information the product already holds. Confirmed by testing the cross-studio scenario with at least five client records across at least three studio pairings, including at minimum one pairing between cities.

The authorship audit log records no mismatch between author and signature across the full duration of the Phase 2 pilot period. If a breach occurs during the pilot, the protocol failure is resolved and the pilot period restarts from zero. Phase 3 does not begin with an unresolved breach in the audit history.

Stylist note completion rates, defined as approved post-visit notes sent within 24 hours of an appointment, reach a minimum of 70% across pilot studios within four weeks of the chair note interface launching. This figure is a floor, not a target. It represents the minimum viable signal that the draft-and-approve workflow is making honest note-writing easier rather than heavier.

Client reply rates to post-visit notes, measured across the pilot studios, show any upward trend in the four weeks following the Phase 2 note workflow going live. A reply is the most direct available signal that the note landed as a message from a person rather than a system confirmation. Flat or declining reply rates indicate the notes are not reading as genuine, which requires qualitative investigation before Phase 3 begins.

The holding message, sent when a stylist cannot approve a note on the day, is confirmed to never carry a named stylist's or host's signature. Confirmed by a full audit of every holding message sent during the pilot period.

No cancellation reply sent during Phase 2 contains any word or phrase from the banned phrase list, any rebook prompt, or any promotional offer. Confirmed by a complete review of all cancellation replies sent during the pilot period against the voice specification.

The rhythm-based host prompt surfaces correctly for clients with visit histories matching the trigger criteria, and does not surface for clients who have not yet established a discernible rhythm. Confirmed by testing with client records across a range of visit histories before the feature goes live.

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

records is confirmed complete before any client contact is made. The client's experience of recognition is confirmed unchanged through the transition.

The no-availability colleague suggestion surfaces a named colleague with a personal description in every unavailability scenario during the pilot period. Zero instances of the dead-end empty calendar state remaining as the final screen. Confirmed by walkthrough across at least six unavailability scenarios on both platforms before launch.

Key Risk

The largest single risk in Phase 2 is the one the Anti-Principles exercise was designed to surface: the authorship principle does not hold under real operational pressure. The December Saturday scenario is not theoretical. It describes the conditions that will occur within the first month of this phase being live, and it describes them accurately. A high-volume studio, a tired stylist, six notes still unsent at seven in the evening, and a product that is technically capable of filling that gap perfectly and invisibly. The temptation is not abstract. It is specific, it is well-intentioned, and it arrives at the precise moment the right decision is hardest to make.

The mitigation is not a policy document or a training session, though both are necessary. It is design. The honest path, sending a holding message from the studio while Daniel approves drafts the following morning, must be faster and easier to execute than the dishonest one. If approving a draft requires more steps than sending it directly, the workflow will be bypassed under pressure. If the holding message template is harder to find than the name-signed alternative, it will not be used. The interface must be designed so that the correct behaviour at seven o'clock on a December Saturday is the path of least resistance, not the path that requires discipline to choose.

Alongside the design mitigation, the authorship breach alert described in the deliverables table is not optional. A weekly audit that surfaces a breach five days after it happened is not a safeguard. A real-time alert that routes to the product lead and the operations contact the moment a mismatch is flagged is. The alert exists so that if the principle is compromised, it is known immediately and treated as a protocol failure requiring a specific fix, before the client who received the message has had time to mention it at their next visit and notice the blank look on the stylist's face.

The pilot studio selection is the third element of the mitigation. At least one high-volume studio should be in the pilot group specifically so the December Saturday conditions are approximated before the feature reaches full rollout. A pilot that only includes quieter studios will not test the authorship principle under meaningful pressure, which means the pressure test happens at full scale instead of in a controlled environment where a breach can be addressed without broad client impact.

Phase 3: Identity Layer

Phase 3 is the phase that cannot be rushed to, and cannot be skipped past. The features it contains, the Albemarle story, life moment recording, the full client data view, the depth of personalisation that

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

makes a long-standing relationship feel genuinely different from a new one, are the ones that most visibly express what Albemarle means. They are also the ones most likely to cause harm if they arrive before the trust that makes them safe. A product that shows a client their hair history without first having earned their confidence in how that history is held is not offering them something beautiful. It is showing them a file. The order of these phases is not a planning preference. It is a safeguard, and Phase 3 exists on the other side of eight weeks of earning the right to offer it.

By the time Phase 3 begins, three things should be true. The product recognises returning clients before they say anything. Post-visit notes are arriving in stylists' own voices, with an authorship audit that has held clean across the Phase 2 pilot period. And the cross-studio handover works, confirmed by the Edinburgh scenario running against real client records. Only when all three of those things are true does the identity layer make sense to build, because the identity layer is not a set of features. It is a deepening of a relationship that the first two phases have already established. Add depth to a relationship that does not yet exist, and you get something that feels presumptuous. Add it to one that has been carefully built, and you get the thing the workshops described: a client who opens their story and wants to show it to someone.

The hardest thing about Phase 3 is that it requires the most careful design work and the most disciplined restraint at the same time. The life moment feature, done well, is what makes a client feel genuinely known in the way a good host knows a guest. Done carelessly, it is the feature that mentions a wedding that was cancelled, or uses a partner's name that no longer fits, or asks after something the client shared in a tender moment and has since moved on from. The Aspiration Gap session named these emotional landmines with precision, and they cluster almost entirely in the territory this phase is working in. That is not a reason to avoid the territory. It is a reason to move through it with more care than any other phase has required.

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

Deliverables

Deliverable	Priority	Notes
Albemarle story: full client view	Must have	A beautifully presented, chronological record of every visit, drawing on the post-visit note entries created in Phase 2. Each entry shows what was done, what the client was going for, any photograph they have chosen to add, and a line from their stylist. The presentation is generous with space, easy to scroll, and designed to feel worth returning to, not just worth consulting before an appointment
Albemarle story: photograph layer	Must have	Clients can add a photograph to any visit entry. Private by default. Never shared without explicit, informed choice. Easy to remove. The camera framing guidance in the design system applies: close, natural-light, honest rather than produced. No before-and-after framing in any adjacent copy or interface text
Client data view: full version	Must have	The baseline data view from Phase 1 is extended to include preferences, stylist notes written to the client, and the life moment layer when it launches. Every field is in plain language the client would recognise. Every field is editable. The view explains, briefly and in a human voice, what each piece of information is used for and why it is held. This is not a privacy policy. It is a conversation
Life moment recording: host-facing	Must have	A host can add a life moment to a client's record, an event, a transition, something the client has shared that is worth remembering for the right reason. This requires the client's awareness. The prompts make this explicit, and the interface does not allow a life moment to be added silently. Notes are written as the client would recognise them, not as internal shorthand
Life moment accuracy check	Must have	Before any life moment appears in a host-facing communication prompt or a suggested outreach, the product asks the person using it to confirm the detail is still accurate and still appropriate to raise. This is a required step, not an optional one. It applies every time a life moment is surfaced, not only the first time
Preference layer	Must have	Clients can record and edit their own preferences: how they like to be contacted, what they prefer not to be asked, specific hair and service notes that they want to carry across all studios and all hosts. These are distinct from stylist-recorded notes and belong to the client. The interface makes the distinction clear

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

Deliverable	Priority	Notes
Long-standing client invitation logic	Must have	For clients who have passed a defined visit threshold, the product surfaces a prompt to their host suggesting a personal invitation, a first look at a new studio, an evening with the creative director, a seasonal appointment the host has thought through for them. The prompt goes to the host. The invitation is written and sent by the host. No automated loyalty messaging is involved at any point
Suggested appointment refinement	Should have	The rhythm-based suggestion engine built in Phase 1 is refined using the fuller history now available. Where a client's rhythm has shifted, the suggestion reflects the new pattern. Where they have expressed a preference for a particular time or day, that is foregrounded. The refinement is quiet, not announced
Albemarle story: shareable card	Should have	A client can choose to share a single story entry as a card, their colour, their stylist, a line about what they were going for. Private, clean, easy to send to a friend asking for a recommendation. This is not a referral mechanic. It is not incentivised. It is an expression of the client's own story, in a format a friend would find useful. The share option sits inside the story, available when the client wants it, never promoted
Lapsed client: personal host outreach	Should have	When a client has not visited in a defined period, the product offers the host a draft outreach message, written warmly and without any guilt or pressure, to send in their own voice. The host reads it, edits it, and sends it, or decides not to. The product never sends this message automatically. If the host does not act on the prompt within a defined window, the prompt closes quietly. No second prompt is generated
Story navigation and filtering	Nice to have	Clients can filter their story by time period, stylist, or studio. The filtering is light and does not impose a data-management feeling on what should feel like a personal record. Available if the history is long enough to benefit from it
Anniversary and milestone prompt	Nice to have	When a client reaches a significant milestone with Albemarle, one year, five years, their tenth visit, the product surfaces a quiet prompt to their host. The host decides whether and how to mark it. No automated message is sent to the client. The milestone is an occasion for the host to do something genuinely personal, not an occasion for the system to congratulate someone

Success Criteria

The following must be confirmed before Phase 3 is considered complete and Phase 4 begins. Several of the most important measures here are qualitative, because the identity layer is the part of the product that is hardest to validate with numbers and easiest to get wrong while the numbers look fine.

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

The Albemarle story view is opened by at least 60% of clients with three or more visits within four weeks of the feature launching. Opening rate alone is not sufficient evidence of success, but a significantly lower rate would indicate the story is not yet worth returning to, which requires qualitative investigation before Phase 4 begins.

At least one life moment added by a host is confirmed accurate and confirmed appropriate to raise, in every case where it has been used in a host communication during the Phase 3 period. Confirmed by a full review of every life moment used in outreach during the period, cross-referenced with the accuracy check audit log. Any instance where a life moment was raised without passing the accuracy check is treated as a protocol failure requiring a specific fix before the feature reaches full rollout.

No photograph added to the Albemarle story is set to visible or shareable by default. Confirmed by automated test across every photograph upload in the Phase 3 period. Any photograph that defaulted to visible without explicit client action is treated as a critical defect and resolved before Phase 4 begins.

The preference layer is used by at least 25% of clients who visit the client data view during the Phase 3 period. This is a proxy measure for whether the control features feel genuinely available rather than technically present but practically buried. A significantly lower rate suggests the preference layer is not easy enough to find or use, which requires a design review before Phase 4.

No long-standing client invitation is sent automatically without host authorship. Confirmed by the authorship audit log for the Phase 3 period. Every invitation sent under a host's name has a confirmed approval record. Zero exceptions.

The lapsed client outreach, in every instance where the host chose to send it, is confirmed to contain no guilt language, no promotional offer, and no urgency framing. Confirmed by a review of all lapsed client messages sent during the Phase 3 period against the voice specification and the banned phrase list.

Qualitative review sessions are conducted with at least eight clients who have used the Albemarle story view for a minimum of six weeks. The review is not a usability test. It is a conversation about how the story makes them feel about their relationship with Albemarle, whether they have shown it to anyone, and whether it has changed how they think about booking their next visit. The findings from these conversations are reviewed by the product lead and the Albemarle client team before Phase 4 is confirmed ready to begin. If the conversations reveal that the story feels like a database rather than a record, or that life moments feel exposing rather than considered, those findings require a design response before the surface polish of Phase 4 begins.

The client data view full version is confirmed to show every field the product holds about a client, with no hidden fields, no internal shorthand, and no data accessible to Albemarle that is not accessible to the client. Confirmed by a data audit conducted jointly by the development team and the Albemarle client team before the phase ships.

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

What Makes This Phase Difficult

The difficulty of Phase 3 is not technical. The engineering work here is less complex than the cross-studio data infrastructure of Phase 2. The difficulty is that this phase asks the product to hold something genuinely personal, and there is very little margin for error in how it does so. The features that land well in Phase 3 are the ones that make a client feel known in exactly the way a trusted host knows a guest. The ones that land badly are the ones that make a client feel watched, or managed, or quietly catalogued by a system that remembered the wrong thing at the wrong time.

The gap between those two outcomes is not a matter of feature design. It is a matter of when each feature is introduced, by whom, in what context, and with what degree of control offered to the client. Life moments added without the client's awareness tip from warmth into surveillance. A story that surfaces a painful period rather than a happy one, because the product did not know the period was painful, tips from a beautiful record into something the client would rather not look at. An invitation that arrives automatically, however warmly worded, tips from personal attention into a loyalty mechanic wearing a different coat. Every one of those failures was described in precise detail during the design phase. The Anti-Principles exercise, the emotional landmines session, the Funeral Speech, all of them circled this territory and named what goes wrong in it.

What makes Phase 3 difficult is that the solutions to those failure modes are not additional features. They are disciplines: the accuracy check that must run before any life moment is raised, the host authorship requirement that must hold for every invitation, the client's control over their own story that must be real and not merely theoretical. Those disciplines are easier to maintain when the team holds the design phase reasoning in view. They are harder to maintain at the end of a long build cycle, when features are being finished under time pressure and the temptation is to treat the safeguards as edge-case handling rather than structural requirements.

The other difficulty is calibration. The features in Phase 3 should feel different depending on how long a client has been with Albemarle. A client on their fourth visit is not ready to receive an invitation or to find a richly annotated life moment in their record. A client of five years should find those things feeling natural, even expected. The product does not have a simple toggle that moves between those states. It has the history of the relationship, surfaced gradually through the first two phases, and Phase 3 is where that history becomes visible. Getting that calibration right requires the team to think about each feature not as a capability that is either on or off, but as something that earns its place through the depth of the relationship it reflects. That is a harder design problem than building the feature itself, and it is the one this phase turns on.

Phase 4: Polish & Pre-Launch

Phase 4 is not about adding things. The core build is done. The recognition works, the authorship holds, the story exists and means something to the people who have used it. What this phase is about is ensuring that every part of the product, every edge case, every loading state, every screen a

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

client lands on when something unexpected has happened, reflects the same standard as the parts that were designed most carefully. The gap between a thoughtfully designed primary flow and a component-library-default error message is the gap between a product that feels like Albemarle and one that feels like Albemarle most of the time. Phase 4 closes it.

The discipline required here is different from the previous three phases. Phases 1 through 3 were about building something that did not yet exist. Phase 4 is about looking at what exists and being honest about where it falls short, then fixing it before anyone outside the team sees it. That requires the team to hold the same standard at the end of a long build cycle as they held at the beginning, which is harder than it sounds, and worth naming plainly as the main challenge of this phase.

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

Deliverables

Deliverable	Priority	Notes
Performance audit: load times across all primary flows	Must have	Every primary screen is measured against the two-second load time standard established in Phase 1. The home screen, appointment suggestion, host messaging thread, and Albemarle story view are the highest-priority targets. Testing must include mid-range Android devices and standard mobile data connections, not only the ideal conditions of a development machine on Wi-Fi. Any screen that fails to meet the standard is returned to engineering before the launch gate is opened
Performance audit: animation and transition smoothness	Must have	Every transition in the product is reviewed for dropped frames, jarring motion, and any animation that feels effortful rather than calm. The product's register is unhurried. Motion that feels rushed, stuttered, or mechanically abrupt contradicts that register at a level below conscious reading. Reviewed on both platforms, on the same mid-range Android device used for load time testing
Full copy review against Voice and Tone Guidelines	Must have	Every client-facing string in the product, without exception, is reviewed against the Voice and Tone Guidelines and the banned phrase list. This includes component labels, placeholder text, empty states, error messages, loading states, confirmation screens, notification templates, and any copy inside the Albemarle story view. The review is conducted by the nominated content reviewer and signed off before any launch gate is opened. No string is exempt on grounds of technical necessity or low visibility
Emotional review: full journey walkthrough	Must have	A structured walkthrough of the complete product, conducted from the perspective of three distinct client types: a new client on their first visit, a returning client of one year, and a long-standing client of four or more years. For each client type, the reviewer moves through every stage of the journey, from first open through booking, waiting, the post-visit note, and the Albemarle story, and assesses whether the emotional quality at each stage matches the arc defined in the Feel Factor Framework. Findings are documented and any stage that falls short of the standard is addressed before launch
Accessibility review: WCAG 2.1 AA compliance	Must have	A full accessibility audit across both platforms against WCAG 2.1 AA. Focus areas: colour contrast across all states, including pressed and disabled, touch target sizing across all interactive elements, screen reader compatibility for the host messaging thread and the Albemarle story view, text scaling behaviour at the largest system font sizes. Any failure against AA standard is a blocking issue

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

Deliverable	Priority	Notes
Regression testing: Phase 1 through 3 success criteria	Must have	Every success criterion confirmed at the end of each prior phase is retested before launch. The authorship audit is re-run across a fresh period of live use. Cross-studio recognition is confirmed against the same studio pairings used in Phase 2. The client data view accuracy check is confirmed against updated records. Regression testing is not a formality. It is the confirmation that nothing built in Phase 3 or 4 has introduced a failure in something Phases 1 or 2 established
Edge case review: emotional landmines	Must have	A structured review of every scenario identified in the Aspiration Gap and Emotional Journey sessions as a potential source of harm. This includes no-availability states, cancellation handling, the lapsed client outreach flow, the life moment accuracy check, name and pronoun change propagation, and any scenario involving a photograph or a sensitive personal detail. Each scenario is tested against the standard established in the design principles, and any failure is a blocking issue regardless of how edge-case the scenario appears
App store preparation: iOS and Android	Must have	Store listings, screenshots, and preview content prepared for both platforms. Screenshots must reflect the actual product, not a marketing idealisation of it. Any copy in the store listing is reviewed against the voice guidelines before submission. The listing does not use any word from the banned phrase list. Preview content is natural-light photography consistent with the visual direction
App store review submission and approval	Must have	Both platforms submitted and approved before the launch date is confirmed. Apple's review process carries the most schedule risk and should be submitted with enough lead time to absorb a rejection and resubmission without pushing the launch date
Notification timing and frequency review	Should have	Every notification type in the product is reviewed for timing, frequency, and the conditions under which it fires. The review confirms that no notification is sent in a way that violates the "only get in touch to help" principle. The lapsed client prompt, the rhythm-based host prompt, and the pre-visit reminder are each confirmed to fire only under the conditions defined in their respective phase specifications. No notification arrives more than once for the same trigger
Copy review: component library defaults	Should have	A specific pass through the product targeting any string that may have been inherited from the component library or left as a development placeholder. This is the category most likely to contain system language that passed through earlier reviews without being caught. "No results found", "loading", "error", "tap to continue" are the patterns to look for. Every instance is replaced with copy that meets the voice standard

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

Deliverable	Priority	Notes
Visual system compliance check	Should have	A systematic review of every screen against the design system: colour usage, type scale, spacing tokens, photography style, and the brass accent rule. Any component or screen that has drifted from the visual specification during build is documented and either corrected or escalated before launch. The review is conducted by a designer with direct access to the design system documentation
Soft launch plan: limited studio rollout	Should have	Before full launch across all eleven studios, the complete product is released to a defined subset of studios and clients. The selection criteria are the same as Phase 2's pilot logic: at least one high-volume studio, at least one cross-city pairing, a sample of returning clients across a range of visit histories. The soft launch period is a minimum of two weeks. Findings from the soft launch inform any final fixes before full rollout
Launch communications: host and stylist briefing	Should have	Every host and stylist at every studio receives a clear briefing on how the product works, what has changed, and what their responsibilities are. The briefing is specific about the authorship rules, the life moment process, and the host prompt logic. It is written in a voice consistent with the brand: warm, plain, honest about what the product can and cannot do. No briefing material uses any word from the banned phrase list
Post-launch monitoring plan	Nice to have	A defined plan for the first four weeks post-launch, covering which metrics are reviewed daily, which are reviewed weekly, and what thresholds trigger an escalation. The plan is documented before launch and assigned to named individuals. The authorship audit is on the daily list

Success Criteria

The following must all be confirmed before the product launches. These are the launch gates, and they do not have exceptions.

Every primary screen, defined as the home screen, the appointment suggestion and booking flow, the host messaging thread, the Albemarle story view, and the client data view, loads within two seconds on a mid-range Android device on a standard mobile data connection. Confirmed by measurement across a minimum of three device types before the gate is opened.

No client-facing string in the shipped product contains any word or phrase from the banned phrase list established in the Voice and Tone Guidelines. Confirmed by a completed copy audit, signed off by the nominated content reviewer, covering every string on both platforms including component labels, error messages, empty states, loading states, and notification templates. The sign-off is a named individual, not a team, and it is on record before the launch date is confirmed.

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

the audit that cannot be resolved before launch is a reason to delay the launch, not to reclassify the failure.

Every success criterion from Phases 1, 2, and 3 is confirmed still met after the Phase 4 build. The authorship audit shows no mismatches across a fresh two-week period of live use. Cross-studio recognition is confirmed against at least three studio pairings. Name and pronoun changes are confirmed to propagate within 24 hours. The client data view shows no hidden fields. These confirmations are documented and signed off before the launch gate is opened.

The emotional review walkthrough, conducted for all three client types, produces no finding rated as failing against the Feel Factor Framework. Any finding rated as falling short of the framework standard is addressed and re-reviewed before the launch gate is opened.

Both app store listings are approved by their respective platforms before the launch date is confirmed. The launch date is set only after approval is in hand, not in anticipation of it.

The soft launch period produces no finding that changes the launch readiness assessment. If the soft launch surfaces a failure against any of the above criteria, the finding is resolved and the relevant criterion is re-confirmed before full rollout proceeds.

The Pre-Launch Emotional Review

The pre-launch emotional review is not a usability test. It is not a QA pass. It is a structured attempt to experience the product as a client would, before those clients are real people whose trust the product will need to earn. It is conducted by people who understand the design intent, have read the Feel Factor Framework and the Design Principles, and are willing to say honestly when something falls short.

The review is organised around three client archetypes, each representing a meaningfully different relationship with the product and a different emotional starting point.

The first archetype is a new client. Someone who has heard about Albemarle from a friend, whose expectations have been set by a warm personal description, and who is opening the app for the first time on a quiet Tuesday evening. The reviewer enters the product as this person and moves through every screen up to and including the booking confirmation and the pre-visit message. The question at each stage is not whether the screen works. It is whether this person still feels welcomed, still feels that the place they were described to is the place they have found, and whether anything in the experience has asked them to do something that a good host would have done on their behalf.

The second archetype is a returning client of approximately one year. Someone who has visited four or five times, who knows their host by name, and who opens the app to rebook something before a work event next week. The reviewer enters the product at the home screen, as a client whose history is populated, and moves through the suggested appointment flow, the host messaging thread, and their Albemarle story. The question at each stage is whether everything that has been learned about this person across a year of visits is quietly and correctly present, whether the experience requires

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

any effort from someone who should not have to introduce themselves again, and whether the warmth of the product at this stage feels earned or performed.

The third archetype is a long-standing client of four or more years. Someone for whom the Albemarle relationship has depth, whose story is rich, who has a life moment in their record, and who has been to two studios. The reviewer enters the product at the Albemarle story view and moves through the full depth of what Phase 3 has built. The question here is whether the experience of being known at this depth feels good, whether the life moment is handled with the care the design principles require, and whether the client data view, showing everything the product holds about this person, feels like a gift or a file. The reviewer also moves through a cross-studio booking scenario and confirms that the Edinburgh visit feels like coming home rather than starting over.

The findings from each archetype walkthrough are documented against the specific stages of the Emotional Journey Map and the three layers of the Feel Factor Framework. The documentation names each stage, the emotional quality experienced, the standard it should meet, and whether it meets it. Where it does not, the finding is specific rather than general: not "the booking flow feels cold" but "the confirmation screen for a first booking contains no host name and closes with a secondary CTA, which contradicts the principle that care is shown rather than claimed."

Findings are presented to the product lead and the Albemarle client team jointly. Any finding that the team rates as falling below the Feel Factor standard is a reason to delay the launch gate. The purpose of presenting findings jointly is to ensure that the decision about whether a finding is a launch blocker is made by people who understand both the technical cost of fixing it and the experiential cost of shipping without fixing it. That conversation belongs before launch, not after.

The review is conducted a minimum of two weeks before the intended launch date. That window exists so that any finding that is a launch blocker has time to be addressed, re-reviewed, and confirmed resolved before the gate is opened. A review conducted in the final 48 hours before launch is not a pre-launch emotional review. It is a post-build rationalisation, and it should not be confused with one.

Post-Launch

The launch is not the end of the sequencing. It is the beginning of a different kind of work, one that is less about building and more about listening. The product that ships at the end of Phase 4 has been designed carefully and tested against the Feel Factor Framework, but it has only ever been tested in conditions the team could control. What happens in the first thirty days of real use, with real clients across all eleven studios, is the first honest signal the product will produce. Some of what it reveals will be expected. Some of it will not, and the reviews defined below exist to catch both.

The cadence is deliberate. Each review has a different purpose, a different set of questions, and a different relationship to the delivery cycle that follows it. None of them are retrospectives in the conventional sense, a chance to reflect and move on. Each one feeds forward. The findings from the

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

30-day review shape the first post-launch sprint. The findings from the 90-day review confirm whether the core principles are holding under real operational conditions. The six-month review is the first moment the team can honestly assess whether the emotional arc is working at scale. The annual review is where the product is measured not against its own specifications but against the standard it was always meant to reach: a client who opens the app and feels, without thinking about it, that someone was already looking after them.

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

Review Cadence

Review	Timing	Output
30-day post-launch review	30 days after full rollout across all studios	Operational stability report covering load times, error rates, and authorship audit findings across the full period. Qualitative findings from host and stylist feedback sessions at each studio. A ranked list of friction points identified by hosts, stylists, and clients, prioritised by emotional impact rather than frequency. Specific fixes for any Phase 4 success criterion found not to be holding in live conditions. A confirmed re-baseline of the authorship audit.
90-day post-launch review	90 days after full rollout	The first proper signal on the five measurable success criteria from the Design Principles Builder: repeat information requests per studio, authorship audit status, banned phrase audit, notification mute and opt-out rates, and client-initiated profile additions. Qualitative review of post-visit note reply rates per studio. Assessment of whether the rhythm-based host prompt is firing correctly and being acted on. First look at cross-studio usage patterns. Any finding that indicates a principle is not holding at scale is treated as a design problem requiring a response, not an operational variance to monitor.
6-month review	6 months after full rollout	The first quantitative assessment of the emotional arc metrics: the proportion of returning clients who complete a booking without navigating the full service menu, the proportion of clients who have used the Albemarle story view for something other than pre-appointment reference, the proportion who have initiated a conversation with their host outside a booking context. A structured qualitative review with clients across at least three studios, using the same three-archetype format as the pre-launch emotional review, to assess whether the experience of being known has deepened over time or plateaued. Host and stylist retention data reviewed for any signal that the product is making their work better or harder. Life moment accuracy check audit reviewed for the full period.
Annual review	12 months after full rollout	Full reassessment of the product against the Feel Factor Framework. The emotional review is conducted from scratch, as if for a new product, with no inherited assumptions about what works. The five measurable signals from the Design Principles Builder are presented as year-one trends. The authorship audit is reviewed in full for the year. The Albemarle story engagement data is reviewed for depth, not just open rates: how many clients have added a photograph, how many have scrolled back more than three visits, how many have shared an entry. The cross-studio usage data is reviewed against the vision of a care that belongs to Albemarle as a whole rather than to one room. The findings are presented to the Albemarle leadership team alongside a recommendation for the following year's delivery priorities.

How Each Review Feeds Forward

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

The 30-day review produces a small number of specific, high-priority fixes and a confirmed operational baseline. Its output is a focused sprint, not a roadmap revision. The team knows at this point what is not working as intended in live conditions, and the sprint is the response to that. The sprint is scoped tightly: only the fixes identified in the 30-day review, plus any authorship breach remediation. Nothing is added.

The 90-day review is the first moment the team can distinguish between things that are not working and things that are working but not yet visible in the data. The difference matters. A host prompt that is firing correctly but being ignored by hosts is a training and communication problem. A host prompt that is surfacing at the wrong rhythm is a product problem. The 90-day review is where those two things are separated, and the response is calibrated accordingly. Its output is a prioritised backlog of improvements, each one traced to a specific finding, with a clear statement of which principle it is intended to hold or restore.

The 6-month review is where the product is first measured against its own ambition rather than its own specifications. The specifications say what the product should do. The ambition, expressed in the Aspiration Gap, the Funeral Speech, and the Feel Factor Framework, describes how clients should feel as a result of a sustained relationship with the product over time. The 6-month review is the first honest opportunity to assess whether that is happening. Its output is a recommendation to the Albemarle team on what the second year of the product's life should focus on. That recommendation may confirm that the current direction is right and the team should continue building depth in the same areas. It may reveal that something in the core design needs revisiting. Either outcome is a valid finding, and the review should be conducted without a presumption of which one it will be.

The annual review is the most consequential of the four, and the one that requires the most independence from the people who built the product. By twelve months post-launch, the team will have lived with the product long enough to have developed assumptions about what it does and does not do. The annual review should be designed to surface what those assumptions are hiding. The three-archetype walkthrough format, used fresh rather than against the memory of the last walkthrough, is what provides that distance. The findings from the annual review are not a list of improvements. They are a reckoning with whether the product is doing what it was designed to do, at the scale it is now operating at, for the people who are now using it. That reckoning is what makes the annual review the input to the following year's strategy rather than a summary of the year just completed.

One thing that should not change between any of these reviews and the work that follows them: the sequencing principles that governed the original build. Emotional dependencies come first. Structural decisions come before surface ones. The highest-risk assumptions are tested before the features that depend on them are added. Those principles do not expire at launch. They apply to every subsequent delivery cycle, and the review cadence exists, in part, to ensure that the team and the Albemarle leadership continue to reason from them when new priorities arrive.

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

Dependencies & Risks

The build sequence defined in this roadmap rests on a set of assumptions. Some of them are technical. Some of them are about how people behave under pressure. All of them carry consequences if they turn out to be wrong, and the right time to name them is before the build begins rather than when a sprint is running late and a shortcut is becoming attractive.

What follows is a risk register for the full delivery programme. It is not exhaustive, because a truly exhaustive risk register is a document nobody reads. It covers the risks that are most likely to arrive, most likely to be rationalised away under pressure, and most likely to damage something the design phase worked hard to protect. For each one, the mitigation is specific. A mitigation that says "monitor closely" or "raise in the next review" is not a mitigation. It is a note that the risk exists. The mitigations here name what changes structurally, who is responsible, and what the early warning looks like.

Risk	Likelihood	Impact	Mitigation
Studio management system integration takes longer than estimated, delaying the cross-studio client history feature and the rhythm-based host prompt	High	High. The cross-studio recognition promise, that a Mayfair client visiting Edinburgh is welcomed as though at home, is central to Phases 2 and 3. If the data pipeline from studio systems to the product cannot be established reliably, those features cannot ship with confidence.	Map all studio management system data contracts and authentication methods in week one of Phase 1, before any dependent feature is scoped into a sprint. Commission a technical spike lasting no more than two weeks to confirm the integration approach and surface any undocumented API constraints. If the integration proves more complex than the spike indicates, Phase 2's cross-studio deliverables are split: the data infrastructure is built to the boundary of what the integration allows, and the host-facing surfaces are released in a point release once the integration is confirmed stable. The cross-studio recognition feature does not ship with a fallback that shows partial data. It ships complete or it waits.

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

Risk	Likelihood	Impact	Mitigation
Cold start performance on mid-range Android devices does not meet the two-second load time standard for the home screen in live conditions	Medium	High. The home screen is the first second. If a returning client's host panel and suggested appointment take four seconds to render on the device they actually use, the product's first act is making them wait, which contradicts everything the design is built around. The standard was set against mid-range Android because that is where the gap between native and cross-platform performance is most visible. It still needs to be confirmed in live conditions.	Establish a continuous performance testing baseline on a defined set of mid-range Android devices, specifically the Samsung Galaxy A series and similar, from the first Phase 1 sprint. Run performance tests on every pull request that touches the home screen, the appointment suggestion engine, or the Redis cache layer. If the two-second standard is not met in staging conditions, the feature does not proceed to production. The mitigation for a persistent performance miss is not to adjust the standard. It is to investigate the cache hit rate, the payload size of the suggestion response, and the rendering path of the host panel, in that order, before any other explanation is considered.
Navigation patterns differ between iOS and Android in ways that create inconsistent emotional experiences at key journey moments	Medium	Medium. The moments where navigation behaviour matters most are the transition from the home screen to the booking confirmation, the deep-link routing from a push notification to a specific host message, and the back-navigation behaviour within the Albemarle story. If those moments feel different enough between platforms that a client on Android is doing something meaningfully different from a client on iOS, the product is not delivering a consistent experience.	Define the navigation specification for every primary flow before build begins, stating the expected behaviour on both platforms at every branch point. Where platform conventions diverge, the design decision about which convention to follow is made explicitly rather than by default. Navigation testing is part of the QA checklist for every flow, on both platforms, using the same device matrix used for performance testing. Any divergence between platforms at a critical journey moment is raised as a defect, not a platform difference, and resolved before the flow ships.

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

Risk	Likelihood	Impact	Mitigation
Copy review becomes a bottleneck as the volume of strings increases across Phase 2 and Phase 3	High	Medium. The copy review requirement is non-negotiable: no client-facing string ships without review against the Voice and Tone Guidelines and the banned phrase list. In Phase 1 the volume is manageable. As the post-visit note workflow, the Albemarle story, the life moment interface, and the preference layer are built, the volume of strings requiring review will increase significantly. If the review is treated as a gate that one person processes sequentially, it will slow delivery.	Establish a copy review workflow at the start of Phase 1, before the volume pressure arrives. The workflow separates strings into three tiers: primary copy, the content on screens a client sees during their core journey, which requires full review; system copy, error messages, empty states, loading states, which requires review against a checklist rather than a full guidelines read; and notification copy, which requires review and an additional motive test check. The content reviewer is involved in sprint planning, not just sprint review, so copy requirements are visible before features enter build rather than after they are functionally complete. A shared copy review queue is maintained in the team's project management tooling, with status visible to the full team. Any string that has been in review for more than 48 hours is flagged automatically. The content reviewer has the authority to request a one-day hold on a feature ship if copy has not been reviewed. That authority is confirmed before Phase 1 begins, not negotiated under pressure in Phase 3.

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

Risk	Likelihood	Impact	Mitigation
Quality gate failures at the Phase 1 to Phase 2 transition cause schedule compression that shortens the Phase 2 pilot period	Medium	High. The Phase 2 pilot is the only opportunity to test the authorship principle, the chair note interface, and the cross-studio recognition under real operational conditions before full rollout. If the pilot is shortened because Phase 1 ran long, the highest-risk assumption in the product is tested against less evidence than the design demands.	Build four weeks of buffer into the overall programme schedule, allocated specifically to the phase transition gates. If Phase 1 completes on time, the buffer is released back. If Phase 1 runs long, the buffer absorbs the slip without compressing the Phase 2 pilot below its minimum of four weeks. The minimum pilot duration is defined in the Phase 2 success criteria, not in the schedule, so it cannot be traded away as a schedule optimisation. If the buffer is exhausted and the pilot still cannot reach four weeks, the conversation is about which Phase 2 features are deferred to a point release, not about shortening the pilot itself.
Life moment recording, once available to hosts, is used without following the accuracy check and client awareness steps defined in the design principles	Medium	High. The life moment feature is the one with the most potential to cause genuine harm if it is used carelessly. A host who adds a life moment without the client's awareness, or who raises a life moment in outreach without confirming it is still accurate, is not making a minor process error. They are using personal information in a way the client did not consent to, against a backdrop where the product's entire promise rests on the client being in control of what is known and how it is used.	The accuracy check is a required step in the product interface, not a recommended one. The interface does not allow a life moment to be surfaced in a host communication without the accuracy check completing first. Host training at the Phase 3 launch briefing includes a specific scenario demonstrating what happens when a life moment is out of date, using the wedding cancellation example from the Aspiration Gap session. The briefing is not a one-time event. It is the documentation that a named Albemarle operations lead is responsible for maintaining and re-issuing when new hosts join. The life moment audit, confirming that every life moment used in outreach during the period passed the accuracy check, is part of the Phase 3 success criteria and is reviewed again at the 90-day post-launch review.

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

Risk	Likelihood	Impact	Mitigation
App store review rejection, particularly on iOS, delays the launch date	Medium	Medium. Apple's review process is not predictable, and a rejection for a privacy or data handling reason, given the volume of personal client data this product holds, is a realistic outcome. A rejection two days before a planned launch date affects studio briefings, host communications, and client expectations in ways that are disproportionate to the technical fix usually required.	Submit to both app stores a minimum of three weeks before the intended launch date, not two. Allocate one week of the Phase 4 schedule specifically to addressing a first-round rejection on iOS, so that a resubmission does not compress the final quality gate. The launch date is confirmed publicly only after both store approvals are in hand. Any external communications about the launch, to clients, to studio teams, to press, are issued after approval, not in anticipation of it.
The post-visit note reply rate does not show the expected upward trend during Phase 2, suggesting notes are not landing as personal	Low	High. A flat or declining reply rate is the most direct signal available that post-visit notes are being read as system confirmations rather than personal messages. If that signal appears, it means either the notes are not genuinely personalised, the authorship principle is being compromised in ways the audit is not catching, or the voice of the notes is drifting from the standard established in the design.	Monitor reply rates per studio and per stylist, not only as a programme-level aggregate. A studio with a significantly lower reply rate than its peers is a signal worth investigating. The investigation starts with a manual review of the notes sent from that studio during the measurement period, assessed by the content reviewer against the voice specification. If the notes are well-written and the reply rate is still low, the investigation moves to the authorship audit for that studio. If the audit is clean, the finding is shared with the Albemarle client team as a qualitative research question, not a product defect.

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

Risk	Likelihood	Impact	Mitigation
Browser or OS updates on iOS or Android introduce breaking changes in navigation, authentication, or push notification handling after launch	Low	High. Platform updates that change the behaviour of URLSessionWebSocketTask, the APNS delivery model, or the system permission flows for biometric authentication are outside the team's control and can introduce failures that were not present at launch.	Maintain a subscription to the iOS and Android developer release notes, with a named engineer responsible for reviewing each beta release for changes that may affect the product's core flows. Establish a smoke test suite that covers the WebSocket connection, the push notification routing, and the biometric authentication flow, run automatically against each new platform beta. Any breaking change is surfaced as a defect before the platform update reaches the client base, not after.

The pattern across this register is consistent. The risks that matter most are not the ones that stop the build. They are the ones that allow the build to continue while something important quietly fails. The authorship breach that goes undetected for three weeks. The copy that ships with a banned phrase because the review queue backed up. The life moment that surfaces in a host message without the accuracy check running because the interface made it possible to skip. Each of those is a failure that looks like a success from the outside, right up until a client notices, and by that point the trust that was damaged is harder to rebuild than the feature would have been to protect.

The mitigations are designed around that pattern. They do not aim to make failure impossible, which is not a realistic goal. They aim to make the honest path easier than the shortcut, the breach detectable before it compounds, and the standard visible to everyone on the team at the moment the pressure to compromise it is highest.

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