



Developer Handoff

Last Updated: 6 July 2026, 14:19

Prepared for Drift
Project Drift App

Contents

Introduction	3
Design Phase Deliverables	3
What Has Been Built	6
What Matters Most	8
Known Constraints & Decisions	13
System Copy Requirements	18
Acceptance Criteria	22
Emotional Criteria	22
Brand Criteria	24
Interaction Criteria	25
Copy Criteria	26
Performance Criteria	27
Pre-Launch Emotional Review	27
What the review covers	28
Who is involved	28
What the review produces	29
The sign-off process	29
Post-launch review cadence	30

Introduction

This document exists at the handoff point between what Drift was designed to be and what gets built. Design strategy without build fidelity is just aspiration. This document carries the intent forward, not the functional requirements, those live elsewhere, but the decisions behind the decisions. Why a screen looks the way it does. What would be lost if a particular interaction were simplified. Where the emotional logic of the product lives and how to protect it when the pressures of a build cycle start to pull things toward expedience.

To be clear about what this is not: it is not a functional specification, not a project plan, and not a design system. It does not define component libraries, user stories, or technical architecture. Those documents exist and this one sits alongside them, not above or below. Its job is narrower and, in some ways, harder. It holds the reasoning that does not naturally survive the translation from design to engineering, the copy rules that cannot be reduced to a style guide, the interaction principles that have no obvious ticket to live in, the emotional experience that a QA checklist will not catch.

The development team, QA, and anyone responsible for reviewing build fidelity against design intent should read this document before work begins and return to it when decisions arise that were not anticipated in the functional spec. If something is being simplified, deprioritised, or reinterpreted during the build, this document is the right place to check whether that change is genuinely neutral or whether it quietly undoes something that mattered. The design phase produced a body of work that is internally coherent. That coherence does not maintain itself automatically. This document is one of the mechanisms by which it survives.

Design Phase Deliverables

These documents did not arrive independently of one another. Each one emerged from a process of structured inquiry, and together they form a single coherent argument about what Drift is, who it is for, and how it should behave. A developer reading only the component specifications or the functional requirements will have the mechanics of the product but not its meaning. These documents supply the meaning. They answer the question that sits underneath every build decision: why was this done this way, and what would be lost if it were done differently.

The table below lists every document produced during the design phase. Where a document belongs to more than one audience, the primary audience reflects who most urgently needs to read it before work begins.

Document	Category	Primary Audience	Summary
Brand Strategy & Positioning	Research & Strategy	Product, Design, Development leads	Defines Drift's market position as the only focus tool built around variability rather than consistency. Establishes the core philosophical commitments that must not be diluted during build.
Emotional Journey Map	Research & Strategy	Development leads, QA	Maps the full user journey from first open to returning visit, identifying where emotional experience peaks, where it drops, and where the design is currently unresolved. The gap between sessions and the post-summary moment are named here as open problems.
Aspiration Gap Analysis	Research & Strategy	Product, Design	Documents the emotional state users arrive in, what they need to feel by the end of a first session, and the single biggest risk: a first session that does not work.
User Identity Shift Framework	Research & Strategy	Design, QA	Defines the transformation Drift is trying to create, from self-blame to self-knowledge, and the specific design conditions required to make it real. Any feature decision that works against this shift should be flagged here.
Voice & Tone Guide	Brand & Experience	All	Defines the complete verbal identity of Drift including words that are off-limits, rhythm rules derived from read-aloud testing, and the distinction between performing warmth and being warm. Applies to every string of copy in the product.
Design Principles & Checklist	Brand & Experience	Development, QA	Five core principles with rationale, measurable signals for each, and an 18-point design checklist to be applied to every screen before sign-off. This is the document QA should run against built screens.
Anti-Principles Reference	Brand & Experience	Development, QA	Documents the worst version of the product in detail. Any design decision that moves toward that version, whether in copy, interaction, or data surfacing, should be caught using this document.
Visual Direction Brief	Brand & Experience	Development, Design	Specifies the visual register derived from moodboard and spectrum exercises. Lists directions that are definitively off the table, including dark mode as default, gamification aesthetics, dashboard layouts, and high-contrast alert colours.

Document	Category	Primary Audience	Summary
Heartbreak Scale Report	Brand & Experience	Product, Design	Scores the product across five emotional dimensions: trust, delight, anxiety, confusion, and pride. Identifies the body doubling introduction as the single highest-anxiety moment and the single most confusing element. Specifies what would raise each score.
First 60 Seconds Design Brief	Brand & Experience	Development, Design	Defines the emotional experience required in the product's opening minute, the specific design decisions that create it, and where the experience currently breaks down. The task input scaffold, the mood check-in's visible reactivity, and the body doubling introduction screen are all specified here.
Dinner Party Brand Character	Brand & Experience	Design, Copywriting	Describes Drift's brand character through a social scenario. Used primarily to calibrate tone decisions and to resolve edge cases where the voice guide alone is insufficient.
Rejection Letter & Inversion	Research & Strategy	Product leads	Documents what Drift will never be and why each boundary matters. Useful for evaluating features or partnership decisions against the brand's core commitments.
Technical Constraints Summary	Technical	Development	Documents the constraints that shaped design decisions, including the rationale for interaction choices that may appear unconventional from a pure engineering perspective.
Component Specifications	Technical	Development	Defines behaviour, states, and properties for each UI component. Should be read alongside the Design Principles & Checklist to understand the emotional intent behind each component's specified behaviour.
Interaction & Animation Brief	Technical	Development	Specifies motion behaviour, timing, and the principle that animation in Drift should feel unhurried and never demand attention. Includes the distraction parking animation and session summary transition as priority items.

Developers should prioritise three documents before build begins. The Design Principles & Checklist is the most operationally useful: it translates the strategy into a concrete screen-by-screen test. The Voice & Tone Guide applies to every piece of copy in the product and contains hard rules, not

Everything else in this list supports and deepens those three. The Emotional Journey Map is the right place to return to when a build decision affects a transition or a state that was not fully specified in the functional requirements. The Heartbreak Scale Report is the right reference when evaluating whether a proposed change increases user anxiety or confusion. None of these documents replaces conversation between design and development. They exist to make that conversation more precise.

What Has Been Built

Drift is a focus app built around a single, honest observation: starting is harder than working. The brand exists for people whose brains experience the gap between intention and action as something closer to a wall than a speed bump, and it was designed by people who know that gap from the inside. Its personality is steady, warm, and plain-spoken. It does not perform care. It does not raise its voice to be heard. It sounds like someone who has sat with a lot of people who felt like they were failing at ordinary things and has never once implied the failure was the person's fault.

Who the brand is

Drift is the only focus tool built around variability rather than consistency. It does not ask users to become more organised, more disciplined, or more like the person the productivity category has always assumed they were. It asks one question, how are you feeling right now, and then it builds from there. Its voice is unhurried, witnessing, and honest. It never motivates, never compares, never shames, and never performs warmth it has not earned.

Who the user is

The user arrives carrying a long history with tools that did not work for their brain. They are not unmotivated. They are often intelligent, frequently creative, and genuinely want to do the thing. What they cannot reliably do is start it. They have cycled through enough systems that the failure has started to feel like a verdict about them rather than a mismatch between them and the tool. They fear that Drift will be another source of evidence for that verdict. Success, for them, is not productivity. It is the feeling, even once, that their brain was never the problem.

The emotional arc

At first contact, the user should feel one thing: safe enough to be honest. Not impressed, not motivated, just the quiet recognition that this space is not going to judge what they say next. The visual tone, the absence of pressure, and the copy on the first screen should communicate this before a single interaction occurs.

During the journey, the arc moves through cautious curiosity into the first crack of relief when the mood check-in responds visibly to their answer, then into the small but chemically real moment of activation when the first subtask completes, then into the quiet steadiness of a session with a body double present, and finally into the soft landing of a summary that names what they did without scoring it.

After sustained use, the arc arrives somewhere much more significant. The user develops self-knowledge rather than self-surveillance. The old story, "I'm lazy, I can't follow through, something is wrong with me", loses its grip. Not because Drift fixed them, but because Drift gave them enough small, accumulated evidence that it was never true.

Critical moments

These are the five moments where emotional experience is at its most consequential and where a build decision made for pragmatic reasons can do the most damage.

- The first screen. The user is scanning fast, in mild threat-detection mode, looking for the first sign that this is like everything else. Nothing on this screen should ask anything of them. No input field, no metrics, no feature highlights. The copy must sound like a person wrote it on a considered day. If anything on this screen feels demanding or clinical, the user begins to disengage before they have consciously decided to.
- The mood check-in. This is Drift's core philosophical commitment made visible. The question must feel genuinely curious rather than diagnostic. More critically, the answer must visibly change what Drift suggests next. If the response is invisible to the user, the question becomes a lie. The moment the user sees that their answer shaped what came next, something shifts. That is the first crack of relief.
- The body doubling introduction. This is simultaneously Drift's most powerful feature and its most fragile moment. Users who encounter the pairing screen without prior explanation consistently pause, hesitate, and in some cases leave. The introduction screen must come before the pairing flow, every time, without exception. It need only be one screen. Two sentences on what body doubling is and why it works. Explicit clarity on what the other person can and cannot see. Without this screen, Drift loses users at the exact moment it should be earning them.
- The first subtask completion. This is the pivot point between anxiety and motion. The brain that has been frozen gets its first piece of evidence that it can move. The acknowledgment must be immediate and quiet. Not a celebration, not a score, just a recognition. Something that says: that happened, it counted. The interaction should feel like a hand placed briefly on a shoulder, not a notification.
- The session summary. This is the most important piece of copy in the product. It is the moment the new belief either lands or does not. The language must speak to effort, not output. Not "3 tasks completed" but something that names what the user actually did in terms that acknowledge what it cost them. No scores, no percentages, no comparisons to previous sessions or other users. If a user completed very little, they should still feel seen by this screen. The summary is followed by a single soft optional prompt, one question, no consequence for leaving it blank, and then Drift lets them go.

Design principles

These five principles governed every design decision made during the design phase. They are stated here plainly so that any build decision that touches them can be evaluated clearly.

- The door back in must always feel safe. The app opens on the mood check-in regardless of how long the user has been away. No copy references time elapsed since the last session. The home screen after a ten-day gap looks identical to the home screen after one day. Any element that implies the user has something to make up for violates this principle.
- Ask before assuming, then respond visibly. The mood check-in is not a feature. It is the product's belief made concrete. But asking is only half of it. If the answer does not visibly change what appears on screen, the question is dishonest. Visible reactivity is not optional.
- The smallest completable unit is always visible. Every task is broken into subtasks before the session begins. The first subtask is always the smallest possible version of beginning. Drift never presents a blank field and waits. It offers a scaffold. The identity shift begins at the first completed action, not at the end of the task.
- Presence must feel like company, not surveillance. Body doubling is explained before it is offered. The body double panel during a session is ambient. It does not demand attention. What the other person is doing, how long they have focused, and how they are performing are never visible. Privacy is stated explicitly at the pairing moment, not buried in settings.
- The session summary speaks to effort, not output. No scores, no productivity percentages, no comparisons. The summary names what the user did in language that acknowledges what it cost them. It is followed by a single soft optional prompt and nothing more.

What Matters Most

These are not preferences. They are not design ideals that survive if the build has time for them. Each item below is load-bearing. Remove or compromise any one of them and something structural in the user experience collapses, usually quietly, usually in exactly the moment where the user was most vulnerable and most likely to decide whether Drift is different from everything else they have tried.

Non-negotiable	What it is	Why it matters	What happens if it is compromised
<p>The mood check-in response must be visibly reactive</p>	<p>When a user answers the mood check-in, something on screen must change in a way they can see. The session length, the suggested first step, the tone of the scaffold, some visible element must shift based on what they said.</p>	<p>The mood check-in is not a data collection mechanism. It is the moment Drift proves it was listening. The first crack of relief that users described, "this is different", happens here, and only if the response is visible. An invisible algorithmic adjustment is indistinguishable from no adjustment at all. The user asked to be seen. If Drift sees them and says nothing, the question becomes a lie.</p>	<p>The user completes the check-in, nothing visibly changes, and the session that appears looks identical to what any other answer would have produced. Trust does not break dramatically. It just quietly fails to form. The user continues, finishes a session, possibly comes back once. But the thing that was supposed to make Drift feel different from every other timer app never happened. They file it accordingly.</p>
<p>The body doubling introduction screen must appear before the pairing flow, every time</p>	<p>A single screen explaining what body doubling is, why it works, and what the other person can and cannot see must appear before any name or status is shown to a new user. This is not an optional onboarding step. It precedes the pairing flow unconditionally.</p>	<p>Workshop evidence was unambiguous on this point. Users who encountered the pairing screen without prior context paused, hesitated, and in several cases closed the app entirely. The feature that is most powerful in Drift is also the most socially loaded. Without context, a name appearing on screen reads as surveillance. With thirty seconds of warm, plain explanation, it reads as company. The difference is entirely in the framing, and the framing must be present before the feature is.</p>	<p>A user arrives at the pairing screen with no understanding of what they are agreeing to. The social anxiety this triggers is not a minor friction point. It is the specific anxiety that most prevents the people Drift was built for from engaging with anything new. They close the app at the feature that would most likely have made them return. The design debt identified in the Heartbreak Scale, body doubling as the most confusing element in the product, becomes a live defect rather than a documented risk.</p>

Non-negotiable	What it is	Why it matters	What happens if it is compromised
<p>No element, copy or visual, may reference time elapsed since the last session</p>	<p>The home screen after a ten-day gap must be identical to the home screen after one day. No copy, badge, counter, or visual state may communicate how long a user has been away. This applies to notifications, session states, and any longitudinal data surfacing.</p>	<p>The people Drift was built for are not forgetting to return. They are often avoiding returning because opening the app means confronting something they feel bad about. The avoidance that keeps them from their work starts keeping them from Drift. Every day that passes makes re-entry feel heavier. The design has one job at that moment: make the door feel light. A single line of copy that references their absence, even a warm one, even a well-intentioned one, places weight on that door at precisely the moment when the user is most fragile. The Anti-Principles work named this failure mode in detail. It takes four words to close the app for good.</p>	<p>A user returns after six days, opens Drift, and sees something that tells them Drift noticed they were gone. Not harshly. Perhaps warmly. It does not matter. The sensation is of having something to explain, of re-entry as confession rather than return. They close the app. They do not delete it, they just do not open it again for a while, and then longer, and then the app sits in a folder with the other things that made them feel like this. The principle that was supposed to protect them was not in the room when the copy was written.</p>

Non-negotiable	What it is	Why it matters	What happens if it is compromised
<p>The session summary must use effort language, not output language</p>	<p>The summary screen must name what the user did in language that acknowledges what it cost them. No scores, no productivity percentages, no task counts presented as measures of worth, no comparisons to previous sessions or other users. If a user completed one small thing on a hard day, the summary must reflect that in terms that make it feel real rather than inadequate.</p>	<p>This is the most important piece of copy in the product. It is the moment the new belief either lands or does not. The user has spent the session accumulating small evidence that their brain can move. The summary is where that evidence becomes a story they tell about themselves. Output language, "3 tasks completed," "62% focus quality", turns that evidence back into a measurement, and these users have been measured against standards they did not set for long enough that measurement alone can undo everything the session built. Effort language, something closer to "you kept going", gives the evidence a form that changes self-perception rather than quantifying performance.</p>	<p>The summary reads like a report. The user looks at what it says and feels assessed rather than seen. The emotional peak of the session, the quiet satisfaction of having done something, is replaced by the mild deflation of having produced a number. For users who completed very little, the number confirms the story they were trying to escape. For users who completed more, the number turns the experience of trying into a metric, which is exactly the reduction Drift was built to resist. The identity shift does not land. The session was good. The summary made it smaller.</p>

Non-negotiable	What it is	Why it matters	What happens if it is compromised
<p>The first screen must contain no input field and no element that implies consequence</p>	<p>The opening screen of Drift must not ask the user to do anything before the mood check-in. No task input, no setup prompt, no streak indicator, no notification badge, no feature highlight requiring interpretation. The user lands somewhere that simply exists, calm and unhurried, before any demand is made.</p>	<p>Users arrive in a mild threat-detection state. They are scanning fast, looking for the first sign that this is like everything else. Their nervous system is not in a position to receive complexity. Any element that reads as a demand, an empty field waiting to be filled, a counter showing what has or has not happened, a tooltip pointing at a feature, registers as cognitive load before the user has consciously processed it. The emotional promise of Drift is that the door is light. The first screen is the door. If it is heavy, the promise is already broken.</p>	<p>The user opens Drift for the first time, or for the first time in a week, and the screen they land on contains something that requires them to know something, decide something, or confront something before they are ready. The specific failure described in the Aspiration Gap, "another system to fail at," "another thing I have to manage my feelings about", happens in this moment. They do not necessarily close the app immediately. They just begin the session already slightly contracted, already slightly defended, and that is not the state from which the rest of the experience was designed to be received.</p>

Non-negotiable	What it is	Why it matters	What happens if it is compromised
<p>The body double's task, timer, and progress must never be visible to the paired user</p>	<p>During a body doubling session, the only information visible about the other person is their name and a status indicator confirming they are present. What they are working on, how long they have been focused, and how much they have completed must not appear anywhere in the session UI, including in any aggregated or anonymised form.</p>	<p>Body doubling works because of presence, not comparison. The feeling that creates quiet momentum in a session is "someone else is here, also trying." The moment the user can see how long the other person has been focused, or what they have completed, that feeling becomes "someone else is doing better than me." For users who already carry a heavy load of self-comparison, that shift is not minor. It is the difference between feeling accompanied and feeling outperformed. The Heartbreak Scale identified comparison as a landmine dressed as a feature. This is where that landmine lives in the product.</p>	<p>A user in a body doubling session glances at the presence panel and sees that the other person has been focused for two hours. They have been going for four minutes. The session that was creating quiet momentum becomes a source of social anxiety. They either close the app or complete the session feeling worse about themselves than when they started. The feature designed to be Drift's most powerful differentiator becomes another mechanism of shame. Not through malice, not through carelessness, but through a single data point surfaced in the wrong context.</p>

Known Constraints & Decisions

Every decision captured in this document was made with full context: the emotional reality of the user, the principles established during the design phase, and the specific failures that Drift was built to avoid. That context does not automatically travel into the build phase. The decisions looked obvious when the reasoning was in the room. Without the reasoning, some of them will look like constraints worth reconsidering, places where a small simplification seems harmless or an added piece of information seems helpful. This section exists to prevent that. Each row below names a decision, explains why it was made, and describes what not to do, because the most common way these decisions get reversed is not through explicit disagreement but through well-intentioned changes made by people who were never told why the original choice mattered.

Decision	Rationale	What Not To Do
<p>The app always opens on the mood check-in, regardless of how long the user has been away</p>	<p>Re-entry after a gap is the most emotionally loaded moment in the product. The user is already carrying low-level dread. The only job of the first screen is to place them in the present tense, not the past. The mood check-in does this: it asks about right now and nothing else. It communicates that Drift has no opinion about how long it took to return.</p>	<p>Do not open on a task list, a home dashboard, or any screen that surfaces incomplete items from a previous session. Do not open on a screen that shows the date of the last session, a streak counter, or any visual indicator of time elapsed. Do not treat the mood check-in as a skippable step that can be bypassed to get the user "into the product faster." It is not a gate. It is the product.</p>
<p>The mood check-in response must be visibly reactive</p>	<p>Asking how someone feels and then ignoring the answer is worse than not asking. The mood check-in is Drift's core philosophical commitment made tangible. The user's answer must visibly change what appears next: session length, task size, the structure of the first suggested step. The change must be perceptible to the user, not only to the algorithm.</p>	<p>Do not implement the mood check-in as a data collection step that feeds a backend recommendation engine while the screen remains visually unchanged. Do not use the same session template regardless of what the user selects. Do not surface the check-in and then route every answer to an identical next state. If the response is invisible, the question is a lie, and users will feel that before they can articulate it.</p>
<p>No copy, badge, counter, or visual state may reference time elapsed since the last session</p>	<p>The people Drift was built for are not forgetting to return. They are avoiding it, because returning means confronting the thing they feel bad about. The avoidance that keeps them from their work starts keeping them from Drift. A warm sentence that mentions their absence lands, in the chest of that user, as: we noticed, we kept count, you have some explaining to do. The Anti-Principles work identified this as the most dangerous failure mode in the product, because it does not feel like a failure. It feels like care.</p>	<p>Do not include "welcome back" copy on the re-entry screen. Do not surface "last session: Tuesday" or any equivalent. Do not add a notification or in-app message that references the gap. Do not build a streak counter, even one framed positively. Do not add a badge or visual state to the app icon that implies something has been missed. Any element that tells the user how long they have been away places weight on the door at the moment the user is most fragile.</p>

Decision	Rationale	What Not To Do
<p>The body doubling introduction screen must appear before the pairing flow, every time</p>	<p>Users who encounter the pairing screen without prior explanation consistently hesitate, and some close the app entirely. The feature that is most powerful in Drift is also the most socially loaded. Without context, a stranger's name appearing on screen reads as surveillance. The introduction screen need only be one screen: warm language, two sentences on what body doubling is and why it works, explicit clarity on what the other person can and cannot see. That is enough. Without it, Drift loses users at the exact moment it should be earning them.</p>	<p>Do not skip the introduction screen to reduce onboarding length. Do not move the privacy explanation to a settings page or a help section. Do not present the pairing flow before the user has been told what they are agreeing to. Do not assume that users who have encountered body doubling in another context understand what it means here. The explanation must be present, every time, for every new pairing experience, because the trust it builds cannot be assumed in advance.</p>
<p>The body double's task, timer, and progress must never be visible to the paired user</p>	<p>Body doubling works through presence, not comparison. The feeling that creates quiet momentum is "someone else is here, also trying." The moment a user can see how long the other person has been focused, that feeling becomes "someone else is doing better than me." For users who already carry significant self-comparison, that shift is not minor friction. It is the difference between company and competition. The Heartbreak Scale identified this as a live landmine, and the product's entire design argues against comparison as a mechanism.</p>	<p>Do not surface session duration, task names, task completion counts, or focus quality scores for the body double, even in anonymised or aggregated form. Do not add a leaderboard, a shared progress view, or any visualisation that allows one user to assess how they are performing relative to the other. Do not build a "mutual accountability" feature that shares completion data between paired users. The presence panel must show a name and a status indicator. Nothing more.</p>
<p>The session summary must use effort language, not output language</p>	<p>The summary is the most important piece of copy in the product. It is the moment the new belief either lands or does not. Output language turns the session back into a measurement, and these users have been measured against standards they did not set for long enough that a number alone can undo everything the session built. Effort language gives the evidence a form that changes self-perception. The difference between "3 tasks completed" and "you kept going" is not aesthetic. It is the difference between data and being seen.</p>	<p>Do not include a productivity score, a focus quality percentage, or a task count presented as a measure of worth. Do not compare the session to previous sessions or to any aggregate standard. Do not add a celebratory animation or tone that implies the user should feel proud in a loud way. Do not build a shareable result card. Do not end the summary on a prompt that feels like a demand. The summary is followed by one soft, optional question and then Drift lets them go.</p>

Decision	Rationale	What Not To Do
<p>The first screen must contain no input field and no element that implies consequence</p>	<p>Users arrive in mild threat-detection mode. They are scanning fast, looking for the first sign that this is like everything else. Any element that reads as a demand, an empty field, a counter, a badge, a tooltip pointing at something requiring interpretation, registers as cognitive load before the user has consciously processed it. The emotional promise is that the door is light. The first screen is the door.</p>	<p>Do not add a task input field to the home screen. Do not surface incomplete tasks from a previous session as the default state. Do not show notification badges, streak counters, or progress indicators on the opening screen. Do not include feature callouts, tooltips, or onboarding prompts that require the user to understand the product before they have experienced it. Do not use the opening screen to upsell, prompt a subscription, or request permissions before the user has felt any value from the product.</p>
<p>The word "just" must not appear before any ask or instruction</p>	<p>"Just" implies the thing being asked is easy. Drift's users know that starting is not easy. They have tried to "just start" more times than they can count. Every time the word appears before an instruction in Drift's copy, it quietly communicates that the struggle is the user's failure rather than the design's responsibility. The Voice Sort exercise named this as an explicit off-limits construction. It is not a preference. It is a commitment.</p>	<p>Do not write "just open the app," "just start with one task," "just take the first step," or any equivalent. Do not use minimising language of any kind before an instruction. If a step is genuinely small, describe it specifically rather than characterising it as easy. The size of the step speaks for itself. Drift does not need to reassure the user that it is small, because doing so implies it should not feel hard, and for this user it sometimes will.</p>
<p>Copy must never use the words: streak, optimise, track, consistency, improve, performance, or "stay on track"</p>	<p>These words belong to the world Drift was built against. Each one carries a set of implicit expectations, about regularity, about measurability, about the right way to show up, that Drift's users have already been failed by. Even a single instance in a notification, a tooltip, or a settings label reactivates the emotional register of every productivity tool that made them feel like they were doing it wrong. The Voice Sort exercise produced a clear list of wrong-world language. This is it.</p>	<p>Do not use these words anywhere in the product: in copy, labels, notification strings, empty states, onboarding, or settings. Do not introduce synonyms that carry the same register: "maintain your habit," "keep up your routine," "build momentum," "power through." Do not frame the body doubling feature as an "accountability" mechanism. Do not describe the session summary as a "performance review" or similar. When in doubt, run the copy through the 18-point design checklist in the Design Principles document before shipping.</p>

Decision	Rationale	What Not To Do
<p>Notifications, if they exist at all, must be framed as invitations to the present, never reminders of absence</p>	<p>The gap between sessions is the most vulnerable moment in the user journey. Their brain fills that silence with avoidance and self-doubt. A notification that references what they have not done, however gently framed, adds to that weight. The only acceptable notification is one that creates an opening to the present moment without any implication of what should have happened before it.</p>	<p>Do not send "you haven't used Drift in X days" notifications. Do not frame notifications around returning, picking up where they left off, or not losing progress. Do not schedule notifications at intervals that imply a correct usage frequency. Do not add push notification prompts early in the onboarding flow before the user has experienced the product's value. If notifications are implemented, the copy must pass the same test as all other Drift copy: read it as someone returning after a bad week and ask whether it makes the door feel lighter or heavier.</p>
<p>The visual register must be low density, unhurried, and never resemble a dashboard</p>	<p>The Opposite Ends exercise placed Drift at the sparse end of the visual density spectrum and the calm end of energy level. The moodboard responses confirmed this direction: the product's visual tone should feel like a quiet morning, not a cockpit. Any screen that surfaces multiple data panels simultaneously communicates that the user is being monitored and assessed. That is the opposite of what Drift is.</p>	<p>Do not design home screens or session screens with multiple simultaneously visible data panels. Do not use dark mode as the primary or default visual register. Do not use neon or high-saturation colour anywhere in the product. Do not use gamification visual language: XP meters, achievement badges, flame icons, trophy graphics. Do not use bold geometric sans-serif typography that reads as tech or corporate. Do not use gradient backgrounds. Do not use high-contrast red or orange as primary action colours. Do not use split-screen or multi-column layouts on mobile. One thing at a time, always.</p>
<p>Onboarding must not ask the user to configure anything before they have felt value</p>	<p>The Aspiration Gap analysis identified onboarding that asks too much before giving anything as the primary risk for immediate drop-off. These users have already spent activation energy downloading the app. Asking them to fill out a profile, select a productivity style, or configure a timer before their first session confirms their fear that this is another system to learn before it will work for them.</p>	<p>Do not include a productivity style selector in onboarding. Do not ask users to set goals, create task categories, or configure notification preferences before their first session. Do not present a setup wizard with multiple steps before value is delivered. Do not make the mood check-in dependent on prior configuration. Drift picks sensible defaults based on the check-in answer and adjusts over time. Configuration that requires explicit user input should be available, but it must be earned through use rather than demanded before it.</p>

Decision	Rationale	What Not To Do
<p>The distraction parking feature must feel like a hand on the shoulder, not a filing cabinet</p>	<p>The act of catching a wandering thought and setting it aside without judgment is genuinely meaningful for an ADHD brain in a session. The feature must honour that with a design that feels warm rather than functional. A gentle animation, a single word of acknowledgment, something that communicates "got it, keep going" without demanding the user's full attention. The Heartbreak Scale identified this as a missed opportunity for quiet delight.</p>	<p>Do not implement distraction parking as a plain input field that accepts text and returns the user to the session without acknowledgment. Do not use language that implies the thought has been filed or stored for later review in a way that sounds administrative. Do not add a notification or follow-up prompt about parked thoughts during the session. The interaction should take no more than two seconds of the user's attention, close itself, and return them to the session with a single small signal that the thought has been held.</p>
<p>The post-summary moment must hold the user in the feeling before releasing them</p>	<p>The session summary is the emotional peak. Immediately after it, the inner critic fills whatever silence the design leaves. The post-summary moment must be a single soft optional prompt, one question, no consequence for leaving it blank, that holds the warmth of the session alive for one more beat. It is not a feature. It is the difference between an experience that sets the user down gently and one that drops them.</p>	<p>Do not route the user directly from the session summary to the home screen or task list. Do not add multiple prompts, a rating request, a share button, or a subscription upsell to the post-summary screen. Do not make the optional prompt feel obligatory through visual weight or positioning. Do not design the closing moment of the session as a gateway to the next one. The user has just done something. Let that be true for a moment before anything else is asked of them.</p>

System Copy Requirements

The table below covers every category of system-generated copy in Drift. These are the strings that live outside designed screens: the messages that appear when something goes wrong, when a state is empty, when a process is running, when an action completes. They are easy to treat as functional text. They are not. For a user whose relationship with productivity tools has been shaped by years of quiet shame, the copy that appears when nothing is happening, or when something has failed, carries as much emotional weight as any designed screen. Possibly more, because those are the moments nobody thought to design for carefully.

Every string in the table below must pass the same test as all other Drift copy before it ships. Read it aloud. Ask whether it sounds like a person wrote it for this specific user on a hard day. Ask whether it implies the user has something to make up for. Ask whether the last word lands on something concrete. If any of those answers give pause, the copy is not ready.

Category	Examples	Tone Guidance	Review Required
Error messages	Network failure during session; pairing failure during body double match; session data not saved; mood check-in response not registered	Errors must be stated plainly and without apology that edges into self-flagellation. Do not use "oops," "uh oh," or any register that performs casualness to soften a real problem. Do not use language that implies the user did something wrong. State what happened in one short sentence and what, if anything, they can do. If nothing can be done, say so directly. Errors must never carry exclamation marks. They must never frame a technical failure as something the user needs to manage their feelings about.	Design lead and copywriter sign-off before any error string ships. QA to test error states against the 18-point design checklist, specifically questions 16 and 17.
Empty states	No tasks added yet; no body double available right now; no sessions completed this week; distraction parking is empty	Empty states in Drift carry a particular risk. An empty task list can read as accusation. "Nothing here yet" can feel like a quiet verdict. Empty states must be framed as openings, not absences. They should acknowledge the present moment without implying that something should already be there. They must never use the word "yet" in a way that implies the user is behind. They must never reference what the user has or has not done previously.	Copywriter review before any empty state ships. Any empty state that surfaces data about previous sessions or previous use must be escalated to the design lead before implementation.
Loading states	Session loading after mood check-in; body double pairing in progress; session summary generating; first-time app load	Loading states must feel unhurried. They must not perform the product's sophistication by surfacing technical language about what is happening underneath. They must not use progress bars that imply the user is waiting for something complex. If a loading state requires copy, it should be short, plain, and present-tense. It should feel like the product is simply getting ready, not processing the user.	Copywriter review. Development to flag any loading state that exceeds three seconds to the design lead, as extended loading states require specific copy treatment to prevent the user from feeling that something has gone wrong.

Category	Examples	Tone Guidance	Review Required
Confirmation messages	Session started; subtask marked complete; thought parked; body double session joined; session ended	Confirmation copy is where Drift's witnessing voice is most exposed. These are the moments of acknowledgment. They must not celebrate loudly. They must not use exclamation marks. They must not use the word "great," "amazing," "well done," or any equivalent. A subtask completing should feel like a hand on the shoulder: brief, warm, and gone before it becomes a performance. The language must be effort-facing, not output-facing. "You did something" is closer to right than "task complete."	Copywriter and design lead review before any confirmation string ships. QA to specifically test completion states to confirm no exclamation marks, scoring language, or comparison language is present.
Onboarding copy	Mood check-in introduction; body doubling explanation screen; first session prompt; task scaffold copy	Onboarding copy carries the full weight of the first impression. Every string must pass the First 60 Seconds test: does this sound like a person wrote it for this specific user, on a hard day, without expecting anything of them yet? The body doubling introduction screen is the single most critical piece of onboarding copy in the product. It must explain what body doubling is, why it works, and what the other person can and cannot see, in plain language, in no more than two sentences per idea, before any pairing flow appears. The mood check-in introduction must frame the question as genuinely curious rather than diagnostic. No onboarding copy may use the words: streak, optimise, track, consistency, improve, performance, or "stay on track."	Full copywriter and design lead review before onboarding ships. The body doubling introduction screen requires a separate sign-off from both the design lead and a representative user, given the evidence of its fragility in testing.

Category	Examples	Tone Guidance	Review Required
Notification copy	Session reminder; return invitation; body double available; session summary ready	Notifications, if they exist at all, must be framed as invitations to the present. They must never reference time elapsed since the last session. They must never imply the user is behind, has missed something, or has a streak to protect. The test for every notification string is simple: read it as someone returning after a bad week. Does it make the door feel lighter or heavier? If the answer is heavier, the copy is not ready. Notifications must not end with exclamation marks. They must not use urgency language of any kind. The acceptable register is warm, unhurried, and entirely present-tense.	Copywriter review before any notification string is registered in the system. Design lead to approve the full notification copy set before push notification permissions are requested from users. Any notification that references session history, usage frequency, or elapsed time must be escalated and reconsidered before shipping.
Success states	Session completed; first session completed; body double session completed; first subtask checked off	Success states must not celebrate loudly. They must not score. They must not compare. They must not tell the user they "crushed it," "smashed it," or "kept their streak alive." The session summary is the primary success state in Drift and it is covered in full in the non-negotiables section of this document. For smaller success moments, the subtask check, the parked thought, the session joined, the register must be quiet and witnessing. Something happened. Drift noticed. That is enough. The language of success in Drift is always effort-facing: "you kept going" rather than "3 tasks completed."	Copywriter and design lead review. QA to test all success states specifically for output language, scoring language, comparison language, and exclamation marks. Any success state that surfaces a number as a measure of worth must be flagged before shipping.

Category	Examples	Tone Guidance	Review Required
Help and support text	Body doubling FAQ; mood check-in explanation; distraction parking guidance; session summary guidance; privacy information	Help text in Drift must not sound like a manual. It must not use passive constructions, bullet points that feel like instructions, or language that implies the user should already know this. The register is the same as the rest of the product: plain, warm, unhurried. Help text exists because something was unclear, and the appropriate response to that is honest explanation, not hedging. Where help text touches body doubling, it must include explicit clarity on what other users can and cannot see. Where it touches the mood check-in, it must explain why the question is asked, not just that it is asked. Privacy-related help text must be direct and specific, never vague.	Copywriter review before any help or support text ships. Any copy that touches privacy, data, or the body doubling feature specifically must be reviewed by the design lead and, where relevant, a product lead before going live.

On the review process

System copy review is not a final check before release. It is a condition of release. No string in any of the categories above should be treated as a developer placeholder that gets reviewed when someone has time. These strings are part of the product. They carry the brand. They speak to users at the moments when users are least defended and most likely to decide whether Drift is different from everything else they have tried.

The review process is as follows. The copywriter produces or approves every string before it enters QA. QA runs the 18-point design checklist against system copy states in the same way it runs against designed screens. Any string that fails the checklist, or that contains prohibited language from the Voice Guide, is returned before the build progresses. The design lead holds final sign-off on the body doubling introduction, the session summary, and all notification copy, given the evidence base for their fragility. Where strings are generated dynamically, the copywriter must review the full set of possible outputs, not a sample, before the feature ships.

The most common way system copy fails Drift is not through carelessness. It is through copy written quickly by someone who understood the product well but was not in the room when the voice was established. The result is usually warm, usually well-intentioned, and occasionally exactly the four words that make a user close the app for the last time. The review process exists to keep that person's good intentions from becoming the product's most damaging moment.

confirm them, the build is not done.

- **First screen, before any interaction.** A user opening Drift for the first time, or returning after any gap, must be able to land on the opening screen without encountering anything that asks something of them, implies something about them, or surfaces anything about their history with the app. The feeling must be: this is a space where nothing is being held against me. If a tester reports any sensation of demand, consequence, or deficit on the opening screen, this criterion has not been met.
- **Mood check-in, after answering.** The user must be able to see that their answer changed something. Not infer it, not trust that it probably did. See it. The session shape, the suggested first step, or some visible element of what follows must shift in a way that confirms the question was real. The feeling must be: I was heard. If a tester selects different mood states and the screen that follows is visually identical each time, this criterion has not been met.
- **Body doubling introduction, before pairing.** A user who has never experienced body doubling must be able to arrive at the concept through a single warm, plain screen before any name or status appears. The feeling must be: I understand what I am agreeing to and it sounds like company, not surveillance. If a tester reports confusion, hesitation, or social anxiety at the pairing screen without having first seen the introduction screen, this criterion has not been met.
- **First subtask completion, in session.** The moment a user checks off their first subtask must be acknowledged immediately, quietly, and without performance. The feeling must be: that happened, it counted. Not: I am being celebrated. Not: I have been scored. If a tester reports that the acknowledgment felt loud, competitive, or absent entirely, this criterion has not been met.
- **Session summary, after completing a session.** A user reading their session summary must feel seen rather than assessed. This applies regardless of how much they completed. A user who finished one small thing on a hard day must be able to read the summary and feel that what they did was real and enough. The feeling must be: Drift noticed what I did and it cost something, and that counts. If the summary contains a score, a percentage, a comparison, or language that quantifies output as a measure of worth, this criterion has not been met.
- **Re-entry after a gap, first screen seen.** A user returning after five or more days must experience re-entry as return, not confession. The opening screen must feel identical in emotional tone to the screen they would see after one day away. The feeling must be: nothing has changed, nothing is being held against me, the door was still open. If a tester returning after a simulated gap encounters any copy, visual state, or interface element that communicates the passage of time, this criterion has not been met.
- **Post-summary moment, after the session summary.** The feeling that follows the session summary must be held for one beat longer before the user is released. There must be a single soft optional prompt. No consequence for leaving it blank. The feeling must be: I am being set down gently, not dropped. If the user is routed directly from the summary screen to the home screen or task list without a transitional moment, this criterion has not been met.

Brand Criteria

These criteria define the baseline for voice, tone, and visual consistency. They apply to every screen, string, and state in the product.

Voice and copy

- No copy anywhere in the product, including error messages, empty states, loading states, notifications, onboarding strings, and help text, contains any of the following words: streak, optimise, track, consistency, improve, performance, "stay on track," "get back on track," "pick up where you left off," "you've been away."
- No copy uses the word "just" before any ask or instruction.
- No copy ends on a vague or abstract word. Every sentence closes on something concrete.
- No exclamation marks appear anywhere in the product.
- No copy uses the following registers: motivational coaching ("you've got this," "keep going, you're doing great"), clinical or diagnostic ("your attention score," "focus deficit"), or comparative ("you're ahead of your average").
- All copy has been reviewed by the copywriter and passed the 18-point design checklist before shipping. System copy, including dynamically generated strings, has been reviewed in full, not by sample.
- The body doubling introduction screen copy has received separate sign-off from the design lead before the pairing flow is accessible to users.
- All notification strings have been approved by the design lead before push notification permissions are requested from users.

Visual language

- The default visual register is light, not dark. Dark mode is not the primary aesthetic of any screen.
- No screen contains neon, high-saturation, or pulsing colour used as a primary visual element.
- No screen uses a dashboard layout with multiple simultaneously visible data panels.
- No gamification visual language appears anywhere in the product: no XP meters, achievement badges, flame icons, trophy graphics, or progress bars styled as reward mechanisms.
- No high-contrast red or orange appears as a primary action colour.
- No split-screen or multi-column layout is used on mobile.
- Typography reads as considered and human, not as tech-corporate or geometric-bold.
- Every screen has been reviewed against the Opposite Ends visual spectrum positioning before sign-off: low visual density, calm energy, generous spacing, one primary element per screen.

Interaction Criteria

These criteria define specific interaction patterns that must function as designed. Where an interaction was designed to create a specific emotional effect, that effect is named alongside the pattern.

- **Mood check-in response.** The session shape presented after the mood check-in must visibly differ based on the user's answer. The difference must be perceptible without explanation. A tester selecting "foggy" must see a meaningfully different suggested session to a tester selecting "ready to go." The mechanism of change must be visible on screen, not only in the underlying recommendation logic.
- **Re-entry experience.** The app must open on the mood check-in regardless of how long the user has been away. No logic may route returning users to a different opening state based on time elapsed since last session. A tester simulating a ten-day gap must encounter the identical opening experience as a tester returning after one day.
- **Body doubling introduction flow.** The introduction screen must appear before any pairing flow, every time, for any user who has not previously been paired. The introduction screen must include explicit confirmation of what the body double can and cannot see. The pairing flow must not be accessible without passing through this screen.
- **Body double presence panel.** During a session, the presence panel must display only a name and a status indicator. No task information, timer duration, focus quality, or completion data belonging to the body double may appear anywhere in the session UI. A tester must be able to confirm that none of this information is accessible, including through any secondary interaction such as tapping the presence panel.
- **First subtask acknowledgment.** The moment a subtask is marked complete, the acknowledgment interaction must be immediate, quiet, and self-closing. It must not require a tester to dismiss it. It must not contain a score, a rating, or comparative language. The interaction should complete within two seconds and return the user to the session without requiring any further input.
- **Distraction parking.** The parking interaction must accept input, acknowledge it, and close without requiring more than two seconds of the user's attention. It must return the user to the session with a single small signal of acknowledgment. It must not trigger a follow-up notification or prompt about the parked thought during the session.
- **Post-summary prompt.** The screen following the session summary must contain a single optional text field. It must be clearly optional: no visual weight, no consequence signalling for leaving it blank. The screen must not contain a share button, a rating request, a subscription prompt, or a route to starting a new session immediately.
- **Task scaffold.** The task input step must not present a blank field with a cursor waiting. It must present a soft prompt, framed as a question about intention rather than an instruction to produce. An optional path to breaking a task into smaller steps must be accessible without being the default view.

- **No input field on the opening screen.** A tester arriving at the opening screen must not encounter any element that requires input before the mood check-in. This includes task fields, setup prompts, configuration steps, and permission requests.

Copy Criteria

All copy must pass tone review before shipping. The process is as follows.

Categories requiring review before QA

- All onboarding strings, including mood check-in introduction, body doubling introduction, task scaffold copy, and first session prompt.
- All notification strings.
- All error message strings.
- All empty state strings.
- All confirmation and success state strings.
- All session summary strings, including any dynamically generated language.
- All help and support text, including any copy that touches privacy, body doubling, or the mood check-in.

The copy review test

Every string must pass all five of the following before it is considered ready.

- Read it aloud. Does it sound like a person wrote it for a specific user on a hard day, or does it sound like a product team wrote it for everyone?
- Does the last word of every sentence land on something concrete?
- Does any word in the string appear on the prohibited list in the Voice and Tone Guide?
- Does the string, in any framing, imply the user is behind, has missed something, or has something to make up for?
- Would a user returning after a bad week feel the door get lighter or heavier after reading it?

Categories requiring design lead sign-off in addition to copywriter review

- Body doubling introduction screen copy.
- Session summary copy, including all dynamic variants.
- All notification strings, before push permission is requested.
- Any copy that surfaces user history, patterns, or longitudinal data.
- Any copy that touches privacy or data visibility.

Dynamically generated strings

- The copywriter must review the full set of possible outputs for any dynamically generated string, not a representative sample. Sign-off on one variant does not constitute sign-off on the category.

Performance Criteria

These thresholds exist because emotional pacing is a design property, not only an engineering one. A loading state that lasts four seconds breaks the feeling of being in safe, unhurried hands. A session that stutters at the subtask completion moment undercuts the only piece of feedback the user has received that their brain is capable of moving. These are not aspirational targets. They are emotional requirements expressed as performance constraints.

- **App open to mood check-in visible.** The opening screen must load within 1.5 seconds on a standard mid-range device on a 4G connection. A loading state lasting longer than 1.5 seconds must trigger specific copy treatment, reviewed by the copywriter, that maintains the unhurried register without implying something has gone wrong.
- **Mood check-in response to session suggestion.** The transition from mood check-in submission to the session suggestion screen must complete within 2 seconds. If it exceeds this, a loading state must appear. The loading state must not surface technical language or imply processing complexity. It should feel like the product is simply getting ready.
- **Subtask completion acknowledgment.** The acknowledgment interaction following a subtask completion must appear within 300 milliseconds of the tap. Any delay beyond this breaks the immediacy that makes the moment land. The interaction must complete and close within 2 seconds without user input.
- **Session summary generation.** If the session summary requires any generation time, that time must not exceed 2 seconds. The transition from session end to summary screen must feel like a natural pause, not a wait. If generation consistently exceeds 2 seconds, this must be escalated to the design lead before the feature ships, as it affects the emotional pacing of the most important moment in the product.
- **Body double pairing.** If no body double is immediately available, the waiting state must not feel like rejection or inadequacy. The copy and visual treatment for a waiting state must be reviewed by the copywriter before shipping. The wait must not be presented as an error.
- **Distraction parking.** The parking interaction must accept input and close within 2 seconds of submission. Any lag in this interaction pulls the user's attention from the session at the exact moment the design is trying to return it there.
- **Animation timing.** All animations in the product must feel unhurried. No transition should feel rushed or abrupt. The interaction and animation brief specifies timing in detail. QA must confirm that no animation, under normal device conditions, plays faster than the specified timing, as compression makes interactions feel urgent in a way that contradicts the product's emotional register.

Pre-Launch Emotional Review

The pre-launch emotional review is a structured walkthrough of the built product conducted before any release candidate is approved for shipping. It is not a QA pass. QA confirms that the product

functions correctly. This review confirms that it feels correct, that the emotional intent of the design phase has survived the build, and that the moments that matter most have not been quietly diluted by decisions made under time pressure.

It happens once per release candidate, before the build is approved. If a release candidate fails the review and is revised, the review runs again before the revised candidate ships.

What the review covers

The walkthrough follows the user's emotional journey from first open to post-summary moment, checking each critical stage against the Feel Factor layers established during the design phase. It does not run through every screen systematically. It runs through every emotionally consequential moment, which is a different list.

The stages covered are, in sequence: the opening screen before any interaction; the mood check-in question and its visible response; the task scaffold and first input moment; the body doubling introduction screen and pairing flow; the body doubling presence panel during a session; the first subtask completion acknowledgment; the distraction parking interaction; the session summary screen; the post-summary prompt; and the re-entry experience after a simulated gap of five or more days.

For each stage, the review asks three things. Does the emotional experience match the intent documented in the design phase? Are there any elements, copy or visual, that introduce signals that should not be there, urgency, comparison, absence, inadequacy? And does the stage transition into the next one without dropping the user?

System copy states are also reviewed in this session, not by reading through a spreadsheet of strings but by triggering them live in the build. Error states, empty states, loading states, and confirmation messages are tested as experienced, not as documented. The review team reads each one aloud. If it does not pass the read-aloud test, it is flagged.

Who is involved

The review requires four people to be present. The design lead, who holds the emotional intent of the design phase and is the primary arbiter of whether a moment has landed correctly. A development lead, who can speak to what is technically feasible to adjust within the remaining build window and who flags anything the review identifies as requiring engineering time. A copywriter, who reviews all copy encountered during the walkthrough against the Voice and Tone Guide and the 18-point design checklist in real time. And a client stakeholder with decision-making authority, who must be present for the sign-off conversation at the end.

A fifth participant is strongly recommended but not required: someone who represents the emotional experience of the end user, whether that is a researcher, an advisor from the ADHD community, or a team member who has been in that role throughout the design process. Their presence is most

valuable during the body doubling introduction and re-entry stages, where the review is assessing something that is genuinely hard to evaluate from the inside.

What the review produces

At the end of the walkthrough, the review team produces a sign-off document with three sections.

The first section records what is working. These are moments where the emotional intent of the design was fully realised in the build, where the feel matched the intent without qualification. This section matters because it establishes a baseline and because it protects resolved decisions from being reopened during the adjustment process.

The second section records what needs adjustment. These are moments where the emotional experience is close but not correct, where a piece of copy is almost right, where a transition is slightly abrupt, where the visible response to the mood check-in is perceptible but not clear enough. Each item in this section is given a category, blocking or post-launch, and an owner.

The third section records anything that blocks launch. These are moments where the emotional intent has not been met in a way that would actively damage the user relationship. The non-negotiables documented elsewhere in this handoff represent the criteria for this section. Any finding that maps to a non-negotiable is automatically blocking unless the review team unanimously agrees otherwise, which requires a written rationale.

The sign-off process

Sign-off requires agreement from the design lead and the client stakeholder. The development lead must confirm that blocking findings can be addressed within the agreed timeline before sign-off is granted. If blocking findings cannot be addressed before the intended release date, the release date moves. The product does not ship with a known violation of a non-negotiable.

Issues are categorised as follows.

Blocking issues must be resolved and re-reviewed before the release candidate is approved. The re-review covers only the affected moments, not the full walkthrough, unless the change made to resolve a blocking issue has plausible knock-on effects elsewhere in the journey.

Post-launch issues are recorded, assigned, and placed on the immediate post-launch backlog. They do not prevent launch but they are not optional work. They have owners and they have target dates agreed before launch, not after.

The sign-off document is dated, signed by the design lead and client stakeholder, and stored alongside the build documentation for the release. It becomes the reference point for the post-launch review cadence.

Post-launch review cadence

Three reviews follow launch. The 30-day review checks whether the emotional experience is holding under real-world use, drawing on early session data, any qualitative feedback received, and the measurable signals defined in the Design Principles document. The body doubling introduction, re-entry experience, and mood check-in response are the primary focus at 30 days, as these are the moments identified during the design phase as most likely to reveal gaps between intent and reality once real users encounter them.

The 90-day review assesses the identity shift, whether the experience is accumulating into the self-knowledge it was designed to produce, and whether any post-launch adjustments made in response to the 30-day findings introduced new issues. The session summary language and longitudinal self-knowledge layer are the primary focus here.

The 6-month review is broader. It considers the full emotional arc of sustained use, whether the product is still meeting users at the re-entry point without shame, whether the body doubling feature has reached the level of warmth identified as a missed opportunity during the design phase, and whether any features added since launch have shifted the product toward the anti-principles documented during the design phase. This review includes a conversation between the design lead, the client, and where possible, users who have been with the product from the beginning.