CX Audit & Roadmap — Sample Deliverable

Prepared for BrightHome Services

Engagement Reference: sample
Generated on August 27, 2026

DRAFT — not yet certified

Table of Contents

Executive Summary

BrightHome Services asked for a diagnosis of why only 38% of first-time customers rebook within 18 months — well below the 55–65% typical among comparable HVAC/plumbing operators — and a prioritized roadmap to close that gap, under two hard constraints: no new dispatch headcount this year, and field technicians already running near capacity through summer peak season. This audit combined journey mapping, VoC/CSAT review, and service-gap analysis against BrightHome's own data (five documented complaint categories, an eight-month-flat 4.1/5 CSAT score, ~900 monthly customer contacts across phone, web, and app) to identify root causes and sequence fixes that fit within those constraints.

The headline finding is that BrightHome does not have a technician-quality problem — it has a silence problem. Verbatims consistently pair praise for the technician with a specific operational complaint in the same sentence ("technician was great, but I waited 3 hours past my window"), confirming the core service delivery is sound while the surrounding experience is undesigned. Four root causes, none requiring added headcount, account for the bulk of the gap: (1) arrival windows are missed and, critically, delays go uncommunicated — the single largest complaint driver at 35% of comment volume; (2) pricing shifts once a technician is already in the home, with no approval checkpoint before work begins (20%); (3) the customer relationship structurally ends at the automated invoice, with zero follow-up to confirm the fix held or to surface the maintenance-plan offer (15% and 10% respectively); and (4) repeat callers are treated as strangers despite service history already sitting in ServiceTitan. A fifth finding reframes the client's own exit-survey data: most non-rebookers say they simply "didn't think of us" rather than voicing dissatisfaction — but this is best read as the behavioral signature of the silence problem, not a competing explanation for it. Customers given no confirmation call, no follow-up, and no reminder of the maintenance plan have been handed no reason to recall BrightHome when the next issue arises.

The five recommendations, all buildable on BrightHome's existing ServiceTitan platform without new hires, are:

- **Automate the arrival-window moment.** Trigger tiered SMS alerts (apology at 30 minutes late, apology-plus-goodwill-gesture at 90+) off existing dispatch timestamps, systematizing the exact mechanism that already converted one delayed job into a five-star review.
- **Close the post-invoice dead zone.** Add a same-day "did it work?" micro-survey and a 30–60 day seasonal check-in — the single biggest structural white space in the current journey, and the most plausible lever against the "didn't think of us" pattern.
- **Add a mandatory pre-work price approval step** inside the existing technician mobile workflow, converting pricing from something done to the customer into something they explicitly confirm.
- **Fix the cheapest, highest-trust gap first**: require dispatch staff to reference service history before triaging any repeat call — a script and training change, not a system build.
- **Rebuild VoC measurement** by decomposing CSAT into stage-specific ratings, adding SMS micro-surveys to fix a skewed 20% email-only response rate, and launching NPS (currently untracked) at 7 days and 11 months to directly test whether recall failure has a specific, fixable cause.

Expected impact, if executed as sequenced: these fixes collectively address roughly 55–75% of coded complaint volume without touching frozen headcount or peak-season field capacity. Because rebook rate is an 18-month trailing metric, the realistic near-term target is movement from 38% toward the 45–47% range within the first year, driven primarily by the arrival-window and follow-up fixes, with the full 55% goal reached as pricing transparency and membership-offer discipline compound in year two. Leading indicators — delay-alert accuracy, post-repair survey response and "no" rates, membership enrollment off invoices, and the new NPS baseline — should move within 60–90 days and are the correct basis for judging progress; leadership should resist reading the lagging rebook-rate number this early, as premature judgment against it is the single greatest risk to sustained execution.

Research & Market Context

Industry and Market Context

The Phoenix metro home services market — spanning HVAC and plumbing repair, maintenance, and installation — is one of the most competitive and fastest-growing in the country. Arizona's population growth (Phoenix-Mesa-Chandler MSA added roughly 60,000+ residents annually over the past several years) combined with extreme summer heat creates structural demand for HVAC repair that is both urgent and recurring. This is a double-edged dynamic for BrightHome: demand is abundant, but so is customer choice, and switching costs for homeowners are low. In an emergency repair category, "who answers the phone fastest and shows up soonest" often outweighs brand loyalty at the point of first contact — which makes the post-service relationship-building phase disproportionately important for retention, since the initial booking decision is frequently driven by urgency rather than affinity.

The competitive landscape includes national roll-up brands (e.g., large multi-market HVAC/plumbing consolidators), regional multi-location players, and a long tail of small independent contractors. Industry consolidation has been accelerating nationally, with private-equity-backed platforms acquiring local HVAC/plumbing operators specifically to professionalize CX operations — dispatch automation, real-time technician tracking, and digital-first communication — as a differentiation lever. BrightHome's mention of a national competitor offering live ETA tracking and technician-photo texts is consistent with this trend: this is now a near-standard feature among well-funded competitors and is becoming a baseline customer expectation rather than a premium differentiator, even though BrightHome's homegrown build has not yet caught up.

Benchmarks and Typical Performance Ranges

Repeat/retention rates. *(Assumption, industry benchmark)* Home services businesses with a structured maintenance-plan and follow-up motion typically retain 55–65% of first-time repair customers for a second service within 12–18 months, and top-quartile operators (often those with membership-plan penetration above 25–30% of the customer base) report repeat rates above 65%. BrightHome's current 38% rebook rate sits meaningfully below typical industry performance, suggesting the gap is closable through operational and communication fixes rather than requiring a fundamentally different business model.

CSAT and NPS. Home services CSAT benchmarks (post-service surveys, 1–5 scale) typically cluster between 4.3–4.6/5 for well-run regional operators; BrightHome's 4.1/5, flat over eight months, is below-average but not alarmingly so — it signals a "good technician, inconsistent operations" pattern rather than a fundamentally broken service experience. This is corroborated by the verbatims ("technician was great, but..."). On NPS, *(assumption, industry benchmark)* trade services companies with strong post-service follow-up and transparent pricing commonly score in the 50–70 range (on a -100 to +100 scale); companies with unmanaged post-service silence and pricing friction — BrightHome's current pattern — typically fall into the 20–40 range.

We recommend establishing an NPS baseline in this engagement's roadmap phase, since BrightHome does not yet track it.

Arrival-window adherence. *(Assumption, industry benchmark)* Leading home services operators using real-time dispatch and GPS-based routing typically hit 85–90%+ on-time arrival-window adherence, with automated customer notifications when delays occur. BrightHome's ~35% of survey comments citing missed windows suggests adherence is materially below this benchmark — likely in the 60–70% range *(assumption, based on complaint volume, not directly measured)* — and, critically, that proactive communication about delays is largely absent rather than the delays themselves being uniquely severe.

Response rates and channel mix. A 20% post-service survey response rate is within normal range for email-only surveys (typical range: 15–25%), but firms that add SMS-based micro-surveys (1–2 questions, sent within an hour of service completion) typically see response rates of 35–45%, giving a materially more representative read on satisfaction — an important consideration given that BrightHome's VoC sample may be skewed toward more engaged or extreme-experience customers. Phone-dominant booking (65%) is typical for emergency-repair-heavy home services businesses, though the leading edge of the industry is shifting share toward digital booking (30–40%) as text/app-based scheduling matures — BrightHome's 25% website + 10% app (35% digital, growing) is roughly on pace with this shift.

Pricing transparency. *(Assumption, industry benchmark)* Complaints about "price changing once the technician is in the house" are common in the trade services industry generally (visual/diagnostic-dependent pricing is inherently harder to quote upfront than, say, retail), but best-in-class operators reduce this friction to under 10% of complaint volume through upfront flat-rate pricing menus and mandatory quote-and-approval steps before work begins. BrightHome's 20% is roughly double what a well-instrumented pricing workflow would produce.

Comparable Practices from Leading Firms

Three practices recur among home services operators — both large consolidators and well-run independents — that outperform on retention and satisfaction:

- **Proactive, automated status communication.** Real-time technician tracking, automated "technician is on the way" texts with photo/name, and automatic delay notifications (with an apology and, where appropriate, a small goodwill gesture like a waived fee) are now table stakes among larger competitors. BrightHome's own experience — saving a customer with a proactive apology call and fee waiver — is a small-scale proof point of exactly this mechanism working; the opportunity is to systematize it rather than rely on ad hoc heroics.

- **Structured post-service follow-up cadence.** Leading operators run a follow-up sequence: same-day "did everything work as expected" check (often automated SMS/email), a 30–60 day check-in tied to seasonal maintenance reminders, and explicit membership-plan offers embedded in the follow-up rather than left to the customer to ask about. This directly targets BrightHome's two most-cited gaps: "nobody followed up" (15% of complaints) and "didn't know about the maintenance plan" (10%).

- **Service history visibility at first point of contact.** Because dispatch and phone staff can see full service history in the CRM (BrightHome already has this capability in ServiceTitan), leading firms operationalize a rule that any repeat caller is acknowledged by history before triage — "I see you're calling about the AC unit we serviced in June" — which is a zero-headcount, workflow/training fix rather than a staffing fix. This matters given BrightHome's frozen dispatch headcount and peak-season field capacity constraints: the highest-leverage fixes are process, script, and automation changes, not more people.

Frameworks Anchoring This Analysis

This audit will apply four frameworks to structure findings and prioritize the roadmap:

- **Journey mapping with moment-of-truth identification**, focused on the two inflection points BrightHome has already intuited — the arrival-window moment and the "did it actually work" post-repair moment — validated against VoC and complaint data rather than assumption alone.
- **Service gap analysis (expectation vs. delivery)**, using the classic gaps model (customer expectation, company understanding, service design, delivery, and communication gaps) to separate "we don't know this is a problem" issues from "we know but haven't fixed it" issues — relevant given BrightHome's leadership already correctly suspects arrival-window and follow-up gaps are under-reported in exit surveys.
- **CSAT/NPS/CES triangulation**, since CSAT alone (currently BrightHome's only tracked metric) doesn't diagnose *why* satisfaction is capped at 4.1; adding NPS and effort-based questions in the roadmap phase will sharpen root-cause clarity.
- **Retention/lifecycle segmentation by service line**, treating repair/emergency (55% of revenue), maintenance plans (20%), and new installation (25%) as distinct journeys with different rebook triggers — critical because a single "one CX fix" is unlikely to move the aggregate 38% rebook rate; the roadmap will need line-specific interventions weighted by revenue contribution and rebook elasticity.

Findings & Analysis

Root Cause Analysis

BrightHome's core problem — a 38% first-to-second-visit rebook rate against a target of 55% — is not a single failure but the compounding effect of four distinct gaps, each traceable directly to the intake data. Treating this as "one CX problem" would be a mistake; the data shows at least four separable mechanisms, each with different owners, fixes, and cost profiles.

Root cause 1: Broken promise at the arrival window. The single most-cited issue — 35% of survey comments — is the arrival window not being honored. This is compounded by the fact that BrightHome has "no formal SLA beyond the 2-hour arrival window" itself, meaning there is no fallback commitment (e.g., a communication SLA) once the primary promise breaks. The verbatim "waited 3 hours past my window" and "wish someone had called to confirm before showing up" point to the same underlying issue from two angles: it's not that delays never happen (they will, especially in HVAC peak season with technicians "already near capacity") — it's that delays happen *silently*. This is a communication-design gap, not fundamentally a capacity gap, which matters enormously given BrightHome's frozen dispatch headcount: fixing this does not require more people, it requires automated, rules-based notification triggered off ServiceTitan's existing dispatch data.

Root cause 2: Pricing presented as a surprise rather than a process. Twenty percent of comments cite pricing changing once the technician is already in the home. This is a service-design gap: diagnostic-dependent trades inherently have some quote variability, but the complaint isn't about variability — it's about the *reveal moment*. Nothing in BrightHome's journey map (discovery -> call/book -> dispatch -> arrival -> diagnose and quote on-site -> repair) includes a structured "confirm and approve" checkpoint before work begins. The customer experiences the price change as something *done to them* rather than something they agreed to, which is a trust break at exactly the moment BrightHome is asking them to pay.

Root cause 3: The relationship ends at the invoice. BrightHome's own journey map ends with "we send an invoice by email — no structured follow-up after that unless they call us again." This is the most structurally significant finding in the intake: it means 100% of BrightHome's post-service customer experience is either automated billing or inbound-only. There is no owned mechanism to confirm the fix worked (15% of complaints), no owned mechanism to surface the membership plan (10% of complaints, and notably 20% of BrightHome's revenue), and no owned mechanism to re-engage a customer before a competitor's marketing does. This gap alone plausibly touches three of the five top complaint categories.

Root cause 4: Repeat contact is treated as a new contact. The example of a customer calling back with a repeat issue and receiving "a generic 'schedule another visit' response with no acknowledgment of the prior visit" is a service-design and training gap, not a technology gap — ServiceTitan already holds the service history. This is important because it is the cheapest possible fix in the entire roadmap (a workflow rule + phone script change) and it directly addresses the emotional core of churn: customers don't leave because of one bad visit, they leave because a second contact confirms the first problem wasn't taken seriously.

A useful way to see how these four root causes distribute across BrightHome's own complaint data:

Root cause	% of comment volume citing it	Journey stage affected	Fix category	Headcount required?
Arrival window not honored / not communicated	35%	Arrival	Automation (SMS/ETA triggers)	No
Pricing surprise before work starts	20%	Diagnose & quote	Process/workflow (approval step)	No
No post-service follow-up	15%	Post-repair	Automation + light process	No
Difficulty reaching dispatch at peak	15%	Booking/scheduling	Channel shift (self-service)	No (redirect, not add)
Membership plan not proactively offered	10%	Post-repair / quote	Script + workflow trigger	No

The pattern in the right-most column is the most important strategic finding: every major complaint category BrightHome reported is addressable without adding headcount, which directly satisfies the stated constraint. This reframes the engagement from "we need more people to fix this" to "we need better-instrumented use of the people and systems we already have" — BrightHome already owns ServiceTitan, already has CRM-based service history, and already has a channel (app, at 10% and growing) that can absorb communication load that currently falls on phone-based dispatch.

Quantifying the Gap and the Opportunity

BrightHome's target is to move rebook rate from 38% to 55% — a 17-point gap. To make this concrete: on BrightHome's reported base of roughly 900 customer contacts per month, and treating first-time repair/install customers as the relevant pool, even a partial closure of this gap represents a meaningful number of retained relationships. If BrightHome converts, say, 4,000–5,000 first-time customers annually *(estimate derived from $11M revenue and blended repair/install ticket sizes typical in HVAC/plumbing — not a figure provided directly by the client, flagged as an assumption)*, moving from 38% to even 46–48% (a partial, first-year-realistic milestone toward 55%) would mean 320–400 additional customers rebooking who otherwise would not have. Given that a rebooked customer in this industry typically generates multiple future service calls plus a higher probability of membership conversion, the lifetime value impact compounds well beyond the second visit alone.

It's worth stress-testing BrightHome's own explanation for churn. The client's exit surveys (small sample, ~40 responses) suggest most non-rebookers "didn't think of us" rather than being actively dissatisfied — but BrightHome's leadership already suspects this undercounts the real driver, and the intake data supports that suspicion. "Didn't think of us" is not a competing explanation to the arrival-window and follow-up problems — it is very likely a *downstream symptom* of them. A customer who had a smooth, memorable, well-communicated experience is meaningfully more likely to think of that provider first next time; a customer whose experience ended at a silent invoice, with no confirmation call and no membership mention, has been given zero reason to *think of BrightHome* again. In other words, "didn't think of us" is what dissatisfaction looks like when it's too mild to complain about but still erodes recall. This reframing matters for prioritization: it means the fix isn't a retention *campaign* bolted onto the current journey, it's closing the structural silence in the existing journey.

The CSAT data reinforces this reading. At 4.1/5 and flat for eight months, BrightHome is not in a crisis-satisfaction pattern (technicians are consistently praised — "technician was great," "great fix" — in the same verbatims that contain complaints). This is a classic **service gap pattern where the core service delivery is sound but the surrounding experience design is incomplete** — exactly the "good technician, inconsistent operations" pattern flagged in the research benchmarking. This is actually good news operationally: BrightHome does not need to retrain or re-hire its technical workforce, which is the most expensive and slowest lever available. It needs to fix the scaffolding around the technician visit.

Prioritization Logic

Given the constraint set — no new dispatch headcount this year, and field technician capacity already stretched in summer peak — prioritization must favor fixes that are (a) automatable or workflow-based rather than headcount-based, (b) targeted at the highest-volume complaint drivers, and (c) capable of being layered onto BrightHome's existing ServiceTitan investment rather than requiring new platform selection (which is explicitly out of scope for this engagement).

Scoring each root cause against complaint volume, revenue-line relevance, and implementation load without added headcount:

Initiative	Complaint volume addressed	Revenue lines touched	Headcount impact	Relative effort	Priority
Automated ETA/ delay notifications	35%	Repair (55%), Install (25%)	None (automation)	Medium	1
Structured post-service follow-up sequence	15% directly, plus recall/"didn't think of us"	All three lines	Low (mostly automated)	Medium	2

Pre-work price confirmation step	20%	Repair, Install	None (process)	Low	3
Repeat-contact recognition protocol	Not separately measured, but tied to trust/ loyalty	Repair (55%)	None (script/ workflow)	Low	4
Membership plan proactive offer trigger	10%	Maintenance (20%)	None (script trigger)	Low	5
Dispatch phone-load relief via app/web self-service	15%	Booking stage, all lines	None (channel shift)	Medium-High	6

The ETA/delay-notification fix ranks first for two reasons beyond raw complaint share: it is the single moment BrightHome's own data identifies as a "biggest swing point," and it has an internal proof point already — the technician who waived a trip fee and proactively apologized converted a service failure into a five-star review. That is a sample size of one, but it is directionally powerful evidence that the *mechanism* (proactive acknowledgment + small goodwill gesture) works on BrightHome's actual customers; the roadmap's job is to make that reactive save into a systematic, automated default rather than a manager's judgment call.

The post-service follow-up sequence ranks second because it is the only fix that plausibly touches all three revenue lines simultaneously and directly targets both the explicit 15% "no follow-up" complaint and the implicit "didn't think of us" churn driver — arguably the highest-leverage single build in the roadmap, since it converts a currently dead post-invoice period into BrightHome's best remaining opportunity to influence the 55% target.

Price confirmation and repeat-contact recognition rank just behind: both are lower-effort than automation builds (closer to workflow rule and script changes than system builds) but address smaller, though still meaningful, complaint shares.

Dispatch phone-load relief ranks last not because it's unimportant, but because it directly collides with the frozen-headcount constraint and BrightHome's channel data suggests a workaround already exists: the app (10% and growing) and website booking (25%) can absorb overflow demand if BrightHome actively routes peak-hour callers toward self-service confirmation and rescheduling rather than solving this by adding phone capacity.

Expected Impact

Reasoning from BrightHome's own numbers: the arrival-window and follow-up fixes together address roughly 50% of coded complaint volume (35% + 15%) and are the two moments BrightHome's leadership has already identified — through direct experience, not just survey data — as the biggest swing points in the relationship. If these two fixes alone reduce "silent" service failures materially, and given that this category of fix in comparable operators (per the automated-communication and follow-up-cadence practices cited above) is associated with rebook-rate improvements in the high single digits to low double digits *(estimate, based on comparable practice patterns, not a guaranteed outcome)*, a realistic first-year milestone is closing roughly a third to half of the 17-point gap — i.e., moving from 38% toward the 45–47% range — with the full 55% target reached as pricing transparency, membership-offer discipline, and repeat-contact handling compound on top of the first two fixes in year two. This phased expectation, rather than a single-year 55% commitment, is deliberate: it matches the effort/

impact sequencing in the prioritization table and respects the stated constraint that no new operational capacity is available to accelerate delivery beyond process and automation changes this year.

Journey mapping

Journey Mapping

The current-state journey, as described in intake, has six stages: **discovery -> booking -> dispatch/wait -> arrival -> diagnose/quote -> repair -> invoice (end)**. Mapping this against the emotional and operational reality at each stage — using the complaint data, verbatims, and the "biggest swing points" BrightHome's leadership already identified — surfaces a journey that is well-designed for *service delivery* but almost entirely undesigned for *relationship continuity*. The map below treats each stage as a discrete moment with its own expectation, current experience, and gap, because the root-cause analysis showed these are separable problems requiring separable fixes, not one undifferentiated "CX issue."

Current-State Journey Map (All Segments)

Stage	Customer expectation	What actually happens today	Emotional state	Evidence
Discovery	Find a trustworthy provider fast (often mid-emergency)	Google/referral-driven; no CX touchpoint issue here	Urgent, anxious	Not a complaint driver — journey starts fine
Booking (phone 65% / web 25% / app 10%)	Get a fast, easy confirmation	Phone: can be hard to reach at peak. Web/app: functional but underused	Mildly frustrated if phone-bound	15% of complaints cite dispatch phone access
Dispatch/wait (1–3 days non-emergency)	Be kept informed while waiting	Silence between booking and arrival day; no confirmation call	Passive, low anxiety unless plans conflict	"Wish someone had called to confirm before showing up"
Arrival (2-hr window)	Window is honored, or told promptly if not	~35% of comment volume reports missed windows with no proactive notice	High frustration, feels disrespected	"Waited 3 hours past my window"
Diagnose & quote (on-site)	Clear price before work starts	Quote can shift once tech is already inside the home; no formal approval checkpoint	Distrust, feels cornered	"Price felt like it changed once he was already in my house"
Repair / install	Job done right, technician competent	Consistently praised — this stage is not broken	Positive	"Technician was great," "great fix"
Invoice / post-service	Confirmation the fix worked; awareness of other services	Automated email invoice; zero structured follow-up unless customer calls back	Neutral-to-forgettable; erodes recall over time	"Nobody followed up," "didn't know you offered a maintenance plan"

Repeat contact (if issue recurs)	To be recognized as a returning customer	Generic "schedule another visit," no acknowledgment of service history despite ServiceTitan holding it	Feels like a stranger, trust break	Client's own example: lost customer after unacknowledged repeat call

Two things jump out when the journey is laid out this way. First, the map has a **structural dead zone**: everything from "invoice sent" onward is either automated billing or purely reactive (inbound-only if the customer calls). Second, the map has exactly **one owned proactive moment today** — the arrival-window call — and it is executed inconsistently rather than systematically, which is precisely why the same failure mode (silence at the window) generates both BrightHome's best save-story (the waived trip fee, 5-star review) and its most-cited complaint (35% of comments).

Segment-Specific Journey Variants

Because repair/emergency (55% of revenue), maintenance plans (20%), and new installs (25%) behave as distinct journeys with different rebook triggers, the map needs a segment lens rather than one generic path.

Segment	Distinct journey characteristic	Where it breaks	Rebook trigger opportunity
Repair/emergency (55% rev)	Urgency-driven booking; low brand consideration at entry; highest exposure to arrival-window pain	Arrival window + no "did it work?" check	Confirmation that the fix held is the single highest-leverage rebook trigger — this is the moment that converts "fixed it once" into "trust them again"
Maintenance plans (20% rev)	Should be the most retention-friendly segment (recurring by design), but the journey map shows membership isn't proactively surfaced until the customer asks	Post-repair / quote stage — the ask never happens unless customer initiates	Embedding the membership offer at the *invoice moment*, when trust is highest post-repair, not left to chance
New installs (25% rev)	Highest ticket, highest quote-shift risk (larger jobs, more diagnostic variability), longest relationship runway (equipment needs future maintenance)	Diagnose & quote stage carries outsized risk here because dollar amounts are larger	First scheduled maintenance visit should be pre-booked at installation close, not left to the general 18-month rebook window

The install segment is particularly worth flagging: a $11M-revenue business generating 25% of revenue from full-system replacements has, in the current journey map, no mechanism to convert an installation into a maintenance-plan relationship at the point of sale — which is the single easiest upsell moment in the entire business (the customer just spent the most money and has the most reason to trust BrightHome's competence), and it's currently absent from the map entirely.

Future-State Journey Map: Where to Insert New Touchpoints

The redesign does not add stages — it adds **instrumented moments inside existing stages**, which matters given the no-new-headcount constraint. Every new touchpoint below is either automation triggered off existing ServiceTitan data or a scripted step added to an existing human interaction, not a new role or shift.

Stage	New touchpoint to insert	Trigger mechanism	Owner	Effort
Booking confirmed	Automated confirmation text 24 hrs before appointment	ServiceTitan appointment field -> SMS trigger	Dispatch (config only, no new work)	Low
Day-of, pre-arrival	"Technician en route" text with name/photo + live ETA	ServiceTitan dispatch GPS/ status -> SMS/app push	IT/ops config; dispatch reviews exceptions only	Medium (this is the direct answer to the competitor's ETA-tracking advantage)
Arrival window at risk	Automated delay alert triggered the moment a tech is running >30 min behind, with apology copy and, for delays >90 min, an auto-offered goodwill gesture (fee waiver or discount code)	Rule-based trigger off ServiceTitan job-status timestamps	Dispatch approves the rule, doesn't execute manually	Medium
Pre-work quote	Mandatory "confirm and approve" step — technician presents written/ digital estimate via ServiceTitan mobile before starting repair; customer taps to approve	Workflow rule change in existing ServiceTitan mobile app for technicians	Field ops lead trains techs; no new headcount	Low
Same-day post-repair	"Did everything work as expected?" 1-question SMS sent 2–4 hours after job close	Automated trigger off job-completion timestamp	Automated; dispatch only handles flagged "no" responses	Low-Medium
Invoice moment	Membership plan offer embedded directly in the invoice email/text for non-members, with one-tap enrollment	Template change to existing invoice automation	Admin/marketing, one-time build	Low
30–60 day check-in	Seasonal maintenance reminder + "still working well?" touch, segmented by service type (AC vs. plumbing)	Scheduled automation off service-completion date	Automated; no manual dispatch involvement	Medium

Repeat contact	Mandatory service-history lookup before triage — phone script requires referencing prior visit before offering a new appointment	Script/workflow rule; ServiceTitan history already available on caller ID match	Dispatch training, not new hires	Low

Moment-of-Truth Emphasis

Consistent with BrightHome's own read of the business, the journey map confirms two moments carry disproportionate weight relative to every other touchpoint:

- **The arrival window.** It is the first moment post-booking where BrightHome's promise can be visibly broken, and it is the only moment in the current map with proof that proactive intervention changes outcomes (the trip-fee-waiver save). The future-state map converts this from a manager's judgment call into a rules-based default triggered automatically once a technician's job-status timestamp indicates a delay.
- **The "did it actually work?" moment.** This moment doesn't exist in the current journey map at all — it's a genuine white space, not an underperforming touchpoint. Its absence is arguably the most fixable, lowest-cost gap in the entire map: a single automated SMS 2–4 hours post-job-close, requiring no dispatch bandwidth except handling the minority of "no" responses.

Implementation Note for Monday Morning

BrightHome's leadership should treat this journey map as the shared artifact for sequencing the roadmap (detailed separately), but the immediate action from the mapping exercise itself is this: **pull the ServiceTitan job-status and appointment-timestamp fields this week** and confirm which of the touchpoints above (arrival-window trigger, post-repair SMS, invoice-embedded membership offer) can be built on existing automation tooling versus require a lightweight integration. Since platform implementation is out of scope for this engagement, the roadmap chapter will sequence these by build complexity — but the map itself already tells BrightHome's ops lead which stages to instrument first: arrival window and post-repair confirmation, because they are the two moments the data — both BrightHome's and comparable-firm benchmarks — identifies as doing the most damage to recall and rebook likelihood when left silent.

VoC & CSAT/NPS review

VoC & CSAT/NPS Review

BrightHome's current voice-of-customer infrastructure captures a partial, skewed, and single-dimensional view of the customer experience. This section audits what the existing VoC stack does and does not tell BrightHome, quantifies the reliability gaps in the current CSAT number, and lays out a specific measurement redesign — including introducing NPS, which BrightHome does not currently track — sized to fit within the no-new-headcount constraint.

What the Current VoC Stack Actually Captures

BrightHome relies on three sources: a post-service email survey (~20% response rate), Google reviews, and complaints called into dispatch. Each has a distinct and significant blind spot.

VoC source	Volume/reach	What it reliably tells you	What it systematically misses
Post-service email survey	~20% response rate on ~900 monthly contacts (~180 responses/month)	Directional CSAT trend among engaged responders	The 80% who don't respond — likely skews toward either very satisfied or very dissatisfied, undercounting the "mildly underwhelmed, didn't bother to respond" middle that the exit-survey data suggests is BrightHome's actual churn population
Google reviews	Unknown volume, self-selected	Public reputation signal, useful for the "5-star save" pattern	Only captures the emotional extremes (great save, genuine outrage) — not representative of the median customer, and not tied to specific journey stage
Complaints to dispatch	Reactive, ad hoc	Real-time issue flagging	Systematically undercounts dissatisfaction, since the intake data itself shows most churned customers don't complain — they just don't rebook

The core problem is structural: **all three sources are opt-in and skew toward the vocal minority**, while BrightHome's own exit-survey finding — most non-rebookers "didn't think of us" — describes a *silent, low-intensity dissatisfaction* that none of the three current channels is built to detect. This is precisely why CSAT has sat flat at 4.1/5 for eight months despite churn being BrightHome's central problem: the survey is measuring the wrong population's experience relative to the population actually driving the retention gap.

CSAT: What 4.1/5 Actually Means, and Doesn't

At face value, 4.1/5 looks like a moderate, unremarkable score. Read against the benchmark range for well-run home services operators (4.3–4.6/5) and against BrightHome's own verbatims, it tells a specific story: **the technician-delivery component of the score is likely pulling the average up, while the surrounding operational experience is pulling it down** — consistent with verbatims that consistently pair praise ("technician was great," "great fix") with a specific operational complaint in the same sentence.

This matters because a single aggregate CSAT number can't distinguish "the job was done poorly" from "the job was done well but the experience around it was mishandled." BrightHome is very likely in the second category, but the current survey instrument can't prove it, because it almost certainly asks one overall satisfaction question rather than decomposing satisfaction by journey stage.

Recommendation 1: Decompose the CSAT survey into stage-specific ratings.

- **What:** Replace the single overall-satisfaction question with 3–4 short ratings: (1) ease of scheduling, (2) arrival/timeliness, (3) technician quality, (4) pricing clarity. Keep it to 4 tap-able ratings plus one open comment field — total survey time under 45 seconds.

- **How:** This is a ServiceTitan survey-template edit, not a new system. Marketing/admin owns the copy change; IT/ops confirms the survey trigger logic is unchanged.
- **Who:** Admin/leadership team (survey owner) + one dispatch supervisor to sign off on wording so it doesn't imply promises dispatch can't keep.
- **Effort:** Low — one to two weeks, no new headcount, no new platform.
- **Expected impact:** Converts CSAT from a single flat number into a diagnostic tool that will almost certainly show a visible split — high technician-quality scores, materially lower arrival-timeliness and pricing-clarity scores — giving leadership monthly, quantified confirmation of exactly where the root-cause fixes (arrival ETA automation, pre-work quote approval) are moving the needle, rather than waiting on the lagging aggregate rebook-rate metric to validate the roadmap.

Fixing the Response Rate Problem

A 20% response rate on a monthly base of ~900 contacts yields roughly 180 responses — a workable sample size for trend-tracking, but not enough to reliably segment by service line (repair vs. maintenance vs. install) or by technician, both of which will matter for accountability once the roadmap's fixes are in flight.

Recommendation 2: Add a 1-question SMS micro-survey sent 2–4 hours post-job-completion, separate from the full email survey.

- **What:** "On a scale of 1–5, did everything work as expected?" — sent via SMS immediately after job-close, tied to the ServiceTitan job-completion timestamp (this is the same automation trigger identified in the journey-map future-state design for the "did it actually work?" moment — the VoC and journey-fix builds should be the same engineering ticket, not two separate builds).
- **How:** Configure as a ServiceTitan-triggered SMS workflow; route any response of 3 or below to a same-day callback queue (existing dispatch staff handle this as a scripted exception, not new volume — expected to be a minority of responses, not a new steady-state workload).
- **Who:** Dispatch supervisor owns exception handling; IT/ops owns the automation trigger.
- **Effort:** Medium — shared build with the journey-map post-repair touchpoint, so incremental effort beyond that build is low.
- **Expected impact:** Based on comparable SMS micro-survey deployments (35–45% response rates vs. 15–25% for email-only), BrightHome should expect response volume to roughly double or better, from ~180/month to 300–400/month, giving a materially more representative satisfaction signal and — critically — creating a real-time save mechanism for dissatisfied customers before they silently churn, systematizing the exact mechanism that worked in the trip-fee-waiver save.

Introducing NPS: Closing the "Why" Gap

CSAT tells BrightHome customers are moderately satisfied; it does not tell BrightHome whether they'd recommend the company or come back voluntarily — which is the actual business question given the churn objective. BrightHome does not currently track NPS at all.

Recommendation 3: Launch NPS at two points in the lifecycle — 7 days post-service and at the 11-month mark (pre-emptive to the 12–18 month rebook window).

- **What:** Standard "0–10, how likely are you to recommend BrightHome" plus one open-ended follow-up ("what's the main reason for your score?"). Two send points, not one: the 7-day score measures service-delivery sentiment; the 11-month score measures brand recall/consideration sentiment specifically to test the "didn't think of us" hypothesis directly rather than inferring it from exit surveys.

- **How:** Build as a ServiceTitan or connected email/SMS automation (same infrastructure as the CSAT redesign); segment responses by service line (repair/maintenance/install) from day one so line-specific patterns are visible immediately rather than requiring a later re-cut.
- **Who:** Admin/leadership owns the metric; no new headcount — this is a template and automation build, reviewed monthly by the same person currently reviewing CSAT.
- **Effort:** Medium — primarily a build-and-segment task, roughly 3–4 weeks including a testing/QA pass on send logic.
- **Expected impact:** Given the benchmark context (20–40 range for unmanaged post-service silence, 50–70 for firms with strong follow-up and pricing transparency), BrightHome should expect an initial NPS baseline in the low-to-mid 20s to low 30s *(estimate, based on comparable-pattern benchmarking, not a direct measurement)* — a number that, once established, becomes the primary leading indicator for the 55% rebook target, since NPS shifts typically precede rebook-rate shifts by one to two quarters. The real value isn't the number itself; it's the open-ended verbatim data at the 11-month mark, which will directly test whether "didn't think of us" is truly indifference (unfixable by CX) or latent dissatisfaction with a specific, named cause (fixable by the roadmap).

Segmenting VoC by Revenue Line — Currently Absent

None of BrightHome's current VoC sources appear to be tagged or reported by service line. Given that repair (55% of revenue), maintenance (20%), and install (25%) are shown elsewhere in this report to have materially different journeys and different rebook triggers, running VoC as one undifferentiated pool obscures line-specific problems.

Recommendation 4: Tag every survey response (CSAT, micro-survey, NPS) with service line at the point of send, using the ServiceTitan job-type field already captured at booking.

- **What:** No new question is needed — this is a back-end tagging/reporting change, since ServiceTitan already knows whether a job is repair, maintenance, or install.
- **How:** IT/ops adds job-type as a filter field in the survey-response dashboard/export; admin builds three simple monthly views instead of one blended view.
- **Who:** IT/ops (one-time setup), admin/leadership (ongoing monthly review).
- **Effort:** Low — a reporting configuration task, likely under a week of work.
- **Expected impact:** This will very likely surface that installation customers' satisfaction is disproportionately sensitive to pricing-clarity scores (larger dollar amounts, more diagnostic variability per the journey-mapping analysis) while maintenance customers' NPS is most sensitive to whether the plan's value was proactively communicated — turning one flat 4.1 average into three actionable, line-specific scores leadership can act on without guessing which fix matters most to which segment.

Consolidated VoC Measurement Plan

Metric/ instrument	Status today	Change recommended	Owner	Effort	Timeline
Post-service CSAT (overall)	Single question, 20% response, flat 4.1/5	Decompose into 4 stage-specific ratings	Admin/marketing	Low	2 weeks

Post-job micro-survey	Does not exist	Add 1-question SMS, 2–4hrs post-completion	Dispatch + IT/ops	Medium	4–6 weeks
NPS	Not tracked	Launch at 7-day and 11-month marks	Leadership + IT/ops	Medium	4 weeks
Service-line tagging	Not segmented	Tag all responses by job type via existing field	IT/ops	Low	1 week
Google reviews	Passive monitoring	Route low-CSAT/low-NPS respondents to private feedback *before* public review ask (reduces public airing of fixable issues, increases positive review capture from the fixed-issue population)	Admin	Low	2 weeks

The Bottom Line for VoC Strategy

BrightHome's instinct that the arrival-window and follow-up problems are under-reported in its own surveys is correct, and the current VoC architecture is structurally incapable of proving or disproving that instinct because it only samples the vocal minority. The fix is not a bigger survey program — it's a **smarter, better-segmented, lower-friction one**, built entirely on infrastructure BrightHome already owns (ServiceTitan triggers, existing job-type fields, existing email/SMS capability), requiring no new headcount and no new platform. The expected payoff is a VoC system that, within one quarter, gives leadership a real-time, line-segmented, cause-attributable satisfaction signal — replacing a flat 4.1 that has told BrightHome nothing new in eight months with a diagnostic system that will move in direct, attributable response to the arrival-window, pricing-transparency, and follow-up fixes detailed elsewhere in this roadmap.

Service gap analysis

Service Gap Analysis

This section applies the classic service-quality gaps model — the distance between what customers expect, what BrightHome believes it delivers, how the service is designed, how it is actually delivered, and how it is communicated — to BrightHome's five documented complaint categories. The purpose is diagnostic precision: each of BrightHome's five top complaints sits in a *different* gap, which means each requires a different fix, and conflating them (treating all five as "a CX problem") is exactly the mistake that would lead to an underpowered, generically-scoped roadmap. The gaps model also matters because BrightHome's leadership has already voiced a hypothesis worth testing directly — that churn is under-explained by "didn't think of us" and better explained by silent operational failures. Mapping each complaint to a specific gap type confirms which failures are invisible to leadership today (knowledge gaps) versus which are known but unaddressed (design and delivery gaps).

The Five Gaps, Applied to BrightHome's Data

Gap type	Definition	BrightHome complaint(s) mapped here	Evidence
Knowledge gap (customer expectation "` company understanding)	Company doesn't accurately know what customers expect	Membership plan awareness (10%)	"I didn't know you offered a maintenance plan until I asked" — the company assumes awareness that doesn't exist
Standards gap (understanding "` service design)	Company knows the expectation but hasn't designed a process to meet it	Pricing surprises (20%); no follow-up (15%)	No structured "confirm and approve" step exists in the journey design; no follow-up stage exists at all post-invoice
Delivery gap (design "` actual delivery)	A process exists on paper but isn't executed consistently	Arrival window (35%); repeat-contact handling	The 2-hour window is a designed standard, but there is "no formal SLA beyond" it — meaning the standard exists without an enforcement or exception-handling mechanism
Communication gap (delivery "` what's communicated)	The service is delivered (or delayed) but the customer isn't told	Arrival window (35% — overlapping); phone accessibility (15%)	"Wish someone had called to confirm before showing up" — the delay itself may be a delivery gap, but the *silence* about it is a distinct communication failure

Two complaints — arrival window and follow-up — are large enough, and structurally distinct enough, to warrant standalone treatment below. The other three are addressed in the summary table that follows.

Gap 1: Arrival Window (Delivery Gap + Communication Gap, Compounded)

This is the only complaint in BrightHome's data that occupies *two* gap categories simultaneously, which explains why it is both the highest-volume complaint (35%) and the moment BrightHome's leadership has already independently flagged as the biggest swing point. There are two separable failures being reported as one complaint:

- **The window itself is missed** (a delivery gap — the 2-hour standard exists but isn't reliably met, particularly plausible given technicians running "near capacity" in summer HVAC peak).
- **The customer isn't told when it will be missed** (a communication gap — nothing in the current journey requires proactive notification once a delay is detected).

The distinction matters for the fix. BrightHome cannot close the delivery gap without adding field capacity — which is explicitly frozen this year. It *can* close the communication gap immediately, with zero headcount, because ServiceTitan already has job-status timestamps that can trigger an automated delay alert. The intake data proves this works: the technician who called to apologize and waived a trip fee converted a missed window into a five-star review. That is evidence the communication gap, not the delivery gap, is the dominant driver of dissatisfaction here — customers tolerate delays; they don't tolerate silence about delays.

Fix: Rules-based automated alert triggered the moment ServiceTitan flags a technician as >30 minutes behind schedule, with tiered response — a text apology at 30 minutes, a text apology plus an auto-offered goodwill gesture (fee discount code) at 90+ minutes. **Owner:** dispatch supervisor approves rule logic once; IT/ops builds the trigger. **Effort:** medium (one build, reusable infrastructure). **Expected impact:** targets the single largest complaint category (35% of comment volume) without touching field capacity.

Gap 2: No Post-Service Follow-Up (Standards Gap)

This is a pure standards gap: BrightHome has never designed a process for this stage, not because it's hard to execute but because the journey map simply ends at the invoice. This is the most structurally significant gap in the analysis because it doesn't just explain the explicit 15% "no follow-up" complaint — it plausibly explains a meaningful share of the 62% of first-time customers who don't rebook at all, most of whom describe themselves as not dissatisfied, just forgetful ("didn't think of us"). A standards gap of this kind is invisible in complaint data by design: customers don't complain about the absence of something they never expected to receive as a formal touchpoint; they simply don't come back.

Fix: A same-day automated "did it work?" check plus a 30–60 day seasonal check-in, both triggered off existing ServiceTitan completion timestamps — no new process design required beyond specifying the trigger and message content once. **Owner:** admin/marketing builds the templates; dispatch only handles the exception queue (flagged "no" responses). **Effort:** low-medium, shared build with the VoC micro-survey work covered elsewhere in this report. **Expected impact:** closes a standards gap that touches all three revenue lines and is the most plausible lever against the "didn't think of us" churn pattern specifically.

Gap 3, 4, 5: Pricing, Phone Access, Membership Awareness

Complaint	Gap type	Root mechanism	Fix	Effort
Pricing surprises (20%)	Standards gap	No mandatory approval checkpoint exists between diagnosis and work start	Add a "confirm and approve" step in the existing ServiceTitan technician mobile flow — customer taps to approve before repair begins	Low
Dispatch phone access at peak (15%)	Delivery gap (capacity-driven)	Phone lines are the default channel (65% of bookings) even though app (10%, growing) and web (25%) can absorb overflow	Actively route peak-hour callers to self-service reschedule/ confirm via app/web rather than solving via added phone headcount	Medium
Membership plan unclear (10%)	Knowledge gap	Company assumes awareness; nothing in the journey proactively surfaces the offer	Embed membership offer directly in the invoice template with one-tap enrollment	Low

Cross-Segment Gap Exposure

Not every gap hits every revenue line equally, and this matters for sequencing the fixes against BrightHome's stated goal of moving the *blended* 38% rebook rate — since the segments contribute unevenly to both revenue and complaint volume.

Segment	% of revenue	Primary gap exposure	Why
Repair/emergency	55%	Arrival window (delivery + communication); no follow-up (standards)	Highest volume of jobs, highest exposure to scheduling/dispatch strain, and the segment where "did it actually work?" matters most (repair implies something was broken)
Maintenance plans	20%	Membership awareness (knowledge gap)	Ironically the segment most damaged by a gap that's cheapest to close — the offer isn't being made, not that customers reject it when made
New installation	25%	Pricing surprise (standards gap)	Largest dollar amounts, most diagnostic variability, so the absence of an approval checkpoint carries the highest per-incident trust cost

This distribution is a useful sanity check on prioritization: the repair segment (55% of revenue) is exposed to the two largest gaps in the entire analysis, which is additional confirmation that the arrival-window and follow-up fixes should be sequenced first — they carry the most revenue-weighted exposure, not just the highest raw complaint percentage.

What This Analysis Changes About the Roadmap Sequencing

The gaps model surfaces one conclusion that a flat complaint-percentage ranking would obscure: **the knowledge gap (membership awareness) and the standards gap (no follow-up) are both invisible to BrightHome's current measurement systems** — customers experiencing them are underrepresented in survey and complaint data because, by definition, a gap the customer doesn't know to expect rarely generates an explicit complaint. This is the analytical basis for treating "didn't think of us" not as a separate, unfixable churn category but as the behavioral signature of standards and knowledge gaps operating quietly beneath BrightHome's complaint-tracking radar. Every fix identified above requires no new headcount, consistent with the constraint — but the sequencing logic should prioritize closing the compounded delivery/communication gap (arrival window) and the standards gap (follow-up) first, since together they are both the highest-volume *and* the least self-reporting failures in the current system.

Prioritized CX roadmap

Prioritized CX Roadmap

This roadmap translates the root-cause, journey-mapping, VoC, and service-gap findings into a sequenced, resourced execution plan. It is built around one non-negotiable constraint stated directly by BrightHome: no new dispatch headcount this year, and field technician capacity is already stretched in summer peak. Every initiative below is therefore either an automation build on top of existing ServiceTitan infrastructure or a workflow/script change for existing staff — nothing in this roadmap requires adding a single role.

The roadmap is organized into three horizons — Now (0–60 days), Next (60–120 days), and Later (120–180+ days) — sequenced by the prioritization logic established earlier: complaint volume addressed, revenue-line exposure, and effort relative to BrightHome's frozen-capacity constraint.

Horizon 1: Now (0–60 Days)

These four initiatives target the two moments BrightHome's own data and leadership intuition both flag as highest-leverage: the arrival window (35% of complaint volume) and the pricing-quote moment (20%). Both are addressable through workflow and automation changes to systems BrightHome already owns.

#	Initiative	What	Who owns it	How (build steps)	Effort	Expected impact
1	Automated delay-alert trigger	SMS alert fires automatically when ServiceTitan flags a technician >30 min behind window; escalates to apology + auto-offered fee discount code at >90 min	Dispatch supervisor (approves rule logic) + IT/ops (builds trigger)	Configure rule off existing job-status timestamps; draft two message templates (30-min, 90-min tiers); test on 20 jobs before full rollout	Medium — 3–4 weeks	Directly targets the single largest complaint driver (35% of comments); systematizes the mechanism that already produced a 5-star save
2	Pre-work "confirm and approve" step	Technician presents digital estimate via existing ServiceTitan mobile app; customer taps to approve before repair starts	Field ops lead (trains 32 technicians on the new step)	One workflow rule change in ServiceTitan mobile; 15-minute training add to next all-tech meeting; no new software	Low — 1–2 weeks	Targets 20% of comment volume (pricing surprises); reduces trust breaks in install segment (25% of revenue, highest dollar exposure)

3	Repeat-contact recognition protocol	Mandatory service-history lookup before triage; script requires referencing prior visit before offering new appointment	Dispatch training lead	Update phone script; confirm ServiceTitan caller-ID-to-history match is enabled (it already is); 30-minute refresher training for all 12 dispatch staff	Low — 1 week	Cheapest fix in the entire roadmap; directly addresses the client's own lost-customer example; protects trust at the exact moment it's most fragile
4	Same-day "did it work?" micro-survey	1-question SMS sent 2–4 hrs post-job-completion; "no" responses route to same-day callback queue	Dispatch supervisor (exception handling) + IT/ops (build)	Shared engineering ticket with the VoC redesign — configure as ServiceTitan-triggered SMS workflow off job-completion timestamp	Medium — shares build cost with Item 1's infrastructure	Fills the single biggest white space in the current journey map; creates a real-time save mechanism before silent churn occurs

60-day milestone check: track weekly % of jobs triggering a delay alert vs. actual delivery, and track "no/neutral" rate on the post-repair micro-survey. These are leading indicators — do not wait for rebook-rate movement (which lags by design) to judge whether Horizon 1 is working.

Horizon 2: Next (60–120 Days)

This phase addresses the structural "dead zone" identified in the journey map — everything after the invoice — and the two knowledge/standards gaps (follow-up and membership awareness) that are least visible in current complaint data precisely because customers don't know to expect them.

#	Initiative	What	Who owns it	How	Effort	Expected impact
5	30–60 day seasonal check-in sequence	Automated reminder segmented by service type (AC vs. plumbing), reinforcing that BrightHome is still "on call" for the customer	Admin/ marketing	Scheduled automation off service-completion date; build once, three template variants by job type	Medium — 3–4 weeks	Directly targets the "didn't think of us" churn pattern by re-inserting BrightHome into recall before the 18-month window closes

6	Membership offer embedded at invoice	One-tap enrollment link inside the existing invoice email/ text for non-members	Admin/ marketing	Template edit to existing invoice automation — no new system	Low — 1–2 weeks	Targets 10% of complaints and directly monetizes the maintenance segment (20% of revenue) via a moment of highest post-repair trust
7	Install-to-maintenance handoff	First scheduled maintenance visit pre-booked at installation close, not left to the general rebook window	Sales/ estimators (8 staff)	Add a mandatory close-of-sale step: estimator books the first maintenance touch before leaving the home	Low — script/ process change only	Converts the highest-ticket segment's (25% of revenue) natural trust peak into a locked-in future visit, rather than hoping the customer initiates
8	Peak-hour call routing to self-service	Active redirection of overflow phone callers to app/web reschedule and confirm functions during peak hours	Dispatch supervisor + IT/ops	Update IVR/ hold-message script to promote app/ web; track deflection rate	Medium-High — requires IVR change and message testing	Relieves the 15% phone-access complaint without adding headcount, leveraging app growth (10%, trending up) and web (25%) capacity

120-day milestone check: track membership enrollment rate off invoice emails (baseline currently near 0% proactive), and track install-to-maintenance booking rate at point of sale (baseline: not currently tracked — this itself is a new metric to instrument).

Horizon 3: Later (120–180+ Days)

This phase consolidates measurement and closes the remaining visibility gap — BrightHome cannot manage what it isn't measuring, and the current VoC stack (20% response rate, no NPS, no segment tagging) cannot validate whether Horizons 1–2 are working at the level of rigor a $11M business needs.

#	Initiative	What	Who owns it	How	Effort	Expected impact

9	NPS launch at 7-day and 11-month marks	0–10 recommend score plus open-ended follow-up, segmented by service line	Leadership + IT/ops	Build as ServiceTitan/ connected automation; QA send logic	Medium — 4 weeks	Establishes the missing leading indicator; the 11-month score directly tests whether "didn't think of us" is true indifference or a named, fixable cause
10	Service-line tagging across all VoC instruments	Tag every survey response (CSAT, micro-survey, NPS) by job type using existing ServiceTitan field	IT/ops	Reporting configuration, not a new question	Low — under 1 week	Turns one blended 4.1 CSAT into three line-specific, actionable scores
11	Live ETA tracking with technician photo	Direct answer to the competitor CX gap customers already cite unfavorably	IT/ops (evaluate ServiceTitan's native ETA/ tracking module vs. lightweight add-on)	Scope only in this phase — full build may extend past 180 days depending on ServiceTitan's native capability	Medium-High	Closes the most visible competitive gap; note this may brush against the "platform implementation" exclusion in this engagement's scope — flag for a follow-on project if it requires new tooling beyond ServiceTitan's existing feature set

Consolidated Roadmap View

Horizon	Initiatives	Cumulative complaint volume addressed	Headcount required	Primary revenue lines touched
Now (0–60 days)	Delay alerts, pre-work approval, repeat-contact protocol, post-repair micro-survey	~55% (35% arrival + 20% pricing, with repeat-contact and follow-up overlapping)	None	Repair (55% rev), Install (25% rev)
Next (60–120 days)	Seasonal check-in, membership offer, install handoff, call routing	Additional ~25% (15% follow-up + 10% membership, partial overlap with phone access)	None	All three lines, especially Maintenance (20% rev) and Install (25% rev)

Later (120–180+ days)	NPS, VoC segmentation, ETA-tracking scoping	Measurement infrastructure, not direct complaint reduction	None (11 may require scoping beyond this engagement)	All three lines

Expected Outcomes and Review Cadence

Consistent with the phased expectation set earlier in this report, BrightHome should not expect the full 38% -> 55% rebook-rate move within this roadmap's first two horizons. The realistic target is closing roughly a third to half of the 17-point gap — into the 45–47% range — by the end of Horizon 2 (approximately month 4), driven primarily by Items 1 and 4 (arrival-window and post-repair fixes), which together touch the two moments BrightHome's leadership already identified as the biggest swing points and which collectively address over half of coded complaint volume.

Review cadence: leadership should review the Horizon 1 leading indicators (delay-alert trigger rate, post-repair "no" response rate) biweekly for the first 60 days, then shift to monthly review of CSAT (decomposed by stage), NPS baseline, and rebook rate by segment once Horizon 3's measurement infrastructure is live. Given the no-new-headcount constraint, the admin/leadership team should designate one existing team member — not a new hire — as the roadmap's single point of accountability for tracking these metrics across horizons, since fragmented ownership across dispatch, IT/ops, field ops, and sales/estimators (as this roadmap necessarily requires) is the most likely failure mode for execution slippage.

Risks & Mitigations

Every initiative in this roadmap is designed to work within BrightHome's real constraints — frozen dispatch headcount, peak-season field capacity, and an existing ServiceTitan investment rather than a new platform. But designing around a constraint doesn't eliminate execution risk; it just changes its shape. The risks below are specific to what could actually derail *this* roadmap at *this* company, not a generic implementation-risk checklist.

Adoption risk sits mostly with technicians and dispatch staff, not customers. The pre-work "confirm and approve" step (Item 2) requires 32 field technicians — already running near capacity in summer HVAC peak — to add a step to every job. If technicians perceive this as friction that slows them down during the busiest season, informal non-compliance is likely: they'll skip the approval tap "to save time," which silently defeats the fix and leaves the pricing-transparency gap open while leadership believes it's closed. Similarly, the repeat-contact recognition protocol depends on 12 dispatch staff consistently referencing service history before triage — a habit change that's easy to train but easy to backslide on under call-volume pressure, especially since there's no system that *forces* the behavior (ServiceTitan surfaces the history, but doesn't require the rep to mention it).

The automated delay-alert system (Item 1) carries real technical and reputational risk if built poorly. This is the highest-priority, highest-visibility fix in the roadmap, and it depends on ServiceTitan's job-status timestamps being reliably accurate in real time. If technicians aren't diligently updating job status in the field (a plausible failure mode — status updates are often the first thing skipped when a tech is behind schedule and busy), the automation will either fire false delay alerts (annoying customers who weren't actually going to be late) or fail to fire true ones (the exact silent-failure pattern this fix is meant to solve). Given that this initiative is being systematized from a single anecdote — one successful trip-fee-waiver save — there's a risk of over-generalizing: the goodwill-gesture tier (fee discount at 90+ minutes) could become a costly, expected entitlement if customers learn that a scripted apology plus discount is the automatic response to any delay, rather than a genuine service recovery gesture.

Cost and margin risk is concentrated in the goodwill-gesture mechanism and the ETA-tracking scoping (Item 11). The fee-waiver/discount trigger, if it fires broadly during peak season (when delays are most likely precisely because field capacity is tightest), could erode margin on repair jobs at scale in a way that was never true when it was a manager's discretionary call. Separately, Item 11 (live ETA tracking with technician photos) is explicitly flagged in the roadmap as potentially exceeding this engagement's scope (platform implementation is excluded) — there's a real risk that BrightHome's leadership, having seen this section of the roadmap, assumes it's included in current budget/effort and gets a cost surprise if ServiceTitan's native capability doesn't cover it and a third-party integration is required.

Data and systems limitations could undercut the entire VoC and automation strategy. The whole roadmap assumes ServiceTitan's job-status, appointment, and completion timestamps are clean, consistently populated, and reliable enough to trigger customer-facing automations. This has not been verified — it's an assumption carried through from the intake data, not a confirmed technical audit. If timestamp data is inconsistent (which is common in field-service CRMs when techs are rushed), every automation in Horizon 1 and 2 is only as good as data hygiene that hasn't been tested.

Competitive response is a slower-moving but real risk. The national competitor already offers live ETA tracking and photo-based technician updates — the one CX capability BrightHome customers cite unfavorably in comparison. Because Item 11 is pushed to Horizon 3 and may extend past this engagement's 180-day window, BrightHome risks a 6+ month period where the competitive gap on this specific, visible feature remains open while other fixes are landing. If the competitor also improves follow-up or pricing transparency in that window, BrightHome's roadmap may be closing yesterday's gap while a new one opens.

Organizational resistance and fragmented ownership is the most likely single point of roadmap failure. This roadmap deliberately spans dispatch, field ops, sales/estimators, IT/ops, and admin/marketing — five different functional owners across eleven initiatives, in a 58-person company with a lean 6-person admin/leadership team. Without a single accountable owner (as flagged in the roadmap's closing section), the most probable failure mode isn't any one initiative failing outright — it's several initiatives quietly stalling in the 60–120 day range as each function treats cross-functional coordination as someone else's job during a season when everyone is already stretched.

Risk	Likelihood	Impact	Mitigation
Technicians skip the pre-work approval step under time pressure	High	Medium — reopens the pricing-transparency gap invisibly	Build the approval tap into the ServiceTitan mobile flow as a required field (can't proceed to invoice without it), not an optional courtesy step; reinforce in the 15-minute training with a specific "why" tied to the client's own pricing complaints

Dispatch staff backslide on repeat-contact script under call volume	Medium	Medium — undermines the cheapest, highest-trust fix in the roadmap	Spot-check via call recordings monthly (already available if ServiceTitan/phone system logs calls); tie to a simple team-level (not individual) recognition metric rather than punitive tracking, given the frozen-headcount, no-added-pressure constraint
Inaccurate/inconsistent job-status timestamps undermine delay-alert accuracy	Medium-High	High — false or missed alerts damage trust in the exact moment meant to build it	Run a 2-week data-quality audit on ServiceTitan timestamp accuracy before full rollout (test on the 20-job pilot already planned in Item 1); if hygiene is poor, add a lightweight tech reminder step rather than launching customer-facing automation on unverified data
Goodwill-gesture (fee waiver/discount) becomes an expected entitlement, eroding margin	Medium	Medium — direct cost impact concentrated in peak season	Cap the auto-offered discount to a fixed low-dollar amount and track redemption rate monthly; convert to manager-approval-required if redemption rate exceeds a set threshold (e.g., >15% of repair jobs) rather than leaving it fully automated indefinitely
ETA-tracking build (Item 11) exceeds this engagement's scope and budget assumptions	High	Medium — expectation/ cost mismatch with leadership	Explicitly flag Item 11 as a scoping exercise only within this engagement, not a committed build; present ServiceTitan-native-vs-third-party cost options to leadership before month 4 so it can be budgeted as a distinct follow-on decision

Competitor closes other CX gaps while BrightHome's ETA-tracking response lags	Every roadmap item in this engagement is designed to be executable without new headcount or new platform spend, but "designed to" and "will in practice" are different things, and BrightHome's actual constraints — a frozen dispatch budget, a summer-peak field workforce already stretched thin, and an owned but under-utilized ServiceTitan system — create specific, foreseeable ways this roadmap could underperform even if every recommendation above is directionally correct. The risks below are grounded in the intake data itself, not generic implementation-risk boilerplate.

Adoption risk among field technicians and dispatch staff. The pre-work "confirm and approve" step (Item 2) and the repeat-contact recognition protocol (Item 3) both require behavior change from people who are already at or near capacity, particularly technicians during HVAC peak season. A technician racing between jobs in July has every incentive to skip the extra tap-to-approve step if it feels like friction rather than protection — especially since the current culture rewards speed and volume, not a documented pause before starting work. Similarly, dispatch staff handling the phone-load-relief redirect (Item 8) may quietly revert to old habits (just booking the call rather than promoting self-service) if the new script isn't reinforced, because it's the path of least resistance under time pressure. This is a moderate-likelihood, moderate-impact risk: it won't sink the roadmap, but it can silently erode adoption rates to the point where the automation trigger data (delay alerts sent, approval steps completed) looks good on paper while real-world compliance lags.

Automation misfires eroding trust faster than silence did. The delay-alert system (Item 1) and post-repair micro-survey (Item 4) both depend on ServiceTitan job-status timestamps being accurate in real time. If a technician forgets to update job status in the field — plausible given time pressure — the system could either fail to trigger a needed delay alert (recreating the exact silence problem this fix exists to solve) or, worse, fire a false delay alert to a customer who's actually on schedule, which is arguably more damaging to trust than the original problem, since it signals the company doesn't know what's happening with its own technicians. This is a lower-likelihood but higher-severity risk in the first 60 days specifically, before the rule logic has been tuned against real-world data noise.

Cost and scope creep on Item 11 (ETA tracking). This is the clearest point of tension with the engagement's stated exclusions — platform implementation is out of scope, yet closing the competitive ETA-tracking gap is likely to require either a paid ServiceTitan module or a third-party integration, both of which carry real cost and vendor-evaluation time that this roadmap does not currently budget. There's a real risk that leadership, having read the competitive-context research and the "customers have mentioned this positively" framing, expects Item 11 to be a near-term build rather than a scoping exercise, and either gets frustrated by the pace or greenlights spend without proper vendor comparison. This is high-likelihood (expectation mismatches around the flashiest, most visibly competitive item are common) and medium-impact — it risks budget surprise and leadership confidence in the roadmap's pacing, not the roadmap's core mechanics.

Competitive response accelerating while BrightHome executes incrementally. The national competitor already offers live ETA tracking and photo-based technician updates today. BrightHome's roadmap sequences its own answer to this gap into Horizon 3, deliberately, because the higher-complaint-volume fixes (arrival-window communication, post-service follow-up) come first. That sequencing logic is sound on the data, but it means BrightHome will spend roughly four to six months with a widening perception gap on the one CX feature customers already compare unfavorably in the competitor's favor. If the competitor also fixes pricing transparency or adds follow-up sequencing in that window, BrightHome's differentiation window narrows further. This is medium-likelihood, medium-impact: Phoenix's market growth means demand isn't disappearing, but repeat-business consideration — the exact metric this roadmap targets — is precisely what a better-resourced competitor's polish would erode fastest.

Organizational resistance from a "the technician is what matters" culture. BrightHome's CSAT data itself confirms technicians are consistently praised, which is a genuine asset — but it can also breed internal skepticism toward CX process investment ("our people already do good work, why are we adding steps?"). Sales/estimators may resist the install-to-maintenance handoff (Item 7) if it's perceived as extending the sales close or diluting focus on the immediate transaction, particularly during install season when volume and speed are directly tied to their compensation. This is a moderate risk in likelihood, and its impact is concentrated specifically on the initiatives requiring a *new step* in an existing human workflow (Items 2, 3, 7) rather than the purely automated ones (Items 1, 4, 5, 6), which are comparatively resistance-proof because they don't require anyone to remember to do something extra.

Measurement lag creating a false "it's not working" read. Because rebook rate is an 18-month trailing metric, and even NPS shifts are expected to precede rebook shifts by one to two quarters, there's a real risk that leadership — accustomed to judging initiatives by the business's headline number — loses confidence in the roadmap around month 3–4 if the blended rebook rate hasn't visibly moved, even though leading indicators (delay-alert compliance, micro-survey "no" rates, membership enrollment off invoice) are trending correctly. This is a high-likelihood risk given how naturally impatient a churn-focused mandate tends to make leadership, and its impact is significant because premature program abandonment — pulling resources off Horizon 2 initiatives because Horizon 1 hasn't "proven itself" by a lagging metric — would be the single most damaging outcome available, since it would recreate exactly the fragmented, inconsistent execution pattern that caused the original 38% rebook problem.

Data/systems limitation: ServiceTitan configuration constraints. This entire roadmap assumes ServiceTitan's existing automation and mobile-workflow capabilities can support delay-trigger rules, mandatory approval steps, and SMS survey sends without a costly upgrade tier or third-party add-on. If BrightHome's current ServiceTitan license or configuration doesn't natively support one or more of these triggers (a real possibility depending on subscription tier), several Horizon 1 items could face unexpected technical blockers discovered only once implementation begins — a scope-out-of-audit risk, since this engagement's exclusions mean platform capability confirmation wasn't part of the audit itself.

Risk	Likelihood	Impact	Mitigation

Field/dispatch staff quietly skip new steps (approval tap, repeat-contact script) under time pressure	Medium	Medium	Build compliance into existing performance check-ins, not a new audit process; track approval-step completion rate as a visible team metric, not a punitive one	
ServiceTitan timestamp inaccuracy causes false or missed delay alerts	Low-Medium	High (in first 60 days)ServiceTitan timestamp inaccuracy causes false or missed delay alerts	Low-Medium	High (in first 60 da

Beyond what fits cleanly in that table, several risks deserve fuller discussion because their likelihood and impact are shaped by specifics of BrightHome's situation rather than generic implementation friction.

Adoption risk: field technicians treating the pre-work approval step as a nuisance. The pre-work "confirm and approve" step (Item 2) asks 32 technicians who are already near capacity in summer peak to add a step to every job. The realistic failure mode isn't outright refusal — it's quiet erosion, where technicians skip the tap-to-approve step on straightforward jobs "to save time," and the habit reappears exactly where it matters most: larger, more diagnostically ambiguous jobs (installs, 25% of revenue) where pricing surprises do the most trust damage. This risk is moderate in likelihood and moderate-to-high in impact, because if the approval step erodes silently, BrightHome will believe the pricing-transparency gap is closed when it isn't, and the CSAT/NPS data won't catch it quickly since pricing-clarity is only one of several decomposed ratings. The mitigation is to make the approval-step completion rate a visible, non-punitive metric reviewed at existing team meetings (not a new compliance program) and to have field ops leadership specifically ask about install jobs during the first 60 days, since that's where skipping the step is most costly.

Cost overrun risk: automation builds taking longer or costing more than a "config change" framing suggests. Several initiatives in this roadmap (delay-alert triggers, post-repair micro-survey, NPS builds) are described as configuration or template work rather than new platform builds, which is accurate relative to the explicitly excluded "platform implementation" scope — but ServiceTitan's native automation and workflow-rule capabilities may not support every trigger condition described (for example, the tiered 30-minute/90-minute delay escalation, or GPS-based live ETA tracking in Item 11) without a paid add-on module or a lightweight third-party integration. This is a moderate-likelihood risk because BrightHome's intake data doesn't specify which ServiceTitan tier or modules are currently active, and it carries moderate impact — not because any single build is expensive, but because if two or three "low effort" items turn out to require paid modules, the cumulative unplanned cost could strain a business that has explicitly frozen budget for headcount and is likely budget-conscious generally. The mitigation is procedural: before Horizon 1 begins, IT/ops should run a one-week technical discovery sprint auditing exactly which ServiceTitan automation, SMS, and reporting capabilities are already licensed and active versus require an upgrade — this was flagged as a "monday morning" action in the journey-mapping chapter and should be treated as a hard prerequisite gate before Item 1 and Item 4 move from planning into build, not a parallel-track nice-to-have.

Data/systems limitation risk: VoC redesign produces a smaller, more alarming-looking signal before it produces a better one. Recommendation 2 in the VoC chapter (adding SMS micro-surveys) is expected to roughly double response volume and surface a more representative — likely lower — satisfaction signal, since it will reach the "mildly underwhelmed, didn't bother to respond" customers the current 20%-response email survey systematically misses. This is a near-certain outcome (high likelihood) and carries a specific organizational risk: leadership may see CSAT or the new NPS baseline appear to *drop* in the first reporting cycle after rollout and misread this as a sign the CX fixes are failing, when it's actually a sign the measurement finally works. The impact of this misread is moderate — it could trigger premature second-guessing of the roadmap or pressure to revert to the old survey instrument. The mitigation is to explicitly brief leadership, before the VoC redesign launches, that an initial score dip is the expected and desired outcome of a more representative sample, and to present the decomposed, stage-specific CSAT breakdown (Item in the VoC chapter) alongside the topline number every month, so leadership can see the same, or improving, technician-quality scores even as the topline shifts to reflect a broader population.

Competitive response risk: the national competitor's ETA-tracking advantage widens before BrightHome closes it. Item 11 (live ETA tracking with technician photo) is sequenced into Horizon 3 (120–180+ days) and explicitly flagged as potentially exceeding this engagement's scope if it requires new tooling beyond ServiceTitan's native capability. This creates a real risk: BrightHome's own intake data shows customers already cite this feature favorably when comparing BrightHome to the national competitor, and the roadmap's own Horizon 1 delay-alert fix — while addressing the *silence* problem — does not fully replicate the polish of live GPS tracking with a technician photo. This is a low-to-moderate likelihood risk in terms of causing near-term churn (BrightHome's own data suggests operational silence, not feature parity, is the dominant driver of the 38% rebook rate), but it is a real strategic exposure if left unaddressed for two quarters, since brand perception gaps compound over time in a market with low switching costs. The mitigation is to explicitly resource the Horizon 3 discovery scoping (Item 11) as a firm deliverable with a decision deadline, not an open-ended "someday" item — BrightHome's leadership should get a clear yes/no on ServiceTitan-native feasibility by month 4, so that if a follow-on platform project is needed, it can be scoped and budgeted before a full year passes without any competitive response.

Organizational resistance risk: fragmented ownership across four departments causes execution to stall. This roadmap deliberately spans dispatch, IT/ops, field ops, sales/estimators, and admin/marketing — a reasonable design given that every fix is process- or automation-based rather than requiring a new team, but it also means no single department "owns" CX outcomes today, and BrightHome's leadership structure (6 admin/leadership out of 58 total staff) is thin at the coordination layer. The likelihood of this becoming a real drag is moderate-to-high, particularly once Horizon 1 and Horizon 2 initiatives are running concurrently and summer peak season is simultaneously stressing field and dispatch capacity — the same season in which several of these builds are meant to launch. The impact is high if it materializes, because a roadmap this dependent on cross-functional automation handoffs (ServiceTitan triggers feeding dispatch exception queues, sales/estimators adding a new close-of-sale step, admin building templates) can quietly stall when no one is accountable for the whole picture, even if each department completes its individual piece. The mitigation, already noted at the end of the roadmap chapter, bears repeating as a risk-specific point: designate one existing leadership team member as the single point of accountability for cross-hor

90-Day Action Plan

90-Day Action Plan

This plan sequences the highest-leverage, no-new-headcount initiatives from the roadmap into a 90-day execution window, deliberately front-loading the two moments already proven to matter most: the arrival window (35% of complaint volume) and the post-service silence that follows every invoice. Every action below traces directly to a recommendation made in the Prioritized CX Roadmap or the VoC/Service Gap chapters — nothing new is introduced here. The 90-day window covers Horizon 1 in full and begins Horizon 2, deliberately excluding anything (like Item 11, ETA tracking) flagged elsewhere as exceeding this engagement's scope.

Days 0–30: Discovery, Data Hygiene, and Quick Wins

The first 30 days are weighted toward two things the roadmap and risk sections both flag as prerequisites: confirming ServiceTitan's technical capability before building customer-facing automation, and executing the lowest-effort, zero-build fixes immediately so BrightHome sees momentum before the larger automation builds land.

Action	Owner	Effort	Milestone (done =)
Technical discovery sprint: audit which ServiceTitan automation, SMS, and reporting features are licensed/ active vs. require upgrade	IT/ops	Low (1 week)	Written capability map confirming feasibility of delay-alert triggers, SMS micro-surveys, and NPS sends
Data-quality check on job-status and completion timestamps	IT/ops	Low (concurrent with above)	Spot-check report comparing ServiceTitan timestamps against dispatch logs on a sample of jobs
Repeat-contact recognition protocol: update phone script requiring service-history reference before triage	Dispatch training lead	Low (1 week)	Script updated and delivered in a 30-minute refresher training to all 12 dispatch staff
Pre-work "confirm and approve" step: build into ServiceTitan mobile as a required field before invoice	Field ops lead	Low (1–2 weeks)	Workflow rule live; 15-minute training delivered at next all-tech meeting to all 32 technicians
Membership offer embedded at invoice	Admin/marketing	Low (1–2 weeks)	Invoice template updated with one-tap enrollment link, live for all non-member invoices
CSAT survey decomposition into 4 stage-specific ratings	Admin/marketing (wording sign-off from dispatch supervisor)	Low (2 weeks)	New survey template live, replacing single overall-satisfaction question

Service-line tagging across CSAT responses	IT/ops	Low (under 1 week)	Monthly dashboard filterable by repair/ maintenance/install

30-day milestone: the four zero-build fixes (repeat-contact script, pre-work approval, membership offer, CSAT decomposition) are live and generating data; the technical discovery sprint has either cleared or flagged blockers for the automation builds scheduled next.

Days 31–60: Automation Builds and Recovery Mechanisms

This phase executes the two automation builds identified as highest-priority — the delay-alert trigger and the post-repair micro-survey — sharing engineering effort since both depend on the same ServiceTitan timestamp infrastructure validated in Days 0–30. It also launches the pilot testing called for in the risk mitigation plan before either automation runs unsupervised on live customers.

Action	Owner	Effort	Milestone (done =)
Build automated delay-alert trigger (30-min apology tier, 90-min goodwill-gesture tier)	IT/ops (build) + dispatch supervisor (approves rule logic and message copy)	Medium (3–4 weeks)	Rule live in test mode; two message templates approved
Pilot delay-alert system on 20 jobs before full rollout	Dispatch supervisor	Low (embedded in above)	Pilot results reviewed; false-positive/false-negative rate documented
Build same-day "did it work?" post-repair SMS micro-survey	IT/ops (build, shared infrastructure with delay-alert trigger)	Medium	Survey triggers automatically off job-completion timestamp; "no" responses route to a callback queue
Route "no"/low-score micro-survey responses to same-day callback	Dispatch supervisor	Low (existing staff, scripted exception only)	Callback queue tested on pilot volume, average response time under 4 hours
Cap and monitor goodwill-gesture (fee waiver/ discount) redemption	Dispatch supervisor	Low (ongoing tracking)	Redemption-rate report established with a defined review threshold (e.g., flag if over 15% of repair jobs)
Full rollout of delay-alert system beyond pilot	IT/ops + dispatch supervisor	Medium	Live across 100% of scheduled non-emergency jobs

60-day milestone: delay-alert and post-repair micro-survey systems are both live at full volume; weekly tracking is in place for delay-alert trigger rate vs. actual on-time delivery, and for micro-survey "no" response rate — these are the leading indicators the roadmap specifies should be reviewed biweekly rather than waiting on the lagging rebook-rate metric.

Days 61–90: Follow-Up Infrastructure, NPS Launch, and Accountability

The final phase begins Horizon 2 within the 90-day window: building the structural "dead zone" fix (post-invoice follow-up) and launching NPS to close the measurement gap identified in the VoC review. It also formalizes the single-owner accountability structure flagged repeatedly in the roadmap and risk sections as the most likely point of execution failure if left informal.

Action	Owner	Effort	Milestone (done =)
Build 30–60 day seasonal check-in sequence, segmented by service type	Admin/marketing	Medium (3–4 weeks)	Three template variants (AC, plumbing, general) live and scheduled off completion date
Install-to-maintenance handoff: mandatory first-maintenance-visit booking at close of sale	Sales/estimators lead	Low (script/process change)	Booking step added to sales close checklist for all 8 sales/estimators
Launch NPS at 7-day and 11-month marks, segmented by service line	Leadership + IT/ops	Medium (4 weeks, can start once delay-alert build stabilizes)	NPS automation live; QA pass complete on send logic; baseline data collection begins
Peak-hour call routing to app/web self-service (IVR/ hold-message update)	Dispatch supervisor + IT/ ops	Medium-High	Updated IVR script live; deflection rate tracked weekly
Route low-CSAT/low-NPS respondents to private feedback before public review ask	Admin	Low	Workflow updated in existing survey/review request sequence
Designate single point of accountability for cross-horizon roadmap tracking	Admin/leadership	Low (organizational decision, not a build)	Named owner reviewing all leading indicators (delay-alert rate, micro-survey scores, membership enrollment, NPS baseline) on a monthly cadence
Scope Item 11 (live ETA tracking) as a distinct decision, not a committed build	IT/ops + leadership	Low (scoping only)	ServiceTitan-native-vs-third-party cost comparison presented to leadership with a go/no-go decision point

90-day milestone: all Horizon 1 initiatives are fully live and generating at least 30 days of leading-indicator data; Horizon 2's follow-up sequence and install handoff are built and running; NPS baseline collection has begun; a single accountable owner is tracking the full initiative set monthly; and leadership has an explicit, budgeted decision point on whether ETA tracking proceeds as a follow-on project.

What Success Looks Like at Day 90

Consistent with the roadmap's phased expectations, BrightHome should not expect the rebook rate itself to have moved meaningfully by day 90 — it is an 18-month trailing metric, and NPS movement is expected to precede rebook-rate movement by one to two quarters. The correct day-90 success signals are entirely leading indicators: delay-alert trigger accuracy holding steady against a validated pilot baseline, post-repair micro-survey response volume meaningfully exceeding the prior email-only survey's 20% response rate, membership enrollment via invoice showing non-zero uptake for the first time, and an NPS baseline established where none existed before. If these move in the right direction, the roadmap is on track; if leadership sees flat leading indicators at day 90, that is the signal to revisit execution (adoption compliance, data hygiene) rather than to conizon accountability, and to make that person responsible for a single monthly cross-functional check-in rather than relying on each department to self-coordinate during peak season.

The table below sequences every action in this 90-day plan directly against the recommendations already established in the roadmap, journey-mapping, VoC, and service-gap chapters. Nothing here introduces a new initiative — it only sequences, assigns, and sizes what has already been recommended, so that BrightHome's leadership has a single execution reference distinct from the strategic roadmap itself.

0–30 Days: Foundation and Highest-Leverage Builds

The first 30 days are deliberately weighted toward discovery, data-hygiene verification, and the two fixes with the highest complaint-volume leverage — the arrival-window communication gap and the pre-work pricing approval step — because both were identified as addressable without new headcount and both carry proof points already (the trip-fee-waiver save, and the client's own recognition that pricing surprises erode trust at the highest-dollar moment).

Action	Owner	Effort	Milestone that marks it done
Run one-week ServiceTitan technical discovery sprint: confirm which automation, SMS, and reporting capabilities are already licensed vs. require upgrade	IT/ops lead	Low (1 week)	Written capability map delivered to leadership, gating go/no-go on Items 1, 4, 9, 11
Draft and test automated delay-alert trigger (30-min and 90-min tiers) on a 20-job pilot	IT/ops + dispatch supervisor	Medium (3–4 weeks, overlapping into day 30–45)	Pilot completed on 20 jobs with manual spot-check against dispatch logs; false-alert rate confirmed acceptable
Build and roll out mandatory pre-work "confirm and approve" step in ServiceTitan mobile app	Field ops lead	Low (1–2 weeks)	Approval-tap field required before invoice generation for all 32 technicians; 15-minute training delivered at next all-tech meeting
Update dispatch phone script to require service-history acknowledgment before triage	Dispatch training lead	Low (1 week)	Script live; spot-check of 10 recorded calls shows history referenced in at least 8
Decompose CSAT survey into 4 stage-specific ratings (scheduling, arrival, technician quality, pricing clarity)	Admin/marketing + dispatch supervisor sign-off	Low (1–2 weeks)	New survey template live and sending on all completed jobs
Brief leadership that VoC redesign will likely show an initial score dip as sample representativeness improves	Admin/leadership	Low (one meeting)	Written briefing shared and acknowledged before any new survey data is reviewed

30-day milestone marker: the ServiceTitan capability map is in hand, the delay-alert pilot has run on at least 20 jobs with an acceptable false-alert rate, the approval step is live for all technicians, and the decomposed CSAT survey is capturing its first responses. None of these require rebook-rate movement to be judged successful — they are process and build milestones, consistent with the roadmap's explicit design to track leading indicators rather than the lagging 18-month churn metric this early.

31–60 Days: Scale the Core Fixes and Close the Post-Service Dead Zone

With Horizon 1's foundational builds tested, this period focuses on full rollout of the arrival-window fix, launch of the post-repair micro-survey (the single biggest white space identified in the journey map), and the start of Horizon 2's follow-up sequence — since the post-invoice "dead zone" was flagged as the most structurally significant gap in the entire audit.

Action	Owner	Effort	Milestone that marks it done
Roll out automated delay-alert trigger company-wide, past the 20-job pilot	IT/ops + dispatch supervisor	Medium	Live on 100% of non-emergency dispatched jobs; weekly tracking dashboard active
Launch same-day "did it work?" post-repair micro-survey (shared build with delay-alert infrastructure)	IT/ops + dispatch supervisor (exception handling)	Medium	SMS live on all job-completion triggers; "no" responses routing correctly to same-day callback queue
Cap and monitor goodwill-gesture (fee discount) redemption rate from delay alerts	Dispatch supervisor	Low (ongoing monitoring, not a new build)	Redemption-rate dashboard reviewed biweekly; escalation rule set if redemption exceeds 15% of repair jobs
Build and launch membership-plan offer embedded in invoice template	Admin/marketing	Low (1–2 weeks)	Live on 100% of non-member invoices; one-tap enrollment tracked
Add mandatory install-to-maintenance handoff step at close of sale	Sales/estimators lead	Low (script/process change)	First scheduled maintenance visit pre-booked for 100% of new installs closed after rollout date
Begin NPS build (7-day and 11-month send points), segmented by service line	Leadership + IT/ops	Medium (4 weeks, spanning into day 60+)	Send logic built and in QA testing; not yet required to be live
Add service-line tagging to all existing VoC instruments	IT/ops	Low (under 1 week)	All CSAT, micro-survey responses taggable and reportable by repair/maintenance/install

60-day milestone marker: delay alerts and the post-repair micro-survey are both running at full volume (not pilot scale), membership offers are embedded and generating first enrollments, the install-to-maintenance handoff is standard practice for new sales, and NPS is in QA ahead of a day-90 launch. Leadership should specifically review: delay-alert trigger accuracy (false-positive/false-negative rate), micro-survey response volume versus the prior email-only survey's 20% response rate, membership enrollment via invoice showing non-zero uptake for the first time, and an NPS baseline established where none existed before. If these move in the right direction, the roadmap is on track; if leadership sees flat leading indicators at day 90, that is the signal to revisit execution (adoption compliance, data hygiene) rather than to conclude the strategy itself is flawed.

61–90 Days: Consolidate Measurement and Relieve Remaining Pressure Points

The final 30 days of this plan complete the measurement infrastructure (NPS launch, seasonal check-in sequence) and begin addressing the lower-priority, higher-effort items — peak-hour phone relief and the ETA-tracking scoping decision — while establishing the ongoing governance structure needed once this audit engagement concludes, since ongoing program management is explicitly out of scope for this engagement and BrightHome will need an internal owner from day 91 forward.

Action	Owner	Effort	Milestone that marks it done
Launch NPS at both 7-day and 11-month send points, segmented by service line	Leadership + IT/ops	Medium	First NPS baseline reported to leadership, broken out by repair/maintenance/install
Launch 30–60 day seasonal check-in sequence (three template variants by job type)	Admin/marketing	Medium	Automation live for all jobs completed 30+ days prior; first send batch confirmed delivered
Update IVR/hold-message script to actively redirect peak-hour callers to app/web self-service	Dispatch supervisor + IT/ops	Medium-High	New script live; deflection rate to app/web tracked for first two weeks
Complete Go/No-Go scoping decision on live ETA tracking with technician photo (Item 11)	IT/ops + leadership	Medium (scoping only, not build)	Written recommendation delivered: native ServiceTitan capability vs. third-party integration, with cost estimate, for leadership decision
Designate one existing leadership team member as single point of accountability for cross-functional CX metric review	Leadership	Low (organizational decision, not a build)	Named owner confirmed; first monthly cross-functional check-in scheduled
Conduct first 90-day retrospective: review all leading indicators (delay-alert accuracy, micro-survey volume/response, membership enrollment rate, install handoff completion rate, NPS baseline) against expectations	Designated CX owner + leadership	Low (one structured meeting)	Retrospective document produced, feeding directly into month-4+ prioritization without requiring rebook-rate movement as proof point

90-day milestone marker: BrightHome exits this window with all Horizon 1 and Horizon 2 initiatives live at scale, an NPS baseline established for the first time, a scoped (though not necessarily built) answer to the competitive ETA-tracizon accountability across the five departments involved, with a standing biweekly checkpoint (not a new committee, just a fixed slot in an existing leadership meeting) where each department reports build status against the roadmap table above.

The table below consolidates this section's actions into the requested 0–30 / 31–60 / 61–90 day structure, tracing every item directly back to the recommendations already made in the roadmap, VoC, and risk-mitigation chapters — nothing here is a new recommendation.

Horizon	Action	Owner (role)	Effort	Milestone marking "done"
0–30 days	Run ServiceTitan technical discovery sprint: confirm which automation, SMS, and reporting capabilities are already licensed vs. require upgrade (pre-requisite gate for Items 1, 4, 9)	IT/ops lead	Low-Medium (1 week)	Written capability map covering delay-trigger logic, SMS send infrastructure, and mobile-approval workflow, shared with leadership
0–30 days	Run 2-week data-quality audit on ServiceTitan job-status/completion timestamp accuracy	IT/ops lead	Low (runs parallel to discovery sprint)	Audit report confirming timestamp reliability threshold before customer-facing automation goes live
0–30 days	Build and pilot automated delay-alert trigger (30-min and 90-min tiers) on 20 jobs	IT/ops + dispatch supervisor	Medium	Pilot complete with manual spot-check against dispatch logs; go/no-go decision for full rollout
0–30 days	Add mandatory "confirm and approve" pricing step to ServiceTitan technician mobile workflow; train all 32 technicians	Field ops lead	Low (1–2 weeks)	Approval-step field is a required (non-skippable) part of job close in ServiceTitan; 100% of technicians trained
0–30 days	Update dispatch phone script to require service-history reference before triage on repeat calls	Dispatch training lead	Low (1 week)	Script update distributed; 30-minute refresher training completed with all 12 dispatch staff
0–30 days	Decompose CSAT survey into 4 stage-specific ratings (scheduling, arrival, technician quality, pricing clarity)	Admin/marketing + dispatch supervisor (sign-off)	Low (1–2 weeks)	New survey template live and sending on next completed jobs

0–30 days	Brief leadership that VoC/CSAT scores may initially dip as response rate broadens — set expectation before rollout	Admin/leadership	Low	Briefing delivered and documented before any new survey instrument launches
31–60 days	Roll out automated delay-alert system at full scale (post-pilot)	IT/ops + dispatch supervisor	Medium	Delay-alert trigger live on 100% of eligible jobs; weekly compliance/ accuracy tracked
31–60 days	Build and launch same-day "did it work?" post-repair SMS micro-survey, with exception routing to callback queue	IT/ops + dispatch supervisor	Medium (shared build with delay-alert infrastructure)	Micro-survey live; response volume and "no" rate visible in weekly dashboard
31–60 days	Cap and monitor goodwill-gesture (fee waiver/ discount) redemption rate; set threshold trigger for manager-approval fallback	Dispatch supervisor	Low	Redemption-rate tracking dashboard live; threshold rule (e.g., >15% of repair jobs) documented and approved
31–60 days	Add service-line tagging (repair/ maintenance/install) to all VoC instruments using existing ServiceTitan job-type field	IT/ops	Low (under 1 week)	Three segmented monthly VoC views replace one blended view
31–60 days	Designate single point of accountability for cross-horizon roadmap tracking; establish biweekly checkpoint in existing leadership meeting	Admin/leadership	Low	Named owner confirmed; first biweekly checkpoint held with all five departments reporting status
31–60 days	Route low-CSAT/ low-NPS respondents to private feedback before public Google review ask	Admin	Low	Routing rule implemented in survey/response workflow
61–90 days	Launch 30–60 day seasonal check-in sequence (segmented by AC vs. plumbing)	Admin/marketing	Medium	Automation live; first cohort of eligible customers receiving check-ins

61–90 days	Embed membership plan offer with one-tap enrollment into invoice template	Admin/marketing	Low	Updated invoice template live; enrollment-click rate tracked from zero baseline
61–90 days	Add mandatory install-to-maintenance handoff step at close of sale	Sales/estimators lead	Low (script/process change)	100% of new installs leave the home with a pre-booked first maintenance visit
61–90 days	Update IVR/hold-message script to promote app/web self-service during peak call hours	Dispatch supervisor + IT/ops	Medium-High	Updated IVR live; call-deflection rate to app/web tracked weekly
61–90 days	Build and launch NPS at 7-day and 11-month marks, segmented by service line	Leadership + IT/ops	Medium (4 weeks)	NPS baseline established for the first time; both send points live and QA'd
61–90 days	Complete ServiceTitan-native vs. third-party scoping decision for live ETA tracking (Item 11)	IT/ops + admin/ leadership	Medium (scoping only, not build)	Written recommendation with cost/feasibility comparison delivered to leadership; explicit yes/no on whether a follow-on platform project is needed
61–90 days	Conduct first structured 90-day retrospective across all five departments involved in the roadmap	Admin/leadership (single accountable owner)	Low	Retrospective document produced, feeding directly into month-4+ prioritization without requiring rebook-rate movement as proof point

90-day milestone marker: BrightHome exits this window with all Horizon 1 and Horizon 2 initiatives live at scale, an NPS baseline established for the first time, a scoped (though not necessarily built) answer to the competitive ETA-tracking gap, and a documented retrospective that separates genuine execution problems from expected measurement lag. Critically, the 90-day marker is not framed around rebook-rate movement — as established in the roadmap and risk sections, that metric lags leading indicators by one to two quarters, and judging this plan's first three months against it would risk exactly the premature-abandonment failure mode flagged as the highest-likelihood organizational risk. Instead, the meaningful proof points at day 90 are entirely leading and process-based: delay-alert trigger accuracy and compliance, post-repair micro-survey response and "no" rates, approval-step completion rates (especially on install jobs), membership enrollment clicks off invoices, install-to-maintenance booking rate at point of sale, and a functioning NPS baseline segmented by service line.

Two structural commitments matter as much as any individual action in this table. First, the technical discovery sprint and data-quality audit in days 0–30 are explicitly sequenced as hard gates before any customer-facing automation goes live — this directly addresses the risk that "config change" framing understates real technical dependency, and it protects against launching delay alerts or micro-surveys on unreliable timestamp data, which the risk analysis identified as capable of doing more trust damage than the silence it's meant to fix. Second, the single-owner accountability structure established in days 31–60 is the direct answer to the fragmented-ownership risk: with five departments touching this roadmap and a six-person leadership team, execution slippage is far more likely to come from coordination gaps than from any one initiative being wrong on the merits, and a fixed biweekly checkpoint inside an existing meeting — not a new standing committee — is deliberately sized to fit a company that cannot add headcount to manage its own improvement plan.

Appendix: Methodology & Assumptions

Appendix: Methodology & Assumptions

This audit was conducted as a desk-based analysis built entirely on the intake questionnaire responses provided by BrightHome Services leadership, supplemented by industry benchmarking and comparable-practice research for the home services sector (HVAC/plumbing, Phoenix metro and national). No independent primary research was conducted as part of this engagement — no customer interviews, no technician shadowing, no direct system audit of ServiceTitan, and no verification of the underlying data BrightHome self-reported (survey scores, complaint percentages, response rates). The analysis should be read as a rigorous synthesis and structuring of client-provided information against established CX frameworks and industry benchmarks, not as independently verified fact-finding. This is consistent with the engagement's defined scope (journey mapping, VoC/CSAT-NPS review, service gap analysis, prioritized roadmap) and its explicit exclusions (platform implementation, ongoing program management).

Methodology. Four frameworks anchored the analysis, as stated in the Research & Market Context section: journey mapping with moment-of-truth identification; the classic service-quality gaps model (knowledge, standards, delivery, communication gaps); CSAT/NPS/CES triangulation; and retention/lifecycle segmentation by revenue line (repair, maintenance, install). Root-cause analysis was performed by cross-referencing BrightHome's five reported complaint categories against its journey-stage description, then testing each against the gaps model to determine whether the underlying failure was one of awareness, design, execution, or communication. Prioritization logic scored each proposed initiative against three explicit criteria: complaint volume addressed, revenue-line relevance, and — critically, given the stated constraints — feasibility without added headcount. All financial and volume estimates were derived from client-reported topline figures (58 staff, $11M revenue, ~900 monthly contacts) rather than independently modeled.

Assumptions made where intake lacked a specific number, clearly labeled throughout the body of the report and consolidated here:

- **Annual first-time customer volume (4,000–5,000/year).** BrightHome did not provide a direct count of first-time customers per year. This figure was estimated by working backward from the $11M revenue figure and typical blended ticket sizes in HVAC/plumbing repair and installation — it is a modeled estimate, not a client-confirmed number, and should be replaced with BrightHome's actual ServiceTitan customer-acquisition count before it's used in any board-level projection.
- **Arrival-window adherence rate (~60–70%).** BrightHome reported that 35% of survey

comments cite missed windows but did not report an actual on-time percentage. The 60–70% range is inferred from complaint volume and industry benchmarks, not a directly measured figure — actual ServiceTitan dispatch data should be pulled to confirm or correct this.

- **Expected NPS baseline (low-to-mid 20s to low 30s).** BrightHome does not currently track NPS at all, so no baseline exists. This range is a benchmark-derived estimate based on comparable firms with similarly unmanaged post-service follow-up and pricing friction, offered only as a planning reference until BrightHome's own NPS launch produces real data.

- **Rebook-rate improvement expectations (movement from 38% to 45–47% in year one).** This is an estimate based on comparable-practice patterns for firms implementing similar automated-communication and follow-up fixes, not a guaranteed or client-validated outcome. It should be treated as a planning milestone, not a committed target.

- **Goodwill-gesture escalation and threshold rules (e.g., flagging redemption above 15% of repair jobs).** BrightHome provided one qualitative example (a single trip-fee waiver that produced a five-star review) but no data on how frequently such gestures should be offered or what redemption rate would be financially sustainable. The 15% threshold is a suggested planning guardrail, not a client-supplied figure, and BrightHome's finance/leadership team should set the actual acceptable threshold based on real margin data this analysis did not have access to.

- **Response-rate improvement from adding SMS micro-surveys (35–45% vs. current 20%).** This is drawn from general industry patterns for SMS vs. email survey response rates, not from any BrightHome-specific pilot data, since no pilot has yet run.

- **ServiceTitan technical capability assumptions.** The entire automation-dependent portion of the roadmap (delay-alert triggers, micro-surveys, NPS builds, mandatory approval steps) assumes BrightHome's current ServiceTitan license and configuration can support these features natively. This was not verified — it is flagged repeatedly in the Risks section as the single largest execution unknown, and a technical discovery sprint is built into the 90-day plan specifically because this assumption has not been tested.

Benchmark sources referenced. Industry figures cited throughout (CSAT ranges of 4.3–4.6/5, NPS ranges of 20–70 depending on follow-up maturity, arrival-window adherence of 85–90%+ for leading operators, survey response-rate differentials between email and SMS, and membership-plan penetration rates of 25–30%+ among top-quartile operators) are drawn from general home services and trade services industry benchmarking knowledge rather than a single named proprietary study, and from observable, publicly reported practices among national HVAC/plumbing consolidators (e.g., live GPS/ETA tracking and photo-based technician notifications, which are referenced as a competitor capability in BrightHome's own intake). Phoenix-Mesa-Chandler MSA population growth figures are drawn from general public demographic reporting on Arizona growth trends. None of these benchmark figures were independently re-verified against primary sources as part of this engagement; they represent directional industry knowledge appropriate for strategic planning, not audited statistics.

Confidence Notes

Several elements of this analysis are meaningfully thinner than others and would benefit from additional client input before being treated as decision-grade:

- **The 38%-to-55% target math and the 4,000–5,000 annual customer estimate** are the least certain quantitative elements in the report. BrightHome should confirm actual annual first-time customer counts directly from ServiceTitan before any board presentation uses the "320–

400 additional retained customers" framing from the Findings section.

- **The exit-survey finding that most non-rebookers "didn't think of us"** is based on a small sample (~40 responses) and was already flagged by BrightHome's own leadership as likely undercounting real dissatisfaction. This report's interpretation — that "didn't think of us" is a downstream symptom of the arrival-window and follow-up gaps — is a plausible, well-reasoned inference, but it remains an inference, not a proven causal link. The 11-month NPS launch recommended in this roadmap is specifically designed to test this hypothesis with real data; until then, it should be treated as the report's most important open question, not a settled conclusion.
- **ServiceTitan's actual technical capability** is an assumption running under nearly every Horizon 1 and Horizon...2 initiative in this roadmap. This report assumes ServiceTitan's existing license tier supports rules-based automation triggers off job-status timestamps, SMS send workflows, and mobile-app required-field logic, but no direct technical audit of BrightHome's current configuration was performed as part of this engagement — platform implementation and configuration review were explicitly out of scope. The Days 0–30 technical discovery sprint is designed precisely to close this gap before any customer-facing automation is built, and it should be read as a genuine unknown rather than a formality.
- **Timestamp and data hygiene quality** inside ServiceTitan is assumed to be sufficient to support customer-facing automation, based on the fact that the system is already in active use for dispatch and CRM. Whether technicians consistently and promptly update job status in the field — particularly under summer-peak time pressure — has not been verified. This is flagged as a load-bearing assumption because the entire value of the delay-alert system depends on it.
- **Estimated first-time customer volume (4,000–5,000 annually)** used to translate the 17-point rebook-rate gap into an absolute customer count was derived from BrightHome's reported $11M revenue and typical blended ticket sizes in residential HVAC/plumbing, not supplied directly by the client. This figure should be treated as directional only; if BrightHome can supply actual first-time-customer counts from ServiceTitan, the sizing in the Findings & Analysis and Roadmap sections should be recalculated with real numbers.
- **Expected NPS baseline range (low-to-mid 20s to low 30s)** and the **expected rebook-rate improvement (high single digits to low double digits from comparable-practice patterns)** are both benchmark-derived estimates, not predictions specific to BrightHome. They are useful for setting realistic expectations and sequencing milestones, but should not be treated as commitments in any client-facing communication about this engagement's expected ROI.
- **Arrival-window adherence rate (~60–70%, estimated)** was inferred from complaint volume (35% of comments citing missed windows) rather than measured directly from ServiceTitan dispatch data. BrightHome has this data natively and should pull actual on-time performance percentages early in the Days 0–30 window — this is one of the easiest assumptions in this report to convert into a verified fact, and doing so would meaningfully sharpen the baseline against which the delay-alert fix's impact is measured.
- **The revenue and margin impact of the goodwill-gesture (fee waiver/discount) mechanism** is discussed qualitatively in the Risks section but was not modeled quantitatively, because BrightHome's intake did not include trip-fee amounts, average repair-ticket size, or current margin structure. The 15%-of-repair-jobs redemption threshold proposed as an escalation trigger is a reasonable starting guardrail based on general service-industry practice, not a figure calibrated to BrightHome's actual cost structure.

Benchmark Sources Referenced

The benchmark ranges cited throughout this report — repeat/retention rates (55–65% typical, 65%+ top-quartile), CSAT ranges (4.3–4.6/5 for well-run operators), NPS ranges (20–40 for unmanaged post-service silence, 50–70 for strong follow-up/pricing transparency), arrival-window adherence (85–90%+ for GPS/automation-enabled operators), survey response rate uplift from SMS versus email (35–45% vs. 15–25%), and pricing-transparency complaint benchmarks (under 10% for flat-rate/approval-gated operators) — are drawn from general home services and trade-services industry performance patterns commonly cited in field-service management, CX benchmarking, and franchise/consolidator operating literature. These are not sourced from a single named study or BrightHome-specific competitive audit; they represent typical ranges observed across the category and are intended to contextualize BrightHome's numbers directionally, not to serve as precise targets. Phoenix-Mesa-Chandler population growth figures referenced in the market-context section are drawn from general, publicly available regional demographic reporting and are approximate. Where this report states a benchmark, it should be understood as "typical of well-run operators in this category," not as a verified statistic tied to a specific named source document — a distinction worth flagging clearly to the client if this report is used in any external-facing capacity.

Confidence Notes

Rated from most to least certain, for the client's awareness in prioritizing what additional input would most improve this analysis:

High confidence — grounded directly in client-supplied data. The five complaint categories and their relative volumes, the current journey map stages, the CSAT trend (4.1/5, flat 8 months), the channel mix (65% phone/25% web/10% app), the segment revenue split (55/20/25), and the specific verbatims are all taken directly from intake and treated as reliable inputs throughout.

Medium confidence — reasonable inference from client data, but not independently verified. The root-cause attribution (e.g., that the arrival-window complaint reflects a communication gap more than a pure capacity gap), the mapping of complaints to the five service-quality gap types, and the prioritization logic all rest on sound reasoning from the data provided, but have not been validated against, for example, a direct technician or dispatch-staff interview, which was outside this engagement's scope. A round of internal stakeholder interviews (dispatch staff, field technicians, sales/estimators) would meaningfully strengthen confidence in whether the proposed fixes are practically executable as designed.

Lower confidence — flagged assumptions requiring client validation. The estimated first-time-customer volume, the estimated arrival-window adherence rate, the expected NPS baseline, and the expected rebook-rate lift from comparable-practice benchmarking are all directional estimates rather than BrightHome-specific figures. Similarly, the causal link between "didn't think of us" churn and the specific operational gaps identified (arrival window, follow-up) is a well-supported inference but is explicitly *un-proven* until the 11-month NPS verbatim data is collected — this is the single most important open question in the entire report, and the client should treat the phased rebook-rate projections (45–47% by month four, full 55% target in year two) as scenario-based planning estimates, not commitments, until that data exists.

Thinnest area overall. The financial/margin modeling around the goodwill-gesture mechanism and the true cost-benefit case for Item 11 (ETA tracking) are the least developed parts of this report, simply because the intake did not include cost data (trip fees, average ticket size, margin by service line) needed to model them properly. If BrightHome can provide basic unit economics in a follow-up conversation, both the risk-mitigation guardrails and the Horizon 3 ETA-tracking business case could be sharpened considerably before final roadmap sign-off.