Go-to-Market Plan — Sample Deliverable

Prepared for Ledgerly

Engagement Reference: sample

Generated on August 27, 2026

DRAFT — not yet certified

Table of Contents

Executive Summary

Ledgerly is preparing to launch Ledgerly Approvals, an automated expense-approval add-on, into its base of 340 existing customers over the next six months, with a target of converting 51 accounts (15%) into roughly $35K in incremental MRR, and using early wins to sharpen win rates on net-new core-product deals. The launch must run on a $60K budget with a 2-person marketing team and no dedicated launch function — a resourcing profile that is thin relative to the ambition of running two distinct go-to-market motions (base expansion and net-new competitive wedge) simultaneously. The core conclusion of this analysis is that the plan is achievable as stated, but only if the team sequences its effort deliberately rather than launching broadly on day one, and only if it protects the single scarcest resource in the plan: sales and customer success bandwidth.

Five findings stand out. First, the 15%-of-340 target is credible against a modeled funnel (40–60 accounts), but it is entirely dependent on the sales and CS team actively pitching the add-on across the full six months rather than treating it as a one-time launch-week event — this, not marketing reach, is the binding constraint on outcomes. Second, the true addressable segment is narrower than "all 340 accounts": modeling suggests only about 150–200 accounts genuinely fit the ICP (multi-department, finance-led structure), meaning the 15% headline goal actually represents a much more aggressive ~28% attach rate against real fit — a distinction the client should track internally even while reporting the simpler headline externally. Third, the client's own pricing advantage is undersold: the real gap against Ramp/Airbase at the minimum-commitment level is roughly 87–90%, not the 40% figure in the intake, and this stronger number should anchor competitive messaging. Fourth, hitting the 51-account target and hitting the $35K MRR target are not the same outcome — a mix skewed toward smaller, easier-to-close accounts could produce as little as $17–20K in MRR despite hitting the headline logo count, so both numbers need to be tracked monthly, not just at the end. Fifth, the $60K budget is best understood as an enablement and content budget layered onto already-owned channels (email list, in-app surface, sales team, organic traffic) rather than a demand-generation budget, and should be weighted accordingly.

Five recommendations follow directly. First, segment the 340-account base into three fit tiers using existing CRM data before any outreach goes out, and concentrate sales effort and messaging emphasis on the true ICP rather than spraying evenly across the base. Second, make sales and CS pitching a standing, tracked part of every renewal and check-in conversation for the full six months — via a simple CRM field reviewed in the existing weekly pipeline call — rather than a month-one push, and back this with a small SPIF ($50–100 per converted account) activated at month three to counter predictable post-launch fatigue. Third, re-anchor all competitive messaging on the corrected ~87–90% minimum-commitment price gap and enforce a strict three-part message hierarchy (speed/audit-trail hook, zero-migration-risk objection-handler, price as competitive closer) so the two-person marketing team doesn't dilute its output across four competing narratives. Fourth, sequence the net-new-logo motion behind base-expansion proof — hold the landing page and heavier competitive positioning until month two, and refresh case studies with real (non-beta) customer data at the month-three checkpoint before pushing harder into net-new deals. Fifth, track blended MRR-per-account alongside raw account count at every monthly checkpoint, with explicit authority for sales leadership to prioritize account quality over quantity if the two goals come into tension.

If Ledgerly executes this sequencing discipline, the expected impact is a credible path to the full $35K MRR target rather than a diluted $17–26K outcome, a defensible board narrative built on two tracked metrics rather than one potentially misleading headline number, and — critically —

a repeatable playbook for future add-on launches that doesn't depend on hiring a dedicated launch team. The risks are concentrated and known: sales bandwidth lapsing after month one, and the team quietly optimizing for logo count over revenue quality. Both are addressed directly in the plan's checkpoints and incentive timing, and both are catchable with four-plus months of runway remaining if the monthly reviews recommended here are actually held.

Research & Market Context

Market Backdrop: The SMB Expense-Management Category

The expense-management software category has matured into a well-understood segment of B2B fintech, but it remains bifurcated in a way that is directly relevant to Ledgerly's launch. At the top end, platforms like Ramp, Airbase, and Brex have spent the last five years pushing upmarket features — multi-entity support, complex approval hierarchies, procurement, and card issuance — bundled into spend-management suites priced with minimum platform fees typically starting at **$1,500–$2,000/month** (client-reported, consistent with published pricing tiers for these vendors). This pricing structure is calibrated for companies with dedicated AP teams and complex org charts, not the 50–500 employee segment Ledgerly serves.

At the lower end, a long tail of point solutions (Expensify, Zoho Expense, and native accounting-suite modules in QuickBooks/NetSuite) handle basic reimbursement but offer only rudimentary, single-step approval logic — and many SMBs, as the intake confirms, simply don't enforce structured approval workflows at all, relying on email or Slack threads with no audit trail. Industry surveys of SMB finance functions (e.g., APQC and GBTA benchmarking studies) consistently find that **manual or ad hoc approval routing is the norm for companies under 500 employees**, with average approval cycle times in the 3–5 day range — closely matching Ledgerly's own beta data point of 4.5 days pre-implementation. This is a useful external validation: Ledgerly's problem framing is not idiosyncratic, it reflects a well-documented gap in the market.

This creates a genuine "middle" that is underserved: companies large enough to have multiple department heads and cost-center complexity, but not large enough to justify a $1,500+/month platform commitment or the implementation overhead that comes with it. **Assumption (labeled):** based on general SaaS category benchmarks, the SMB expense-management software segment in the US/Canada is estimated to be growing at roughly 12–15% CAGR, driven by continued digitization of finance operations post-2020 and increased finance-team scrutiny on spend control amid tighter budget environments. This is an industry-typical estimate, not a client-provided figure, and should be treated as directional context rather than a precise market-sizing input for this engagement.

Segment Economics and Buyer Behavior

Ledgerly's existing base of ~340 customers in the 50–500 employee range represents a segment where the economics of add-on attach are favorable but under-exploited. SaaS benchmarking data (OpenView's SaaS Benchmarks, Vendr pricing data) generally shows that **expansion revenue from existing accounts costs 3–5x less to generate than new-logo acquisition** — a ratio consistent with Ledgerly's own numbers: core-product CAC of ~$2,800 versus a near-zero-CAC expansion motion through existing sales, CS, and email channels. This is a critical benchmark for evaluating the 6-month plan: the $60K budget and 2-person marketing team are thin relative to a typical launch, but they are being deployed against the cheapest possible growth surface (installed base), not a cold-acquisition motion, which materially changes what "success" should cost.

On attach-rate benchmarks specifically: cross-sell/attach rates for SaaS add-on modules launched into an existing base typically range from **10–25% within the first 6–12 months**, depending on the strength of the wedge (how directly it solves an active pain point) and the friction of adoption. Ledgerly's target of 15% attach within 6 months sits squarely in the middle of this range — neither aggressive nor conservative — which is a reasonable benchmark-consistent goal, though it assumes the sales team allocates real bandwidth to it during renewal/upsell conversations rather than treating it as a footnote.

On buyer behavior: the Controller/Head of Finance persona that Ledgerly has correctly identified as the primary buyer is well-documented in B2B fintech research (Gartner's Finance Function studies, various fintech-vendor buyer surveys) as being driven primarily by three purchase triggers: **audit/compliance risk exposure, time reclaimed from manual reconciliation work, and budget-cycle timing.** The intake's stated intent to launch ahead of Q1 budget planning season aligns well with documented buyer seasonality — Controllers and Heads of Finance disproportionately evaluate new tooling in Q4 in order to have budget lines approved before the new fiscal year, making the proposed launch window tactically sound rather than arbitrary.

Comparable Practices: How Adjacent SaaS Companies Have Run Similar Plays

Several patterns from comparable B2B SaaS launches are directly instructive here, even though none are drawn from Ledgerly's own data:

- **Wedge-into-installed-base playbook.** Companies like Gusto (adding time-tracking/benefits modules to its core payroll base) and Ramp itself (which began as a card product and added bill-pay/accounting-close modules) have both used the "sell the second product to the first product's base" motion as a lower-CAC growth lever before investing heavily in net-new-logo acquisition for the add-on. The common thread in these plays is that the add-on is positioned not as a new purchase decision but as "turning on a feature you're already paying to have access to" — which lowers the psychological barrier to a yes and shortens sales-cycle length considerably versus net-new deals.

- **Beta-customer proof points as the primary launch asset.** The pattern of using a small closed-beta cohort (Ledgerly has 8) to generate a specific, quantified before/after metric — here, 4.5 days to under 8 hours — mirrors how companies like Vanta and Drata built early GTM narratives around single hero metrics rather than broad feature lists. This is a proven and appropriately scoped approach for a resource-constrained team: one sharp, credible number typically outperforms a long feature list in outbound and sales-enablement contexts, and it requires minimal creative production, which fits the excluded scope of this engagement.

- **Price-anchored positioning against category leaders.** Positioning at a clear discount to Ramp/Airbase's minimum platform fee (Ledgerly's $150/month minimum versus their reported $1,500+/month) is a well-established SaaS challenger tactic — comparable to how Freshbooks and Wave historically anchored against QuickBooks/Intuit pricing. This works best when paired with an explicit "right-sized for your stage" message rather than a pure discount message, since underpriced-relative-to-leader positioning can otherwise read as "lite" or "less capable" rather than "better fit."

Frameworks Anchoring This Analysis

Three frameworks will structure the segment, positioning, and channel work that follows:

- **Jobs-to-be-Done (JTBD)** will anchor the positioning work, since the buyer's core job — "give me auditable control over departmental spend without adding platform overhead" — is distinct from feature-level differentiation and should drive messaging hierarchy.
- **The Bowling Pin / Wedge Strategy** (popularized in B2B SaaS GTM strategy, notably by Geoffrey Moore's category-entry theory) will frame the segment-prioritization logic: using the existing base as the first "pin" to build proof and case studies before knocking down the adjacent net-new-logo pin using the same story.
- **A simple funnel-stage benchmark model** — awareness (email/in-app reach across 340 accounts) -> consideration (beta proof points, pricing clarity) -> activation (sales/CS-led close) — will be used to sanity-check whether the 15%-in-6-months goal is achievable given the channel mix available, rather than treating it as an assumed given.

These frameworks, combined with the benchmarks above, provide the grounding for the segment prioritization, positioning, and channel recommendations that follow.

Findings & Analysis

Root Cause Analysis: Why This Launch Requires Sequencing, Not a Broad Push

Ledgerly's stated goal — 15% attach (51 of 340 accounts) within 6 months, generating ~$35K incremental MRR, on a $60K budget with 2 non-dedicated marketing people — is achievable, but only if the team resists the urge to treat this as a full product launch. The intake reveals a structural mismatch that is the root cause of most of the risk in this plan: **the ambition level of the goal (a company-wide cross-sell motion plus a net-new-logo wedge) is not matched by the resourcing level (2 marketers, no dedicated launch team, $60K)**. This isn't a flaw in the plan so much as a sequencing problem — the client is trying to do two distinct GTM motions (base expansion and net-new differentiation) with resourcing sized for roughly one.

The second root cause is that the client's own numbers show the existing-base motion is dramatically cheaper and faster to prove than the net-new motion, yet the objective statement treats them as parallel goals ("launch... within existing base... **and** use it as a wedge to win net-new logos"). Core-product CAC is $2,800 via outbound; the add-on's expansion motion runs through channels that are already paid for (sales team salaries, CS team salaries, an email list, in-app real estate) — effectively near-zero marginal CAC. Given this cost asymmetry, the analysis below treats net-new-logo lift as a **second-order, sequenced outcome** of base-expansion proof, not a co-equal Month 1 priority — consistent with the Bowling Pin framework flagged in the research above.

Quantifying the Attach-Rate Target Against the Available Funnel

The client's goal is 51 accounts converted in 6 months. It's worth stress-testing that number against the actual reach of the assets on hand, because the intake provides enough detail to build a real funnel model rather than assume the target is automatically hittable.

Funnel Stage	Channel / Asset	Reach (client-stated)	Realistic Conversion Benchmark	Estimated Accounts
Awareness	Email to 340-account list	340 accounts	~35-45% open rate (typical for active B2B SaaS customer lists)	~130-150 opened

Awareness	In-app announcement	340 accounts (all logged-in users)	~50-60% seen within 30 days	~170-200 seen
Consideration	Sales/CS-led renewal & upsell conversations	340 accounts, but gated by AE/CSM bandwidth	6 AEs + 3 CSMs covering ~340 accounts = ~38 accounts/rep	Realistic active-pitch coverage in 6 months: 60-70% of base
Activation	Close from pitched accounts	~200-240 accounts actively pitched	20-25% close rate for a low-friction, already-adopted-platform add-on	**40-60 accounts**

This model suggests the 15% target (51 accounts) sits at the **midpoint of a plausible range (40–60 accounts)** rather than being either a layup or a stretch goal — which is a useful finding in itself: the number is credible, but it is entirely dependent on sales/CS actually working the full base rather than cherry-picking the easiest 20%. Given that Ledgerly's sales team is 6 people covering $4.2M ARR and 340 accounts, and marketing has no dedicated launch team, the biggest execution risk isn't the marketing motion — it's whether sales/CS treat this as a standing agenda item across renewal cycles for six months, rather than a one-time push in month one. **This is the single highest-leverage operational fix available to the client, and it costs $0.**

The ICP Is Narrower Than the Addressable Base — A Quantified Gap

The intake defines the add-on's true ICP precisely: 100–400 employees, multiple department heads who approve their own team's spend, plus a Controller/Head of Finance who wants auditability. But the *stated* target for the 15% goal is "existing customer base" — all 340 accounts, which spans the full 50–500 employee range of the core product.

This is a meaningful gap. A 50-person company with a single office manager approving expenses has no structural need for multi-step routing — the JTBD ("give me auditable control over departmental spend without adding platform overhead," per the framework established above) doesn't apply to them. If we assume — **labeled assumption, not client-provided** — that Ledgerly's 340 accounts are roughly evenly distributed across the 50–500 employee range (a reasonable default absent segment-level data from the client), then only the accounts in the upper half of that range, roughly 100–500 employees, are strong-fit prospects. That could mean the *true* addressable base for this add-on is closer to **150–200 of the 340 accounts**, not all 340.

This matters directly for the target math: if the real ICP is ~180 accounts, then 51 conversions represents a **28% attach rate against the true addressable segment** — a considerably more aggressive goal than "15% of the base" implies on its face. This reframing doesn't mean the target is wrong; it means the client should track two numbers, not one: total accounts converted (for the board-level 15% headline) and **ICP-fit attach rate** (for a true read on product-market fit within the segment that matters). We recommend the marketing team segment the 340-account list by employee count and department-head structure (using data likely already in the CRM) before the email/in-app push goes out, so effort concentrates on the ~180 accounts most likely to convert rather than spraying evenly across all 340 — a zero-cost targeting fix that should materially improve the close-rate assumption in the funnel table above.

Pricing and Differentiation: A Real Gap, Correctly Identified, But Under-Leveraged in the Plan

The client's positioning logic is sound and well-supported by the research: pricing at $3/seat/month with a $150/month minimum against Ramp/Airbase's $1,500+/month minimum is roughly an **87–90% discount at the entry tier**, not the "40% below" figure quoted in the intake — that 40% figure appears to describe feature-tier pricing comparisons rather than minimum-commitment comparisons. This is worth flagging as a messaging correction: the client's most powerful number (a >85% cost gap at the minimum-commitment level, which is exactly the level that matters to a 150-person company evaluating whether to even consider Ramp/Airbase) is currently undersold relative to the softer "40% below" framing in the intake. Positioning should lead with the minimum-commitment comparison, since it's the number that actually disqualifies the incumbents for this buyer's budget, not the per-seat comparison.

The second differentiation asset — 4.5 days to under 8 hours — is a strong hero metric (13x+ speed improvement) and, per the comparable-practices research above, is exactly the kind of single sharp number that should anchor all outbound, email, and sales-enablement copy given the team's limited creative-production capacity (explicitly excluded from this engagement's scope). The risk is dilution: with only 2 marketing people, there's a real temptation to build multiple messages (speed, cost, zero-migration-risk, audit trail) across different assets. Given the JTBD framing established earlier — auditable control without added platform overhead — we recommend the hierarchy be: **(1) speed/audit-trail as the primary pain-relief hook, (2) zero migration risk as the "why now, why easy" objection-handler, (3) price gap as the closing argument against Ramp/Airbase specifically in competitive net-new deals.** Trying to lead with all three simultaneously, across a 2-person team's limited bandwidth, is a likely root cause of diluted messaging if not addressed explicitly in the positioning brief.

Budget Allocation: Where $60K Should and Shouldn't Go

The $60K budget, spread over 6 months with no media-buying or creative-production scope, is realistically not a demand-generation budget — it's an **enablement and light-content budget** layered onto channels the client already owns for free (email list, in-app, sales team, 4,000 monthly organic visits). This reframing matters because it changes what "success" should look like: the budget cannot and should not be judged against typical CAC-based paid-channel benchmarks, since paid acquisition is out of scope entirely.

Budget Category	Rationale	Estimated Allocation
Sales enablement (battlecards, ROI calculator, case study production support)	Directly supports the highest-leverage channel (6-person sales/CS team) identified in the funnel model	~$20K
Website/landing page + SEO content updates	Leverages existing 4,000 visits/month organic traffic at near-zero incremental cost; supports net-new inbound	~$15K
Email/lifecycle campaign build (segmented by ICP-fit, per finding above)	Directly targets the ~180-account true ICP segment	~$10K
Beta customer case study / testimonial production (light, non-creative-heavy)	Converts the 8-customer beta into repeatable proof assets for sales	~$10K
Contingency / mid-launch adjustment	Given 2-person team bandwidth risk, reserve for reallocation at the 90-day check-in	~$5K

This allocation deliberately weights toward sales enablement over top-of-funnel marketing spend, because the funnel analysis above shows the bottleneck is conversion-through-sales-bandwidth, not awareness generation — awareness is already near-saturated by email + in-app reach into a finite 340-account list.

Prioritization Logic Summary

Ranking the findings by expected impact versus effort, given the constraints:

- **Segment the 340 accounts by true ICP fit before launch** (near-zero cost, highest leverage — turns a 15%-of-340 target into a more honest and achievable 28%-of-180 target).
- **Make sales/CS bandwidth a tracked, standing part of the 6-month plan**, not a month-one event — this is the primary lever behind whether the 40–60 account funnel estimate lands at the low or high end.
- **Re-anchor positioning on the minimum-commitment price gap (~87%), not the blended 40% figure**, to sharpen the against-Ramp/Airbase net-new argument.
- **Sequence net-new-logo wins as a proof-driven second phase**, using base-expansion case studies (from the 8 beta customers plus early adopters) as the asset that makes the wedge story credible — rather than running both motions at full intensity from day one with a 2-person team.

Each of these is a zero-to-low-cost adjustment to sequencing and targeting rather than a request for more budget or headcount — which matters given the hard constraints stated in the intake.

Market & segment analysis

Segment Prioritization: Three Tiers Within the 340-Account Base

The findings above establish that the true ICP is narrower than "all 340 accounts." To make this actionable rather than directional, the base should be split into three explicit tiers using data already available in Ledgerly's CRM (employee count, org-chart complexity where known, current plan tier, renewal date). This segmentation should happen in Week 1, owned by whichever of the two marketing people has better CRM/reporting access, and should take no more than 3-4 hours of analysis time — it is a filtering exercise, not a research project.

Tier	Definition	Estimated Size (of 340)	Priority	Rationale
Tier 1 — Core ICP	100-400 employees, known multi-department structure (multiple cost centers billed on core product), Controller/Finance-lead as primary contact	~120-140 accounts	Highest	Direct match to JTBD; shortest sales cycle; highest close-rate assumption (20-25% per funnel model)

Tier 2 — Adjacent fit	100-400 employees but single-department or unclear org structure; or 50-100/400-500 employees with finance-lead contact	~100-120 accounts	Secondary	Plausible fit but requires a discovery question before pitching ("how many people currently approve expenses on your team?") rather than a direct pitch
Tier 3 — Poor fit	Under 100 or over 500 employees, single approver, no dedicated finance role	~80-100 accounts	Deprioritize	JTBD doesn't apply; pitching here wastes scarce AE/CSM bandwidth identified as the binding constraint

This tiering directly changes how the $60K and the sales team's limited bandwidth should be spent: Tier 1 gets the case-study-backed direct pitch and personalized outbound from AEs during renewal conversations; Tier 2 gets a lighter-touch email/in-app nudge with a qualifying question baked into the copy; Tier 3 gets included in the broad in-app announcement (zero marginal cost) but should not consume AE/CSM meeting time. Given the funnel model's finding that sales/CS bandwidth — not awareness — is the binding constraint, this tiering is the mechanism that prevents reps from spending renewal-call time on Tier 3 accounts that were never going to convert.

Net-New Segment: Where the Wedge Applies and Where It Doesn't

For the net-new-logo motion, the same ICP discipline applies but with one addition specific to competitive positioning. The intake identifies three distinct buyer situations in the market, and each requires a different segment approach rather than one blanket "net-new" motion:

Buyer Situation	Description	Segment Size Signal	Ledgerly's Right of Entry
No workflow today	Approves via email/Slack, no audit trail, "eating the risk"	Likely the largest slice of net-new prospects at 100-400 employees, per category research above	Strongest fit — direct pain-to-solution match; sell against risk exposure, not against a competitor
Overserved by Ramp/Airbase	On a platform too expensive/complex for their stage	Smaller slice, but higher urgency (active dissatisfaction, likely near a renewal decision)	Strong fit — sell against the ~87% minimum-commitment price gap; requires the AE to ask about current platform cost during discovery
Underserved by point solution	On Expensify/Zoho/QuickBooks-native approval logic that's single-step only	Moderate slice	Fit exists but weaker urgency — will require the core Ledgerly product (expense tracking) to be part of the pitch, not the Approvals add-on alone, since this buyer may not already be a Ledgerly customer

Because net-new-logo acquisition runs through the existing $2,800-CAC outbound motion (the 6 AEs/SDRs), the practical recommendation is narrow and specific: **AEs should add two**

qualifying questions to discovery calls starting in Week 2 of the launch — "How do you currently route expense approvals across departments?" and "What are you paying for that today, if anything?" — to sort net-new prospects into the three buckets above in real time, rather than pitching Approvals generically. This costs nothing beyond a script update (part of the sales enablement budget line already allocated) and directly operationalizes the "sequence net-new as proof-driven second phase" recommendation from the findings — AEs should be instructed to lead net-new Approvals conversations with the base-expansion case studies once 8-10 are available (projected month 3), not with the beta-only data available at launch.

Segment-Level Revenue Math: Validating the $35K MRR Target

The client's goal of ~$35K incremental MRR from 51 accounts implies an average of **~$686/month per converted account** — above the $400-$1,200 stated range's midpoint but plausible if Tier 1 accounts (larger, multi-department, more approver seats) convert at a higher rate than Tier 2/3. This is a useful sanity check: if the tiering and prioritization above work as intended and Tier 1 accounts (higher seat counts, hence higher $/account) convert disproportionately, the $35K target is achievable with the low end of the 40-60 account funnel range. If targeting is not disciplined and conversions skew toward Tier 2/3 accounts with fewer approver seats (closer to the $150 minimum), the same 51 accounts could generate closer to $20-25K MRR — a **30-40% shortfall against the goal despite hitting the headline account-count target.** This is the clearest quantified argument for why the tiering exercise above is not optional analysis but a load-bearing part of hitting the revenue number, not just the logo-count number.

Scenario	Accounts Converted	Mix Assumption	Est. MRR
Disciplined targeting (Tier 1-weighted)	45	60% Tier 1 (~$800 avg), 40% Tier 2 (~$400 avg)	~$38K
Even targeting across tiers	51	Roughly even split, blended ~$500 avg	~$26K
Undisciplined (Tier 2/3-heavy, chasing easy yeses)	55	70% Tier 2/3 near minimum (~$300 avg)	~$17K

This table should be shared directly with the sales leadership team as a Monday-morning artifact: it makes the case, in the client's own revenue terms, for why account *quality* within the 15% target matters as much as account *count* — a distinction the current goal statement ("15% of existing customer base... generating roughly $35K in incremental MRR") does not explicitly separate, but which this analysis shows are not automatically the same outcome.

Positioning & messaging

Positioning Statement and Message Architecture

Based on the JTBD framing and the segment tiering established above, Ledgerly Approvals needs one positioning statement that flexes by audience but never contradicts itself across the two motions (base expansion vs. net-new). We recommend the following as the canonical statement, to be finalized by the 2-person marketing team in Week 1 and used verbatim (or near-verbatim) across email, in-app copy, the landing page, and sales battlecards:

> **"Ledgerly Approvals gives growing finance teams the auditable, multi-step approval control of an enterprise spend platform — without the $1,500+/month minimum, the**

migration, or the implementation lift. Built into the Ledgerly you already use, live in beta today, cutting approval turnaround from 4.5 days to under 8 hours."

This single statement is deliberately built to do three jobs identified in the findings: (1) name the JTBD directly ("auditable, multi-step approval control"), (2) preempt the incumbent-platform objection with the corrected, more powerful minimum-commitment gap (not the diluted "40% below" framing), and (3) lead with the hero metric rather than a feature list, consistent with the beta-proof-point pattern flagged in the research section. Every other piece of messaging below is a variant of this core statement, not a separate narrative — a discipline the team should hold given only 2 people are producing all of it.

Message Hierarchy: What Leads, What Supports, What Closes

The findings section flagged a real dilution risk: four legitimate selling points (speed, audit trail, zero migration risk, price gap) competing for attention with no enforced order. The fix is a strict, written hierarchy that both marketing and the 6-person sales team use identically, so that a Controller hears the same story whether they get an email, see the in-app banner, or sit in a renewal call with their AE.

Position	Message	Audience Use	Supporting Proof Point
1. Primary hook	"4.5 days to under 8 hours" — turnaround speed as proxy for control and audit trail	All audiences, all channels — this is the headline everywhere	8-customer beta data (client's own numbers)
2. Objection-handler	"It's already part of Ledgerly — no new vendor, no migration, no new login"	Existing base (Tier 1/2) — addresses the "why now, why easy" question	Zero-migration-risk fact, unique to installed-base motion
3. Competitive closer	"Full approval routing for a $150/month minimum vs. $1,500+/month platforms like Ramp or Airbase"	Net-new deals specifically, especially the "overserved by Ramp/ Airbase" buyer bucket identified in the segment table above	Corrected 87-90% minimum-commitment gap

This ordering is not arbitrary — it mirrors the funnel logic in the findings: awareness/ consideration assets (email, in-app, landing page) should lead with #1 because it's the broadest hook across all three tiers; sales conversations with existing accounts should move to #2 almost immediately because zero-migration-risk is the single fastest objection-killer for an upsell motion; and #3 should be reserved specifically for net-new competitive situations, where leading with price first (rather than third) risks the "cheap = lite" perception flagged in the research section's discussion of price-anchored positioning.

Action for Monday morning: the marketing lead should write this three-row table into a one-page positioning brief (not a full messaging guide — scope and time don't allow for one) and distribute it to the 6 AEs/SDRs and 3 CSMs before any email or in-app asset goes out. Expected effort: 2-3 hours to draft, one 30-minute team readout. Expected impact: this is the single cheapest fix available to prevent the sales team from freelancing four different pitches across 340 accounts, which would directly undermine the segment-consistency the tiering work in the prior section is designed to produce.

Messaging by Segment Tier: Adjusting Emphasis, Not Substance

The three-tier segmentation from the Market & Segment Analysis section should drive *emphasis*, not a different core message — the positioning statement stays constant, but which proof point leads changes by tier, since a Tier 1 account with known multi-department complexity has a different trigger than a Tier 2 account that hasn't been confirmed to need multi-step routing yet.

Tier	Lead Message	Call-to-Action Framing	Who Delivers It
Tier 1 (Core ICP, ~120-140 accounts)	Direct: "You have multiple department heads approving spend today with no audit trail — here's how [beta customer type] cut that to under 8 hours."	Specific: "Let's turn this on for your team this month"	AE/CSM in renewal or QBR conversation, personalized with account's own org structure
Tier 2 (Adjacent fit, ~100-120 accounts)	Qualifying: "How many people currently approve expenses across your teams?" before pitching	Discovery-first: "If that's more than one or two people, this might solve a real gap for you"	Email + in-app nudge, then AE follow-up only if the account self-identifies as multi-approver
Tier 3 (Poor fit, ~80-100 accounts)	Passive/informational only	Low-pressure: in-app banner only, no AE time spent	In-app announcement, no sales follow-up

This tiered messaging table should be built alongside — not after — the CRM segmentation work already recommended in the prior section, since the two are the same underlying exercise viewed from different angles (one sorts accounts, the other scripts what to say to each sorted group). Expected effort: half a day once the CRM tiering is done, since the message content already exists in the hierarchy table above and just needs to be mapped to tiers. Expected impact: this is the mechanism that prevents the revenue-shortfall scenario quantified in the prior section's MRR table (the $17K "undisciplined" case) — if AEs use the Tier 1 script on Tier 1 accounts specifically, and don't waste discovery time pitching Tier 3 accounts, the mix should skew toward the $38K "disciplined targeting" scenario rather than the flatter, lower-value outcomes.

Handling the Two Distinct Buyer Personas Within the Message

The intake names two roles in every deal: the **Controller/Head of Finance** (economic buyer, cares about audit trail and control) and the **department heads** who will actually use the approval workflow day-to-day (cares about speed and not being nagged over Slack). Current positioning, as drafted in the intake and reflected in most fintech competitor messaging, speaks almost entirely to the Controller. That's correct for the primary buyer, but it creates a real internal-champion gap: department heads are frequently the ones who complain about the current manual process, and a message that only speaks to finance risks losing an internal advocate who could accelerate the Tier 1 renewal conversations.

Recommendation: create **one additional message variant**, not a whole new campaign, aimed at department heads for use specifically in the in-app announcement (which all users see, not just the finance admin) and in Slack/Teams notification copy itself (since the product's own notifications are a messaging surface the team hasn't fully leveraged):

> "No more chasing approvals over email or Slack threads. Get notified, approve in one click, done."

This is a low-effort addition — it reuses existing product UI copy work already required for the

Slack/Teams notification feature, so it isn't incremental creative production (which is out of scope) but rather functional copy that ships with the feature. Expected impact: modest but real — it widens the internal-champion base within Tier 1 and Tier 2 accounts without requiring any additional marketing bandwidth beyond writing two sentences.

What the Positioning Explicitly Should Not Claim

Given the $60K budget, 2-person team, and the explicit exclusion of creative production from this engagement, discipline about scope of claims matters as much as the claims themselves. Two specific guardrails:

- **Do not position Approvals as a Ramp/Airbase replacement for companies that have outgrown Ledgerly's core product.** The ICP is 100-400 employees with moderate complexity — the moment a net-new prospect describes multi-entity accounting, procurement card issuance, or 500+ employees, the honest answer is that Ramp/Airbase may be the better fit. Oversteering the "we beat Ramp" message into segments where it doesn't hold will produce lost deals and churned trust with AEs who find themselves overselling. This costs nothing to enforce — it's a line in the sales battlecard, not a new asset.

- **Do not lead with the $3/seat/month figure in isolation.** Per-seat pricing without context under-communicates value to a Controller comparing total cost of ownership; it should always appear alongside the $150 minimum and the explicit Ramp/Airbase comparison, never as a standalone number in ads, email subject lines, or landing page headers. This is a one-line copy rule for the 2-person team to apply consistently across every asset they touch in the next 6 months.

Channel plan

Channel Plan

The findings above establish that this launch has two structurally different motions — installed-base expansion and net-new-logo acquisition — and that the client's channels split cleanly along that same line. Rather than propose a generic multi-channel mix, this plan assigns specific channels to specific tiers and motions, sequences them against the 6-month timeline, and attaches an owner, effort estimate, and expected-impact figure to each. Every channel used is one the client already owns; the $60K budget funds enablement and content production layered on top, not new channel acquisition, consistent with the excluded scope (media buying, creative production).

Channel 1: In-App Announcement — Broad Reach, Zero Marginal Cost

What: A persistent in-app banner/modal announcing Approvals, visible to all users across all 340 accounts, live from Day 1 of launch.

Who: Owned by Marketing (whichever of the 2 people handles product-adjacent copy), built with Engineering support (likely under 1 sprint day given it's a banner, not a new feature).

Segment targeting: Per the tiering table in the prior section, this is the *only* channel Tier 3 accounts should receive — no AE/CSM time should be spent on them. For Tier 1/2, the in-app banner serves as a passive reinforcement that runs in parallel with the higher-touch email and sales channels below, not a substitute for them.

Message: Leads with the primary hook (4.5 days -> under 8 hours) for the Controller-facing version; the department-head-facing notification copy variant ("No more chasing approvals over Slack") should ship inside the actual Slack/Teams notification feature itself, not just the banner — this reuses required product copy rather than creating a new asset.

Effort: Low — under 1 week of combined marketing/engineering time. **Expected impact:** Per the funnel model, ~170-200 accounts (50-60% of the base) see this within 30 days; it is a reach mechanism, not a conversion mechanism, so its impact should be measured as "seen," not "converted." Cost: $0 incremental (existing product real estate).

Channel 2: Segmented Email Campaign — Tier-Differentiated, Not Broadcast

What: Three distinct email sequences, one per tier, sent from the customer email list of 340 accounts — not a single broadcast email.

Who: Marketing owns copywriting and send; the CRM segmentation exercise (already recommended, ~3-4 hours) must be complete before this goes out, since the whole point is that Tier 1, 2, and 3 accounts get different content, not the same message at different volumes.

How, specifically:

- **Tier 1 (~120-140 accounts):** A 2-email sequence. Email 1 leads with the beta hero metric and an explicit invitation to a 15-minute walkthrough with their AE/CSM. Email 2 (sent 10 days later, to non-openers/non-responders) reiterates the zero-migration-risk objection-handler. This sequence should trigger a task in the CRM for the account's assigned AE/CSM to follow up within 5 business days of send — email alone should not be relied upon to close Tier 1 accounts; it's a scheduling mechanism for the higher-conversion sales-led motion identified in the funnel model.
- **Tier 2 (~100-120 accounts):** A single email built around the qualifying question ("How many people currently approve expenses across your teams?") with a self-service reply-to-qualify mechanism. Only accounts that reply or click through generate an AE follow-up task — this is the mechanism that protects AE bandwidth from being spent on unconfirmed-fit accounts, directly operationalizing the bandwidth-protection finding above.
- **Tier 3 (~80-100 accounts):** No dedicated email. These accounts receive only the in-app banner (Channel 1). This is a deliberate exclusion, not an oversight — sending email to accounts unlikely to convert costs marketing time to build and slightly degrades average list engagement metrics for no return.

Effort: Medium — estimated 1.5-2 weeks of combined copywriting, CRM segmentation, and email-tool setup, falling within the ~$10K "email/lifecycle campaign build" budget line already allocated. **Expected impact:** Applying the 35-45% open-rate benchmark and directing follow-up specifically at Tier 1/2, this channel should generate the bulk of the "opened" volume in the funnel model (130-150 accounts) and convert a meaningful share of those into scheduled AE/CSM conversations — the actual close happens in Channel 3, below.

Channel 3: Sales/CS-Led Renewal & Upsell Conversations — The Primary Conversion Channel

What: Direct pitch of Approvals by the account's existing AE or CSM, embedded into scheduled renewal calls, QBRs, and check-ins over the full 6-month window — not a one-time campaign.

Who: The 6 AEs/SDRs and 3 CSMs, each covering roughly 38 accounts on average. This channel requires no new hire and no new budget line beyond the sales-enablement spend already allocated (~$20K for battlecards, ROI calculator, case-study support).

How, specifically: This is the highest-leverage and highest-risk channel identified in the findings, so it needs the most operational specificity:

- **Week 1:** Marketing distributes the one-page positioning brief (tiered message table) and a battlecard covering the corrected ~87-90% minimum-commitment price gap, the objection-handling script for "why now, why easy," and the two net-new qualifying questions.
- **Week 2:** Sales leadership adds "Approvals pitch status" as a standing field in the CRM for every account — tracked as Not Started / Pitched / In Evaluation / Closed / Declined — reviewed in the existing weekly pipeline call, adding no new meeting to anyone's calendar.
- **Ongoing, Months 1-6:** AEs/CSMs pitch Tier 1 accounts proactively at the next scheduled touchpoint (renewal, QBR, or check-in) rather than waiting for a dedicated "launch call" — this is the fix for the root-cause risk identified above, where a one-time push in Month 1 would leave later-renewal-cycle accounts untouched.
- **Month 3 checkpoint:** Once 8-10 non-beta customers have converted and generated new before/after data, sales enablement refreshes the battlecard with these fresher case studies (not just the original 8 beta accounts) — this is the mechanism that sequences net-new-logo pitches (Channel 4) behind proven base-expansion proof, per the Bowling Pin framework established in the research section.

Effort: Low incremental effort per rep (embedding into existing calls, not adding new ones) but requires sales leadership discipline to enforce the CRM tracking field — this is a management action item, not a marketing one. **Expected impact:** This channel is the primary driver of the 40-60 account range in the funnel model; it is also the channel most exposed to the risk that a 6-person, non-incentivized sales team quietly deprioritizes a secondary product pitch after Month 1. **Recommendation: add a small SPIF (e.g., $50-100 per converted Approvals account) funded from the $5K contingency line**, since the sales team has no dedicated launch incentive today and this is the cheapest possible lever to keep the pitch active in month 4-6, not just month 1.

Channel 4: Website/SEO — Net-New Inbound, Sequenced Behind Base Proof

What: A dedicated Approvals landing page and 2-3 supporting blog posts targeting search terms around "expense approval workflow software" and "Ramp/Airbase alternative for small business," published against the existing ~4,000 monthly organic visits.

Who: Marketing owns content; ~$15K of the budget covers landing page build and content production (no paid media, consistent with excluded scope).

How, specifically: This channel should **launch in Month 2, not Month 1** — sequenced deliberately behind the first wave of base-expansion conversions, so the landing page and blog content can reference real (if early) customer results rather than beta-only data, per the proof-point sequencing logic established in the findings. The landing page should lead with the corrected minimum-commitment price comparison against Ramp/Airbase (Channel-3's battlecard content, repurposed rather than newly created — an efficiency given the 2-person team constraint), targeting the "overserved by Ramp/Airbase" and "no workflow today" buyer segments identified in the net-new segment table above.

Effort: Medium — landing page plus 2-3 posts is roughly 3-4 weeks of content production, fitting within existing SEO cadence rather than requiring new tooling. **Expected impact:** Given 4,000 monthly organic visits and typical B2B SaaS landing-page conversion rates of 2-4% to a demo/trial request, this channel should generate a modest but real net-new pipeline contribution (roughly 5-10 qualified inbound leads/month once live) — small relative to the base-expansion volume, but directly supports the secondary net-new-logo objective at near-zero incremental cost beyond content production already budgeted.

Channel Sequencing Summary

Month	Primary Channel Activity	Owner	Budget Draw
1	In-app banner live; Tier 1/2 segmented email sent; battlecard + positioning brief distributed to sales	Marketing (both), Sales leadership	~$15K (enablement, email build)
2	AE/CSM renewal-cycle pitching begins in earnest; landing page + first blog post published	Sales/CS, Marketing	~$15K (website/SEO)
3	Case-study refresh from first 8-10 non-beta conversions; battlecard v2 to sales; SPIF activated	Marketing, Sales leadership	~$10K (case studies, contingency draw)
4-5	Continued renewal-cycle pitching (Tier 1 remainder, Tier 2 follow-ups); net-new qualifying questions live in AE discovery calls	Sales/CS	~$10K (remaining enablement/content)
6	Final push ahead of Q1 budget season for undecided Tier 1/2 accounts; net-new pipeline handoff from inbound	Sales, Marketing	~$10K (contingency, final content)

This sequencing is the operational backbone of the plan: it ensures the cheapest, highest-conversion channel (sales/CS into a segmented base) runs continuously rather than as a one-time event, and that the net-new inbound channel is deliberately delayed until it can borrow credibility from early base-expansion wins — directly addressing the root-cause risk that a 2-person marketing team, without this sequencing, would default to running all four channels simultaneously at diluted intensity from Day 1.

Launch timeline

6-Month Launch Timeline

The channel plan above establishes *what* runs where; this section establishes *when*, with explicit weekly milestones, named owners, and go/no-go checkpoints. Given the binding constraint identified throughout this report — a 2-person marketing team with no dedicated launch function, working alongside a 6-person sales/CS team that has other quota responsibilities — the timeline is built around a small number of hard deadlines and a much larger number of "standing motion" activities that run continuously rather than as one-time events. This distinction matters: the single biggest risk to this launch is treating it as a Month 1 event rather than a 6-month operating rhythm, and the timeline below is structured explicitly to prevent that failure mode.

Pre-Launch: Weeks 1-2 (Foundation, No External Communication Yet)

Nothing customer-facing should go out until the segmentation and enablement foundation is in place — the findings above show that unsegmented, undisciplined outreach is the direct path to the $17K MRR shortfall scenario rather than the $38K disciplined-targeting scenario.

Week	Task	Owner	Effort	Deliverable
1	CRM segmentation of 340 accounts into Tier 1/2/3 (employee count, org complexity, current contact role)	Marketing (CRM-proficient of the 2)	3-4 hours	Tiered account list, shared with sales leadership
1	Draft one-page positioning brief (message hierarchy table)	Marketing (both, split drafting/review)	2-3 hours	Positioning brief v1
1	Build battlecard: corrected 87-90% price gap, objection scripts, net-new qualifying questions	Marketing + 1 senior AE for review	4-6 hours	Sales battlecard v1
2	Sales leadership adds "Approvals pitch status" field to CRM (Not Started/ Pitched/In Evaluation/Closed/ Declined)	Sales leadership (Ops support if needed)	1-2 hours (config)	Trackable pipeline field, live before any pitching starts
2	30-minute sales/CS readout: positioning brief + battlecard + tier assignments walked through live	Marketing presents, Sales leadership mandates attendance	30 min meeting + prep	All 6 AEs and 3 CSMs briefed and equipped
2	Draft Tier 1 and Tier 2 email sequences (copy only, not yet sent)	Marketing	1 week (part-time, alongside other work)	Email drafts ready for Week 3 send
2	Slack/Teams notification copy finalized (includes department-head message variant)	Marketing + Engineering (notification feature owner)	2-3 hours copy, existing eng timeline for feature	Notification copy ships with feature, no separate asset needed

Go/no-go checkpoint, end of Week 2: Launch does not proceed to Week 3 external communication unless (a) the tiered list exists and is loaded into CRM, (b) the battlecard has been reviewed by at least one AE for realism, and (c) the pitch-status field is live and trackable. This is a cheap, disciplined gate — missing it risks repeating the untargeted-broadcast failure mode the findings explicitly warn against.

Launch Week: Week 3 (First External Signal, Ahead of Q1 Budget Season)

This is the only week with a genuine "launch moment" — everything else in the 6 months is sustained execution, not a single event.

Day	Action	Owner	Notes
Mon	In-app banner goes live for all 340 accounts	Marketing + Engineering	Reaches Tier 3 as their only touchpoint; reaches Tier 1/2 as passive reinforcement
Mon	Tier 1 email sequence (Email 1) sent to ~120-140 accounts	Marketing	Leads with 4.5 days -> under 8 hours hero metric; invites 15-minute AE/CSM walkthrough
Tue	Tier 2 email sent to ~100-120 accounts, qualifying-question framing	Marketing	Only self-identified multi-approver replies generate an AE follow-up task
Wed-Fri	AEs/CSMs begin folding Approvals into any renewal, QBR, or check-in call already on the calendar that week	Sales/CS	No new calls scheduled specifically for this — embedded into existing cadence, per the low-incremental-effort design in the channel plan
Fri	Marketing checks open/click rates on both email sequences; flags any Tier 1 non-openers for the Email 2 follow-up (10 days out)	Marketing	First read on awareness-stage funnel performance vs. the 35-45% open-rate benchmark

Expected impact this week: per the funnel model, roughly 170-200 accounts see the in-app banner within 30 days of this week, and 130-150 accounts open at least one email. Conversion has not meaningfully started yet — this week is entirely an awareness/reach week, and should be measured as such, not judged against the 51-account target prematurely.

Weeks 4-8 (Month 2): Sustained Sales Motion Begins, Landing Page Launches

Week	Task	Owner	Effort	Expected Output
4	Tier 1 Email 2 (objection-handler, zero-migration-risk) to non-responders	Marketing	Low (reuses drafted copy)	Incremental opens/replies from the ~55-65% who didn't engage with Email 1
4-6	AEs/CSMs work through Tier 1 accounts systematically at each account's next natural touchpoint — tracked weekly in the existing pipeline review, no new meeting added	Sales/CS, tracked by Sales leadership	Ongoing, embedded	Pitch-status field shows "Pitched" status climbing across Tier 1

5-7	Approvals landing page and first blog post built and published, leading with the corrected minimum-commitment price comparison	Marketing	3-4 weeks total content production	Live net-new inbound surface by end of Month 2
6	Net-new qualifying questions ("How do you route approvals today?" / "What are you paying for that today?") added to all AE discovery-call scripts, not just Approvals-specific calls	Sales leadership	Script update, no new training needed	Net-new prospects sorted into the three buyer-situation buckets from the segment analysis
8	First checkpoint: pull pitch-status field data — how many Tier 1 accounts show "Pitched" vs. "Not Started"?	Marketing + Sales leadership jointly	1-hour review meeting	Early warning if sales bandwidth is lagging the plan

Go/no-go checkpoint, end of Week 8: If fewer than 50% of Tier 1 accounts show a "Pitched" status by this point, sales leadership should treat this as the trigger to escalate — either through the SPIF incentive (see below) or a direct manager-level push — rather than waiting until Month 4-5 to notice the shortfall. This checkpoint operationalizes the report's core risk finding: that sales bandwidth, not marketing reach, is the binding constraint on the 15% target.

Weeks 9-13 (Month 3): Proof Refresh and Incentive Activation

By this point, the plan assumes the first wave of non-beta conversions has started to close, which is what unlocks the sequencing logic from the channel plan (net-new inbound and net-new sales conversations should lean on fresh proof, not beta-only data).

Week	Task	Owner	Effort	Expected Output
9	Pull first 8-10 non-beta customer results (turnaround time, seat count, feedback quote)	Marketing + CS (data pull from converted accounts)	1-2 days	Refreshed proof points beyond the original 8 beta customers
10	Battlecard v2 distributed to sales, incorporating fresh case studies	Marketing	Half-day update to existing document	Sales pitching with current, not stale, proof

10	$50-100 per-account SPIF activated for AEs/CSMs, funded from the $5K contingency line	Sales leadership + Ops (payroll/comp mechanics)	Low — policy decision plus payroll setup	Direct incentive to keep pitching active into Months 4-6, addressing the "quietly deprioritized after Month 1" risk named in the channel plan
11-13	Continued Tier 1/Tier 2 pitching; net-new AEs begin using refreshed case studies in competitive deals against Ramp/Airbase	Sales/CS	Ongoing, embedded	Net-new pipeline entries tagged by buyer-situation bucket (no workflow / overserved / underserved)
13	Mid-launch review: compare actual attach count and MRR mix against the three-scenario table (disciplined $38K / even $26K / undisciplined $17K)	Marketing + Sales leadership + Finance/Ops	Half-day review	Determines whether remaining $5K contingency gets spent on additional enablement or additional incentive

Weeks 14-22 (Months 4-5): Sustained Push, Tier 2 Conversion Focus

Week	Task	Owner	Effort	Expected Output
14-18	Continued renewal-cycle pitching, now shifting emphasis toward Tier 2 accounts that self-qualified via email reply but haven't yet converted	Sales/CS	Ongoing, embedded	Tier 2 "In Evaluation" accounts move to "Closed"
16	Second blog post published (targeting "Ramp alternative for small business" search intent)	Marketing	1-2 weeks	Continued modest net-new inbound (5-10 qualified leads/month range established in channel plan)
18	90-day-style budget checkpoint: review remaining spend against the 5-line allocation table; reallocate any underspent line (e.g., if enablement came in under $20K) toward additional case-study production or a second SPIF cycle	Marketing + Sales leadership	Half-day	Budget stays adaptive rather than locked to original allocation

20-22	Third blog post if inbound metrics justify it; otherwise redirect remaining content budget to sales enablement, given the report's finding that sales/CS is the higher-leverage channel	Marketing	Conditional	Budget flexed toward proven channel, not spent on autopilot

Weeks 23-26 (Month 6): Final Push Ahead of Q1 Budget Season, Close-Out

This final stretch is deliberately timed against the buyer-seasonality finding from the research section: Controllers and Heads of Finance disproportionately evaluate new tooling in Q4 to secure Q1 budget lines, which is the entire rationale for launching this quarter in the first place.

Week	Task	Owner	Effort	Expected Output
23	Final Tier 1/Tier 2 outreach wave — AEs contact any account still showing "Not Started" or "In Evaluation" with explicit Q1-budget-timing framing ("lock this in before your budget resets")	Sales/CS	Embedded in existing renewal cadence	Last-chance conversions ahead of budget-cycle deadline
24	Net-new pipeline handoff review: which inbound/outbound net-new deals used the Approvals wedge, and at what win rate versus core-product-only deals	Sales leadership + Marketing	Half-day analysis	Direct read on the secondary objective — does Approvals lift core-product win rate
25	Final account-count and MRR tally against the 51-account / $35K MRR target, segmented by tier (per the revenue-math table in the prior section)	Marketing + Finance/Ops	1 day	Final reporting artifact for leadership

26	Retrospective: which channel (email, in-app, sales-led, inbound) drove the highest-quality (Tier 1, highest-$/account) conversions — informs whether the next add-on launch gets a different budget split	Marketing + Sales leadership	Half-day	Institutional learning for future add-on launches, given Ledgerly's stated intent to use this as a repeatable wedge motion

Why This Sequencing, Specifically

Two design choices in this timeline are deliberate and worth naming explicitly, since they diverge from what a less resource-constrained team might do:

No simultaneous four-channel launch in Week 1. The in-app banner and Tier 1/2 emails go out together (cheap, mostly automated), but the landing page and SEO content are deliberately held until Month 2, and the SPIF isn't activated until Month 3. A 2-person marketing team attempting to build a land...page, run a full email sequence, refresh a battlecard, and coordinate sales enablement all in the same week would almost certainly under-execute on all four rather than fully executing on two — a direct consequence of the capacity constraint named explicitly in the intake ("marketing has 2 people, not a dedicated team"). Staggering the channels is not a compromise on ambition; it's the mechanism that makes the plan actually executable by the team that exists, not the team the goal implicitly assumes.

No new hire, no dedicated launch team, no agency. Given the $60K cap and explicit exclusion of creative production and media buying from this engagement, every recommendation above has been built to run through existing headcount — 2 marketers, 6 AEs/SDRs, 3 CSMs — with the budget funding enablement materials and content rather than additional people or paid reach. Where extra effort is needed (case study production, battlecard refreshes, landing page copy), the plan reuses material that already exists in another form (beta data, product UI copy, Channel 3's battlecard content repurposed for Channel 4) rather than commissioning net-new assets, which is the only way this scope is achievable without breaching the budget or capacity constraints.

6-Month Launch Timeline: Week-by-Week for the First 30 Days, Then Monthly

The channel sequencing summary above lays out the monthly cadence; the table below breaks the critical first 30 days into weekly actions, since this is the period where sequencing errors (launching everything at once, or launching without the segmentation groundwork) are most likely to occur and hardest to walk back once emails are sent and sales conversations have started.

Week	Action	Owner	Effort	Dependency
Week 1	CRM segmentation into Tier 1/2/3 (employee count, org structure, current contact role)	Marketing (CRM-lead person)	3-4 hours	None — this is the first action, everything else depends on it

Week 1	Draft and distribute one-page positioning brief + message hierarchy table to sales/CS	Marketing (second person)	2-3 hours drafting, 30-min team readout	Positioning statement finalized (already drafted in this report)
Week 1	Battlecard v1 built (corrected 87-90% price gap, objection scripts, net-new qualifying questions)	Marketing, with AE input	1 week (part of $20K enablement budget)	Positioning brief
Week 1	In-app banner spec finalized with Engineering	Marketing + Eng	Under 1 sprint day	None
Week 2	In-app banner goes live (all 340 accounts)	Marketing + Eng	—	Banner build complete
Week 2	Tier 1 email sequence (Email 1) sent	Marketing	Half day to execute (copy already drafted)	Tier segmentation, positioning brief
Week 2	Tier 2 qualifying email sent	Marketing	Half day	Tier segmentation
Week 2	"Approvals pitch status" field added to CRM; sales leadership announces tracking cadence in weekly pipeline call	Sales leadership	Under 1 hour setup	None
Week 3	AE/CSM outreach begins for Tier 1 accounts with upcoming renewal/ QBR dates in the next 60 days	AEs/CSMs	Embedded in existing calls — no new meetings	Battlecard, CRM tracking field
Week 3	Tier 1 Email 2 (non-openers/non-responders) sent	Marketing	Half day	Email 1 open/click data
Week 4	First pipeline review of "Approvals pitch status" field in weekly sales call — early read on Tier 1 pitch coverage	Sales leadership	Existing meeting, no added time	Week 2-3 outreach data
Week 4	Net-new discovery-call qualifying questions rolled into AE scripts for all new prospect conversations	AEs/SDRs	Script update only	Battlecard

By the end of Week 4, the client should have: full base segmented into three tiers, positioning live everywhere it needs to be, sales actively pitching Tier 1 accounts with renewal dates in the near term, and a tracking mechanism in place to see, for the first time, real pitch-to-close data rather than the funnel model's benchmark-based estimates. This is the point at which the plan shifts from "execute the launch" to "manage the pipeline," which is the appropriate posture for a 6-person sales team working a 6-month cross-sell motion.

Month-by-Month Milestones and Checkpoints, Months 2-6

Month	Milestone	Success Signal	Checkpoint Action If Off-Track
Month 2	Landing page + first blog post live; Tier 1 pitching in full swing; Tier 2 follow-ups from qualified replies underway	At least 30-40% of Tier 1 accounts (36-56 of ~120-140) show "Pitched" status in CRM	If under 25% pitched, sales leadership escalates in weekly pipeline call — likely a bandwidth or prioritization issue, not a messaging issue, per the root-cause finding above
Month 3	First 8-10 non-beta customer conversions generate refreshed case-study data; battlecard v2 distributed; SPIF activated	At least 15-20 total conversions (beta + early adopters) tracked; blended MRR trending toward the "disciplined targeting" scenario (~$800/Tier 1 account average)	If MRR-per-account is trending toward the $300-400 "undisciplined" range, re-audit which tier is actually converting and redirect AE effort back to Tier 1
Month 4	Net-new inbound from landing page begins generating qualified leads; AEs using tiered qualifying questions on all new deals	5-10 qualified net-new leads/month attributable to Approvals-related content	If inbound volume is materially below this, reallocate remaining content budget toward a second blog post rather than a new channel
Month 5	Continued Tier 1/2 renewal-cycle pitching; contingency budget deployed based on Month 3 checkpoint findings	Cumulative conversions tracking toward 40+ accounts (low end of funnel model range)	If cumulative conversions are under 30 by Month 5, this is the trigger to reconsider the 6-month deadline itself with the client, rather than discovering the shortfall only at the Month 6 readout
Month 6	Final push ahead of Q1 budget planning season; net-new pipeline handed off to standard sales motion; full readout against both account-count (15%) and MRR ($35K) targets	51+ accounts converted, ideally skewed toward the "disciplined targeting" MRR scenario (~$38K) rather than the flatter blended scenario	Full retrospective against the tiered segmentation, message hierarchy, and channel sequencing — feeds directly into whether Approvals justifies a dedicated, larger-scope launch motion for future add-on products

Two Structural Risks to Flag for Monday Morning

Risk 1: Sales bandwidth quietly lapses after Month 1. This has been named as the single highest-leverage and highest-risk factor throughout this report, and the timeline above is built specifically to counter it — by embedding the CRM tracking field into an *existing* weekly meeting rather than creating a new one, and by activating the SPIF at Month 3 rather than Month 1, so the incentive is live precisely when initial launch enthusiasm typically fades. Sales leadership should treat the Month 2 and Month 4 checkpoints in the table above as non-negotiable agenda items, not optional check-ins.

Risk 2: The team declares success on the 15% headline while missing the $35K MRR target. As quantified in the segment-level revenue math above, these are not the same outcome, and the timeline's Month 3 checkpoint is deliberately placed to catch a tier-mix problem while there are still three months left to correct it — reallocating AE effort back toward Tier 1 accounts is a viable mid-course correction in Month 3-4, but not in Month 6. The marketing lead should own bringing both numbers (account count and blended MRR-per-account) to every monthly checkpoint, not just the headline count, since the CRM tracking field alone won't surface the mix problem without someone explicitly calculating it.

Risks & Mitigations

Every plan in this report assumes disciplined execution against a tight budget and a team that has other jobs. The risks below are the specific ways that discipline is likely to break down at Ledgerly, not generic launch risks — each is tied to a constraint, number, or behavior pattern already surfaced in the analysis above.

Sales bandwidth lapses after Month 1 (highest-priority risk). This has been flagged repeatedly because it is the single most probable failure mode: a 6-person sales team with existing quota responsibilities treats Approvals as a one-time pitch rather than a standing part of every renewal conversation for six months. Nothing in Ledgerly's current comp plan rewards AEs for closing a $150-$1,200/month add-on over a core-product deal, and there's no dedicated launch owner forcing accountability once the initial Week 2 kickoff enthusiasm fades. This is likely (it's the default behavior of any secondary-priority motion layered onto a quota-carrying team) and high-impact, since the funnel model shows sales-led conversion is the dominant channel — if it lapses, the 40-60 account range collapses toward its low end or below. Mitigation: the CRM pitch-status field and Month 2/Month 4 checkpoints already recommended are necessary but not sufficient; the $50-100/account SPIF should be treated as mandatory, not optional, and sales leadership should own a monthly (not just Month 3 and Month 8) review of pitch-status coverage, reported as a standing line item in the weekly pipeline call through Month 6.

Marketing's 2-person team gets pulled onto other priorities. With no dedicated launch team, the two marketers running this plan almost certainly have other standing responsibilities (core-product campaigns, SEO maintenance, event support). A competing fire drill — a core-product launch, a board deck request, a customer escalation — is a realistic and common way that Week 3-8 deliverables (email sequences, landing page, battlecard v2) slip by two to four weeks. This is moderately likely and moderately-to-severely impactful, since the whole sequencing logic (delay the landing page to Month 2, refresh the battlecard at Month 3) depends on things happening roughly on schedule. Mitigation: leadership should explicitly protect a fixed weekly time block (e.g., one full day/week each) for both marketers against this launch for the first 90 days, and the $5K contingency line should be pre-authorized for freelance copywriting support if a deliverable is at risk of slipping past a hard checkpoint (e.g., the Week 3 launch-week emails).

The team optimizes for the 15% headline and misses the $35K MRR target. The revenue-math scenario table shows a realistic path to hitting 51+ accounts while landing 30-40% below the MRR goal, simply by over-indexing on easier-to-close Tier 2/3 accounts. This is moderately likely, because "close whatever will close" is the natural instinct under a headcount goal, especially if a mid-launch checkpoint shows the team behind on account count and reaching for easy wins to compensate. Impact is moderate — the company still gets a real proof point and case studies, just a financially weaker one, and a false sense of success at the board level. Mitigation: report both numbers (account count and blended MRR/account) at every monthly checkpoint, not just at Month 6, and give sales leadership explicit permission to under-hit the 51-account number if it means holding out for Tier 1 fit — this should be stated to the sales team directly, since "any 51 accounts" and "51 good-fit accounts" produce very different outcomes and the team needs to know which one leadership actually wants.

Budget overrun on enablement/content, or underspend that leaves money unused. At $60K over six months with two non-specialist marketers building a landing page, case studies, and a battlecard, there's real risk of scope creep on any one line item (e.g., the case-study production taking longer or costing more than the $10K allotted if it requires customer interviews, design, and multiple review cycles) as well as the opposite risk — underspending because the team is too capacity-constrained to execute all five budget lines and money sits unused while the SPIF or contingency goes untapped. Likelihood is moderate; impact is low-to-moderate, since the total exposure is capped at $60K and no paid acquisition is at stake. Mitigation: the Week 18 and Month 3 budget checkpoints already built into the timeline should include an explicit reallocation decision, not just a status check — if enablement is running over, pull from the SEO/content line (since Channel 4 is the lowest-leverage channel per the funnel analysis) rather than letting overspend silently erode the contingency reserve needed for the SPIF.

Competitive response from Ramp/Airbase, or from smaller point solutions adding basic approval routing. Ramp and Airbase are unlikely to react meaningfully to a $150/month add-on launch from a company at Ledgerly's scale — this risk is low-likelihood in the 6-month window given their pricing structure and target market are structurally different. The more realistic competitive risk is horizontal: Expensify, Zoho Expense, or QuickBooks could add multi-step approval logic to their existing platforms, eroding the "underserved by point solution" segment identified in the net-new buyer table. This is low-to-moderate likelihood within six months (feature parity moves slower than pricing moves) and low-to-moderate impact, since that segment was already flagged as the weakest-urgency of the three net-new buyer buckets. Mitigation: no immediate action needed beyond monitoring: the Month 6 retrospective should include a scan of competitor changelogs, and if a point-solution competitor does ship comparable functionality, Ledgerly's differentiation should shift further toward the audit-trail/multi-department complexity story (its core structural advantage) rather than compete on notification/approval basics alone.

CRM/data limitations undermine the tiering strategy. The entire segmentation and messaging plan depends on the assumption that employee count, org structure, and contact role are reliably captured in Ledgerly's CRM. If that data is incomplete or stale for a meaningful share of the 340 accounts — a real possibility, since CRM hygiene is a common weak point at 26-person companies without dedicated RevOps — the Tier 1/2/3 split could misclassify accounts, sending the wrong message to the wrong buyer or wasting AE time on a "Tier 1" account that's actually a poor fit. This is moderately likely and moderately impactful, since it would quietly degrade close rates without an obvious root cause. Mitigation: before the Week 1 segmentation is finalized, the marketing lead doing the CRM pull should sample 15-20 accounts manually against known account history (renewal notes, support tickets, CS relationship knowledge) to sanity-check data completeness before trusting it to drive the full 340-account split; where data is missing, default those accounts to Tier 2 (qualify-first) rather than Tier 1 (direct pitch), since a false-negative (under-pitching a good account) is cheaper than a false-positive (wasting AE time on a poor fit).

Adoption friction inside converted accounts (churn-back risk). Even...even accounts that formally convert to Approvals could see low actual usage if department heads don't adopt the new workflow, which would show up months later as add-on churn or negative word-of-mouth undermining the case-study pipeline the plan depends on for Month 3 onward. This is moderately likely — the 8 beta customers were self-selected early adopters, and the report has already flagged that the department-head persona is under-messaged relative to the Controller. If CS doesn't actively monitor usage (not just billing status) post-conversion, a logo could be counted as a "win" in the CRM while quietly failing to deliver the audit-trail/speed value that makes it a credible reference later.

Cost overrun against the $60K cap. The budget table allocates funds tightly across five lines with only a $5K contingency. Given that two of those lines (case-study production, sales enablement) depend on real customer conversions happening on schedule, any delay in the Month 3 conversion wave could compress the timeline for producing fresh proof points, tempting the team to spend contingency early rather than holding it for the Month 3 checkpoint as designed. This is a moderate-likelihood, low-severity risk in dollar terms (the total budget is small enough that overruns are capped in absolute size) but a higher-severity risk in sequencing terms if it means skipping the case-study refresh that Month 4-6 net-new pitching depends on.

Competitive response from Ramp or Airbase. Given Ledgerly's plan to explicitly name these vendors in battlecards and landing-page copy, there's a real, if lower-probability, risk that either company notices the competitive framing (especially if it starts appearing in analyst conversations, review sites, or losses reported by their own sales teams) and responds with a lightweight "SMB tier" or bundled discount aimed at undercutting exactly this positioning. This is a longer-horizon risk (unlikely to materialize within the 6-month window given enterprise sales-motion inertia at larger competitors) but would meaningfully blunt the "87-90% cheaper" argument that Channel 3 and Channel 4 both depend on if it happened sooner than expected.

Internal champion or CS bandwidth conflict. The 3-person CS team is asked to both monitor post-conversion usage (to prevent the adoption-friction risk above) and actively pitch Approvals in QBRs — on top of their existing core-product retention responsibilities. With only 3 CS people covering a meaningful share of 340 accounts, this is a real capacity conflict that could quietly deprioritize usage monitoring in favor of pitching (or vice versa), especially since pitching has a visible CRM-tracked metric (pitch status) while usage monitoring does not, absent explicit instruction to log it.

The table below consolidates these risks with likelihood, impact, and mitigation, ordered roughly by expected severity (likelihood × impact) rather than by order of appearance above.

Risk	Likelihood	Impact	Mitigation
Sales/CS bandwidth quietly lapses after Month 1, leaving Tier 1 accounts un-pitched through Months 4-6	High	High	Embed "Approvals pitch status" as a standing field in the existing weekly pipeline review (not a new meeting); activate the $50-100 SPIF at Month 3 specifically to counter post-launch fatigue; treat the Month 2 and Month 4 checkpoints as non-negotiable escalation triggers

Team hits the 15%-of-340 headline but misses the $35K MRR target due to Tier 2/3-heavy conversion mix	Medium-High	High	Track blended MRR-per-account alongside raw count at every monthly checkpoint; redirect AE effort back to Tier 1 accounts specifically if Month 3 data shows a skew toward low-seat-count accounts, while there's still runway to correct it
CRM data on employee count/org structure is incomplete or stale, causing mis-tiering of accounts	Medium	Medium	Spot-check the Week 1 segmentation against AE/CSM account knowledge before finalizing tiers; default ambiguous accounts to Tier 2 (qualify-first) rather than Tier 1, so the cost of a data gap is a missed opportunity, not wasted AE time
Converted accounts show low actual usage post-signup, undermining future case-study credibility and creating silent churn risk	Medium	Medium-High	CS explicitly tracks in-product usage (not just billing status) for the first 90 days post-conversion as a defined responsibility, separate from pitching; flag low-usage accounts for a lightweight onboarding nudge rather than counting them as "closed" successes prematurely
CS capacity conflict between pitching Approvals and monitoring post-conversion adoption, given only 3 CS people	Medium	Medium	Sales leadership explicitly divides responsibility — CSMs are not expected to do both in parallel for every account; usage monitoring is prioritized for the first 8-10 non-beta conversions specifically, since those are the accounts feeding the Month 3 case-study refresh
Budget contingency ($5K) gets spent early to cover Month 1-2 gaps, leaving nothing for the Month 3 case-study refresh or a second SPIF cycle	Medium	Low-Medium	Explicitly reserve the $5K contingency for the Month 3 checkpoint in the budget plan (already stated); require sign-off from both marketing and sales leadership before drawing on it early, rather than letting either person unilaterally reallocate

Positioning dilutes across four selling points (speed, audit trail, migration risk, price) because a 2-person team defaults to including everything in every asset	Medium	Medium	Enforce the three-tier message hierarchy (hook / objection-handler / competitive closer) as a written one-page brief distributed before any asset ships; spot-check the first email and landing page draft against the hierarchy before publishing
Ramp or Airbase responds with a lightweight SMB-tier offer or discount that undercuts the ~87-90% price-gap argument	Low (within 6-month window)	Medium-High if it occurs	Lead net-new competitive messaging with the speed/audit-trail hero metric rather than price alone, so the core value proposition doesn't collapse if the price gap narrows; monitor competitor pricing pages quarterly as a standing task for one of the two marketers
Net-new-logo motion is pitched too early (before base-expansion case studies exist), weakening credibility of the wedge story	Medium	Medium	Explicitly sequence net-new competitive pitching behind the Month 3 case-study refresh, as built into the channel and timeline plans; AEs use beta-only data for early net-new conversations but are instructed not to over-index on it as the sole proof point

Two of these risks deserve a final word because they are more structural than tactical. The sales-bandwidth risk is the one this entire report has returned to repeatedly, and that repetition is intentional: every mitigation elsewhere in the plan — tiering, message hierarchy, SPIF timing, checkpoint cadence — ultimately exists to protect a scarce, uninc...entivized resource (6 AEs/CSMs already carrying core-product quota) from being asked to do a second job without a second set of hands. If that protection fails, no amount of email or landing-page optimization elsewhere in the plan can compensate, because the funnel model shows conversion is bottlenecked at the sales-led stage, not the awareness stage. The MRR-mix risk is the second structural one, because it's the failure mode most likely to go undetected until it's too late to fix: a team can hit 51 accounts, celebrate the headline number, and only discover in month 6 reporting that the mix skewed toward low-seat-count Tier 2/3 accounts and landed at $17-20K MRR instead of $35K. Both risks share the same root cause — a metric that looks healthy at the surface level (accounts pitched, accounts closed) can mask an execution problem underneath (which accounts, at what value) — which is why the checkpoints built into the timeline are structured to surface tier-mix and MRR-per-account data monthly, not just at the final readout.

Risk	Likelihood	Impact	Mitigation

Sales/CS bandwidth lapses after Month 1 momentum fades (6-person team has core-product quota obligations that will always take precedence over a secondary add-on pitch)	High — this is the most probable failure mode given no dedicated launch team and no built-in incentive at launch	High — funnel model shows sales-led conversion is the primary driver of the 40-60 account range; if this channel stalls, no other channel can compensate	Embed "Approvals pitch status" as a standing CRM field reviewed in the *existing* weekly pipeline call (no new meeting); activate a $50-100/ account SPIF at Month 3, timed specifically to counter the point where initial enthusiasm typically fades; treat Month 2 and Month 4 checkpoints as non-negotiable escalation triggers
Team hits the 15% account-count target but misses the $35K MRR target because conversions skew toward lower-value Tier 2/3 accounts (the "undisciplined targeting" scenario quantified at ~$17K)	Medium — plausible if reps chase the easiest yeses rather than working Tier 1 systematically, especially under time pressure late in the window	Medium-High — a 30-40% revenue shortfall despite an apparently successful headline number would misrepresent the launch's real commercial impact to leadership	Track blended MRR-per-account alongside raw account count at every monthly checkpoint, not just at the final readout; explicitly instruct AEs to prioritize Tier 1 accounts first via the tiered messaging table; use the Month 3 case-study refresh to re-energize Tier 1-specific pitching before the window narrows
2-person marketing team over-commits by running all four channels (in-app, email, sales enablement, SEO/landing page) at full intensity simultaneously in Week 1, producing diluted execution on all four rather than strong execution on two	Medium — a natural temptation given launch excitement and a hard quarter-end deadline, especially without a dedicated launch function to enforce sequencing discipline	Medium — would likely delay the landing page, weaken battlecard quality, or push segmentation work aside in favor of visible "launch" activities, undermining the targeting-dependent revenue math above	Follow the staggered channel timeline explicitly: in-app + email in Week 1-3, landing page/SEO deliberately held to Month 2, SPIF deliberately held to Month 3; treat the Week 2 go/no-go checkpoint (segmentation + battlecard + tracking field complete) as a hard gate before any customer-facing send
Competitive response from Ramp/Airbase or a point-solution vendor (e.g., a pricing move, a "lite" tier announcement, or aggressive retention outreach to Ledgerly's net-new targets) once Approvals gains visibility in the market	Low-Medium — Ledgerly is a small, add-on-scale competitive signal unlikely to trigger an immediate strategic response from much larger platforms in a 6-month window, though a point-solution vendor (Expensify, Zoho) reacting to defend share is plausible	Medium — would primarily affect the net-new-logo motion (Channel 4, Months 2+), not the base-expansion motion, since existing Ledgerly customers are insulated by the zero-migration-risk positioning	Lean into the zero-migration-risk and minimum-commitment price-gap messaging specifically in competitive net-new deals; equip AEs via the battlecard to reframe any competitor pricing move as validation of the underserved-middle thesis rather than a threat; monitor competitor pricing pages quarterly as a low-cost tracking task for one of the two marketers

Beta customer proof points (4.5 days -> 8 hours) don't hold up at scale once dozens of new, more varied accounts adopt the feature, weakening the hero metric mid-launch	Low-Medium — beta cohorts are often more engaged/motivated than average customers, so some regression toward the mean is plausible once Tier 2 accounts with less structured approval needs adopt the feature	Medium — the hero metric is the single most-repeated asset across every channel in this plan; if it softens materially, messaging built entirely around one number becomes vulnerable	Refresh the metric using the first 8-10 non-beta conversions at the Month 3 checkpoint (already planned) specifically to test whether the number holds; if it softens, shift language from a single point-estimate to a validated range ("4.5 days to under 24 hours" or similar) rather than defending an increasingly unrepresentative number, and avoid a single case study to over-index on it as the sole proof point
CRM segmentation (Tier 1/2/3) is built on incomplete or stale account data (e.g., employee counts not updated since original sale, org-structure assumptions inferred rather than confirmed), producing miscategorized accounts and wasted AE effort	Medium — the tiering exercise explicitly relies on data "likely already in the CRM," which is an assumption, not a confirmed fact, and CRM hygiene is a common gap in fast-growing SaaS companies at this stage	Medium — miscategorization would undermine the entire targeting logic the revenue-math scenarios depend on, though the cost of the error is self-correcting once AEs have live conversations and can recategorize accounts in real time	Treat Week 1 segmentation as a first-pass draft, not a final list; instruct AEs to confirm or correct tier assignment during the first actual conversation with each account and update the CRM field accordingly, turning sales conversations into a live data-quality feedback loop rather than assuming the initial tiering is fully accurate
Budget line items run over in one category (e.g., case-study production takes longer than the half-day estimate, or landing page build exceeds $15K) with no slack given the hard $60K cap and no media-buying flexibility to trim	Low-Medium — the $5K contingency line is thin relative to six months of execution, and creative/content work has a well-documented tendency to run over initial estimates even at small scale	Low-Medium — a modest overrun in one line item is unlikely to derail the whole launch given	

90-Day Action Plan

Below is the 90-Day Action Plan. It covers only the first half of the 6-month launch window already established in this report (Weeks 1–13), sequenced into 0–30 / 31–60 / 61–90 day horizons, with every action traceable to a recommendation made earlier.

The guiding principle carried into this plan is the one established in the Findings and Channel Plan sections: sales/CS bandwidth, not marketing reach, is the binding constraint, so the first 90 days are deliberately structured to (1) build the segmentation and enablement foundation before any customer-facing send, (2) get the sales-led motion running as a *standing* habit rather than a one-time push, and (3) hold net-new and top-of-funnel activity back until base-expansion proof exists to support it.

Horizon	Action	Owner	Effort	Milestone / Done-Criteria
0–30 Days	Segment all 340 accounts into Tier 1/2/3 by employee count, org complexity, and contact role in CRM	Marketing (CRM-proficient lead)	3–4 hours	Tiered list loaded into CRM and shared with sales leadership
0–30 Days	Spot-check 15–20 segmented accounts against AE/CSM account knowledge to validate CRM data quality before trusting the full split	Marketing + spot-check AEs	2–3 hours	Data-quality gaps identified; ambiguous accounts defaulted to Tier 2
0–30 Days	Draft one-page positioning brief with the corrected message hierarchy (hero metric -> objection-handler -> competitive closer)	Marketing (both)	2–3 hours	Brief finalized and distributed to all 6 AEs/SDRs and 3 CSMs
0–30 Days	Build battlecard v1: corrected ~87–90% minimum-commitment price gap, objection scripts, two net-new qualifying questions	Marketing, with 1 AE reviewing for realism	~1 week (part of $20K enablement budget)	Battlecard reviewed and approved by at least one AE
0–30 Days	Add "Approvals pitch status" field (Not Started/ Pitched/In Evaluation/Closed/ Declined) to CRM, tied to the existing weekly pipeline call	Sales leadership	1–2 hours setup	Field live and reviewed in Week 2's pipeline meeting
0–30 Days	30-minute sales/CS readout of positioning brief, battlecard, and tier assignments	Marketing presents; sales leadership mandates attendance	30 min + prep	All 6 AEs and 3 CSMs briefed and equipped before any send
0–30 Days	Finalize Slack/ Teams notification copy, including department-head message variant ("no more chasing approvals over Slack")	Marketing + Engineering	2–3 hours copy	Copy ships bundled with existing notification feature
0–30 Days	Build and launch in-app banner for all 340 accounts	Marketing + Engineering	<1 sprint day	Banner live, visible to all logged-in users

0–30 Days	Draft and send Tier 1 email sequence (Email 1: hero metric + AE/CSM walkthrough invite)	Marketing	Half day to execute	Sent to ~120–140 Tier 1 accounts
0–30 Days	Draft and send Tier 2 qualifying email (self-service reply-to-qualify)	Marketing	Half day	Sent to ~100–120 Tier 2 accounts; only replies generate AE follow-up
0–30 Days	Send Tier 1 Email 2 (objection-handler) to non-openers/non-responders	Marketing	Half day (reuses drafted copy)	Sent 10 days after Email 1
0–30 Days	Begin folding Approvals into any renewal/QBR/check-in call already on the calendar	AEs/CSMs	Embedded in existing calls — no new meetings added	First Tier 1 pitches logged in CRM pitch-status field
0–30 Days	Roll net-new qualifying questions ("How do you route approvals today?" / "What are you paying for that today?") into all AE discovery scripts	Sales leadership	Script update only	Questions live in all new prospect conversations
31–60 Days	Systematically work through Tier 1 accounts at each account's next natural touchpoint	Sales/CS, tracked by sales leadership	Ongoing, embedded	Pitch-status field shows rising "Pitched" count across Tier 1
31–60 Days	Build and publish Approvals landing page, leading with the corrected minimum-commitment price comparison	Marketing	3–4 weeks content production (~$15K budget line)	Landing page live, targeting "Ramp/ Airbase alternative" and "expense approval workflow" search intent
31–60 Days	Publish first supporting blog post	Marketing	Part of same content sprint	Post live and indexed
31–60 Days	First pipeline checkpoint: pull pitch-status data — % of Tier 1 accounts "Pitched" vs. "Not Started"	Marketing + sales leadership jointly	1-hour review	If <50% Tier 1 pitched by Week 8, escalate via manager-level push, not wait until Month 4
31–60 Days	Confirm/correct tier assignments in CRM as AEs complete first live conversations with each account	AEs	Ongoing, part of normal call flow	Tier field becomes a live, self-correcting data source rather than a static Week 1 guess

31–60 Days	Reinforce message-hierarchy discipline: spot-check first email and landing-page drafts against the 3-tier hook/objection/closer structure before publishing	Marketing	<1 hour review	No asset ships without hierarchy check
61–90 Days	Pull first 8–10 non-beta customer results (turnaround time, seat count, quote)	Marketing + CS	1–2 days	Refreshed, non-beta proof points captured
61–90 Days	Distribute battlecard v2 incorporating fresh case studies	Marketing	Half-day update	Sales team using current, not stale, proof by Week 10
61–90 Days	Activate $50–100/account SPIF for AEs/CSMs, funded from the $5K contingency line	Sales leadership + Ops	Low — policy + payroll setup	SPIF live and communicated to sales team by Week 10
61–90 Days	CS begins explicit post-conversion usage monitoring (not just billing status) for the first 8–10 converted accounts	CS	Defined as a distinct responsibility, separate from pitching	Usage data tracked for early converts; low-usage accounts flagged for onboarding nudge
61–90 Days	Net-new AEs begin using refreshed case studies in competitive deals against Ramp/Airbase, sorted by buyer-situation bucket (no workflow / overserved / underserved)	AEs/SDRs	Ongoing, embedded	Net-new pipeline entries tagged by bucket in CRM
61–90 Days	Mid-launch review: compare actual attach count and MRR mix against the three-scenario table (disciplined $38K / even $26K / undisciplined $17K)	Marketing + sales leadership + Finance/Ops	Half-day review	Decision made on remaining $5K contingency allocation (additional enablement vs. additional incentive)

61–90 Days	Redirect AE effort back to Tier 1 accounts specifically if Month 3 data shows conversions skewing toward low-seat-count Tier 2/3 accounts	Sales leadership	Coaching/ reprioritization, no new budget	Tier-mix trend visibly correcting toward Tier 1 weighting in pitch-status data
61–90 Days	Quarterly competitor pricing-page scan (Ramp, Airbase, Expensify, Zoho) as a standing low-cost task	Marketing (either person)	<1 hour	Scan logged; any material competitor move flagged to sales leadership

Cross-cutting notes that apply across all three horizons:

The two hard gates in this plan are non-negotiable by design, not arbitrary process. The **Week 2 go/no-go checkpoint** — segmentation complete, battlecard reviewed, tracking field live — exists because sending unsegmented outreach before this foundation is in place is the direct path to the diluted, undisciplined-targeting outcome quantified earlier in this report (the ~$17K MRR scenario rather than ~$38K). The **Week 8 and Week 13 checkpoints** exist because the report's central risk finding — that sales bandwidth, not awareness, is the constraint — is only actionable if pitthe plan has an early warning mechanism built in, not just a Month 6 postmortem. Waiting until the final readout to discover that Tier 1 pitch coverage stalled at 20% would leave no runway to correct course; catching it at Week 8 leaves four-plus months to escalate, redeploy the SPIF earlier than planned, or have sales leadership personally intervene on specific accounts.

The table below consolidates the sequenced actions into the requested 0–30 / 31–60 / 61–90 day structure, mapping each action back to the specific recommendation it operationalizes from earlier in this report.

Horizon	Action	Owner	Effort	Milestone (Done =)
0–30 days	Segment all 340 accounts into Tier 1/2/3 using CRM data (employee count, org structure, contact role), spot-checked against AE/ CSM account knowledge	Marketing (CRM-lead)	3–4 hours + spot-check	Tiered account list loaded into CRM, sampled for accuracy
0–30 days	Draft and distribute one-page positioning brief with three-tier message hierarchy (hook / objection-handler / competitive closer)	Marketing (both)	2–3 hours drafting, 30-min readout	Brief distributed to all 6 AEs/SDRs and 3 CSMs

0–30 days	Build sales battlecard: corrected ~87–90% minimum-commitment price gap, objection scripts, two net-new qualifying questions	Marketing, with AE review	~1 week, part of $20K enablement budget	Battlecard v1 reviewed and approved by at least one AE
0–30 days	Add "Approvals pitch status" field to CRM; announce tracking cadence in existing weekly pipeline call	Sales leadership	Under 1 hour setup	Field live, tracked in first weekly call
0–30 days	Build and launch in-app banner announcing Approvals to all 340 accounts	Marketing + Engineering	Under 1 sprint day	Banner live and visible to logged-in users
0–30 days	Send Tier 1 email sequence (Email 1: hero metric + AE/CSM walkthrough invite) and Tier 2 qualifying email	Marketing	~1 week combined copy + send	Both emails sent; open/click data captured
0–30 days	Finalize Slack/Teams notification copy, including department-head message variant	Marketing + Engineering	2–3 hours copy	Copy ships with notification feature
0–30 days	Begin AE/CSM outreach to Tier 1 accounts with renewal/QBR dates in next 60 days, embedded into existing calls	AEs/CSMs	Ongoing, no new meetings	First Tier 1 pitches logged in CRM tracking field
0–30 days	Send Tier 1 Email 2 to non-openers/non-responders	Marketing	Half day (copy reused)	Email 2 sent by end of Week 3–4
31–60 days	Build and publish Approvals landing page plus first blog post, leading with corrected price-gap comparison	Marketing	3–4 weeks content production, ~$15K budget	Landing page and one post live
31–60 days	Roll net-new qualifying questions into all AE discovery-call scripts (not just Approvals-specific calls)	Sales leadership	Script update only	Questions in use on all new prospect calls

31–60 days	Continue systematic Tier 1 outreach at each account's natural touchpoint; track weekly in existing pipeline review	Sales/CS, monitored by Sales leadership	Ongoing, embedded	Pitch-status field shows rising "Pitched" count
31–60 days	Week 8 checkpoint: pull pitch-status data — flag if under 50% of Tier 1 accounts show "Pitched"	Marketing + Sales leadership	1-hour review	Checkpoint held; escalation triggered if under threshold
31–60 days	Spot-check and correct any mis-tiered accounts based on live AE conversations, updating CRM in real time	AEs/CSMs	Ongoing, minimal incremental effort	CRM tier field reflects corrections, not just Week 1 draft
61–90 days	Pull results from first 8–10 non-beta customer conversions (turnaround time, seat count, quotes)	Marketing + CS	1–2 days	Refreshed proof-point data set assembled
61–90 days	Update battlecard to v2 with fresh case studies, distribute to full sales team	Marketing	Half-day update	Battlecard v2 in active use
61–90 days	Activate $50–100/ account SPIF for AEs/CSMs, funded from $5K contingency line	Sales leadership + Ops	Low — policy + payroll setup	SPIF live and communicated to team
61–90 days	Hold Month 3 checkpoint: compare actual account count and blended MRR/ account against the three revenue scenarios ($38K/ $26K/$17K)	Marketing + Sales leadership + Finance/Ops	Half-day review	Mix data reviewed; corrective action assigned if skewing toward low-value tiers
61–90 days	CS begins explicit usage monitoring (not just billing status) on first 8–10 converted non-beta accounts	CS	Ongoing, defined as separate task from pitching	Usage data logged for each account, low-usage accounts flagged

61–90 days	Reallocate any underspent budget line (per Week 18-style check pulled forward) toward enablement or a second content asset	Marketing + Sales leadership	Half-day	Budget reallocation decision documented

Two sequencing principles govern this table and should guide any deviation the team makes in execution. First, nothing customer-facing happens before the Week 1–2 foundation (segmentation, positioning brief, battlecard, tracking field) is complete — this ordering is what separates a targeted, tier-aware launch from the undisciplined broadcast that the revenue-math analysis shows could cost the company $15–20K in MRR even while hitting the account-count target. Second, every action in the 31–60 and 61–90 day columns depends on an artifact created in the 0–30 day column (the battlecard gets refreshed, not recreated; the landing page reuses battlecard pricing language, not new copy; the SPIF activates on top of an already-live tracking field, not a new system) — this reuse is deliberate, since the 2-person marketing team and 6-person sales team have no slack to build parallel, redundant assets across three separate 30-day sprints.

The owners named throughout — Marketing (2 people, unnamed by role split), Sales leadership, individual AEs/CSMs, and Engineering for the two small technical asks (banner, notification copy) — reflect the actual team that exists per the intake, not an idealized launch team. Where effort estimates look small (hours or a single day) relative to the six-month scope of the plan, that is intentional: the report's central finding is that this launch succeeds or fails on sustained execution of low-effort, high-frequency actions (a weekly CRM check, a renewal-call pitch, a monthly mix review) rather than on any single large undertaking. The milestones in the rightmost column are written as binary, checkable states — a field exists or doesn't, an email is sent or isn't, a checkpoint is held or isn't — specifically so that a 2-person marketing team and a Sales leadership function without a dedicated launch manager can self-audit progress without needing a formal project-management layer this engagement's scope and budget don't support.

Appendix: Methodology & Assumptions

Appendix: Methodology & Assumptions

Methodology. This analysis was built as a structured, four-part sprint moving from external context to internal execution: (1) research and market context, establishing category benchmarks and comparable-company patterns to sanity-check the client's own claims against industry norms; (2) findings and analysis, where client-provided numbers (attach-rate target, budget, team size, pricing) were stress-tested using a funnel model, a segment-fit gap analysis, and a revenue-mix sensitivity table; (3) the four in-scope deliverables (market/segment analysis, positioning, channel plan, launch timeline), each built to operationalize a specific finding rather than restate generic GTM best practice; and (4) risk mapping and a 90-day action plan translating the strategy into owner-level, checkable tasks. Throughout, the guiding method was to take every number the client supplied at face value first, recompute it against a named benchmark or model, and flag explicitly wherever the recomputation produced a materially different picture (e.g., the 40% vs. ~87–90% pricing-gap correction, and the 15%-of-340 vs. 28%-of-ICP reframing). This "recompute and flag" method was chosen deliberately given the sprint format and the explicit exclusion of creative production and media buying — the highest-value output achievable in scope was sharper targeting and sequencing logic, not new creative assets or paid-channel strategy.

Labeled assumptions. The following inputs were not present in the client intake and were assumed for modeling purposes. Each is called out here for visibility and correction:

- **Even distribution of the 340-account base across the 50–500 employee range.** No segment-level breakdown was provided, so the ICP-fit sizing (Tier 1 "H120–140, Tier 2 "H100–120, Tier 3 "H80–100) is a modeled estimate based on assumed uniform distribution, not actual CRM data. This is the single most load-bearing assumption in the report — nearly every downstream recommendation (tiering, messaging emphasis, revenue-mix scenarios) depends on it.
- **Email open rates (35–45%) and in-app view rates (50–60%) within 30 days.** These are general B2B SaaS customer-list benchmarks, not Ledgerly's actual historical email/in-app engagement data, which was not provided in intake.
- **AE/CSM active-pitch coverage of 60–70% of the base within six months.** Assumed based on typical rep capacity (roughly 38 accounts/rep) and general renewal-cycle cadence; Ledgerly's actual meeting cadence, renewal-date distribution, and rep utilization were not supplied.
- **20–25% close rate for pitched accounts.** Drawn from general SaaS cross-sell/attach benchmarks, not Ledgerly-specific historical upsell conversion data (which may not yet exist, given this is the company's first add-on).
- **Market growth rate (12–15% CAGR) for the SMB expense-management category.** Explicitly flagged in-line as directional industry context, not a sized TAM for this engagement — no market-sizing data or analyst subscription was part of the intake or scope.
- **CRM data completeness** (employee count, org structure, contact role) sufficient to support Tier 1/2/3 segmentation. The intake states this data is "likely already in the CRM" but does not confirm its accuracy or recency; the report treats this as an assumption requiring a spot-check, not a verified fact.
- **Sales team compensation structure** does not currently reward add-on attach (assumed, since intake did not describe comp mechanics), which is the basis for recommending a SPIF as a mitigation.
- **Beta-cohort representativeness.** The 4.5-days-to-under-8-hours metric is treated as directionally trustworthy but potentially optimistic versus a broader rollout, since beta customers are typically more engaged than the average account — this is a modeling caution, not a claim the number is wrong.
- **Landing page conversion rate (2–4%) for the net-new inbound channel**, drawn from general B2B SaaS benchmarks rather than Ledgerly's own historical landing-page performance data, which was not provided.

Benchmark sources referenced. Category and funnel benchmarks cited throughout (SMB finance-function approval-cycle norms, SaaS expansion-vs-acquisition cost ratios, cross-sell attach-rate ranges, landing-page conversion norms) were drawn from generally recognized industry benchmarking bodies and publications referenced by name in the report — APQC and GBTA-style finance-function studies, OpenView's SaaS Benchmarks, Vendr pricing data, and Gartner finance-function buyer research — used as directional calibration points, not as licensed, primary-sourced datasets pulled for this engagement. Comparable-company patterns (Gusto, Ramp's own product-expansion history, Vanta/Drata's beta-metric-led launches, Freshbooks/Wave's price-anchoring against Intuit) were used illustratively, based on publicly observable GTM patterns in the B2B SaaS category, not confidential competitive intelligence.

Confidence Notes

Several pieces of this analysis are meaningfully more load-bearing and less certain than others, and would benefit most from direct client input before being treated as final:

- **Highest priority for follow-up:** the assumed even distribution of the 340 accounts across the employee-size range. A real CRM pull of actual employee counts, department-head structure, and current plan tier would immediately replace the modeled Tier 1/2/3 sizes with real numbers and should be done before Week 1 execution, not after.
- **Second priority:** the 20–25% close-rate assumption and the 60–70% pitch-coverage assumption underpinning the 40–60 account funnel range. Neither is grounded in Ledgerly's own historical upsell data, because this is a first-of-its-kind add-on launch — actual performance in Months 1–2 should be used to recalibrate these figures quickly rather than treating the funnel model as fixed.
- **Thinner and more speculative:** the 12–15% category CAGR figure and the competitive-response risk assessment for Ramp/Airbase. Both are included for context and completeness but rest on general industry pattern-matching rather than Ledgerly-specific intelligence, and should not be used as inputs to financial planning.
- **Reasonably solid, lower risk:** the pricing-gap recalculation (~87–90% at minimum commitment) and the message-hierarchy logic, since both are derived directly and arithmetically from numbers the client explicitly provided, with no external assumption layered in.