SHOPIFY APP
GTM STANDARD

SHOPIFY APP GTM STANDARD

Winning merchants and building recurring revenue.

ALERT COMMERCE

Shopify App GTM Standard: A Practical Operating System for Winning Merchants and Building Recurring Revenue

Copyright © 2026 by Alert Commerce

All rights reserved. No part of this publication may be reproduced, distributed, or transmitted in any form or by any means, including photocopying, recording, or other electronic or mechanical methods, without prior written permission, except in the case of brief quotations embodied in critical reviews and certain other noncommercial uses permitted by copyright law.

Disclaimer

This book describes go-to-market practices for software businesses selling to Shopify merchants. Figures identified as planning examples are arithmetic illustrations chosen to show a method. They are not Shopify benchmarks, industry averages, or promised results. Founder experiences are identified as such and are reported by the people involved rather than independently audited. Platform rules, interfaces, and documentation change; readers should confirm current requirements with the platform before acting.

The author and publisher assume no responsibility for errors or omissions, or for any loss or damage arising from the application of the information in this book. Nothing here is legal, accounting, tax, or financial advice. Readers are responsible for their own decisions and results, and should seek qualified professional advice where appropriate.

Shopify is a trademark of Shopify Inc. This book is not affiliated with, endorsed by, or sponsored by Shopify Inc. Other product and company names are the trademarks of their respective owners.

First edition

Contents

Introduction: Put a number behind the plan 1

1 Give Yourself Enough Runway to Learn 4

2 Pick the Merchants You Can Win 9

3 Choose How Merchants Will Buy 16

4 Give Merchants a Reason to Choose You 22

5 Make Your App Store Listing Sell 29

6 Build a Website That Produces Customers 35

7 Turn Installs into Paying Customers 40

8 Use Advertising to Test and Grow Demand 47

9 Find Merchants with a Reason to Listen 54

10 Run Burst Prospecting 61

11 Build Content That Helps Merchants Buy 68

12 Run a Demo That Earns a Decision 74

13 Keep Deals Moving or Close Them Out 79

14 Make Events Worth the Cost 85

15 Build Partnerships After You Have Traction 92
16 Keep the Customers You Worked to Win 98
17 Run Your GTM Every Week 104
18 Build Your Next 90 Days 111
Appendix A The numbers founders need 117
Appendix B The weekly operating sheet 123
Appendix C The GTM STANDARD assessment 126
Appendix D Templates founders can use 128
Appendix E A 30-minute weekly review 133
Appendix F Operating plans by MRR stage 134
Appendix G Troubleshooting the GTM funnel 140
Sources and attribution 143
About Alert 146

Put a number behind the plan

"We need more installs" is a direction. It is not an operating plan.

How many installs? From which merchants? At what cost? How many become paying customers? How much revenue remains after customers cancel?

You can spend a month rewriting a website, sending emails, and showing up at events without being able to answer a single one of them. The work feels like go-to-market because it looks like what other software companies do. What is missing is any line connecting the effort to the revenue.

This book connects market selection, acquisition, sales, onboarding and retention into one weekly operating system. What you should be holding at the end is a plan you can actually run: a revenue goal, the customer and install counts that support it, the activities you will perform, and a scorecard that tells you what happened.

The method runs backward. Start with the MRR you want to add and estimate how many new paying merchants that takes. Work back from there through paid conversion, activation, installs, and whatever precedes those in your case, whether that is conversations or visits. Only then set an activity plan your time and budget can support. Measure what actually happens, and change the plan.

None of this makes growth predictable, at least not at first. What it does is make the uncertainty visible, which is more useful than it sounds. When the number misses, you can see whether the problem is too few qualified merchants, an offer nobody wants, a signup process that loses people, or customers leaving before you have earned back what it cost to acquire them.

The shape of the answer varies enormously. I have seen a drop-shipping app reach roughly 15,000 downloads with the App Store as its primary channel and six figures spent on advertising, where direct selling barely featured at all. I have also seen a business where every single merchant arrived through a conversation, a demo, and hands-on setup. Both needed a clear buying path. They needed wildly different amounts of human effort to run it.

Your job is not to copy either one. It is to build the motion your product and your economics can actually support.

The GTM STANDARD

FIGURE 1

	STEP AND DECISION	WORKING ASSET
S	**Select your market** Who can you serve well and profitably?	Merchant profile
T	**Tune your offer** Why should they choose you?	Outcome and switching offer
A	**Assemble your foundation** Can they understand, trust, and buy?	Listing, website, and onboarding
N	**Name the signal** Why is this account relevant now?	Evidence-based target list
D	**Drive conversations** How will you create qualified demand?	Weekly acquisition activity plan
A	**Advance the decision** How will interest become commitment?	Demo and follow-up process
R	**Retain the merchant** How will customers get ongoing value?	Activation and retention process
D	**Diagnose and improve** What should change next?	Scorecard and 90-day plan

These are connected decisions, not eight departments to staff. One founder can hold all eight. A product-led app may carry merchants through several of them without a single conversation, while a sales-led app needs an owner standing at every step.

Read it in order once. After that, use the chapters against whatever constraint is in front of you. Every number in a planning example is invented to make the arithmetic followable, so replace them with your own as evidence arrives. The test of whether it worked is simple: at the end you should be holding completed tools, not a list of ideas to consider someday.

1

CHAPTER

Give Yourself Enough Runway to Learn

The first customers are the hardest ones you will ever get. A working product proves that the thing runs. It does not prove that a market exists for it, that the market understands what it does, or that anyone in it will pay a price that keeps your business alive. Those are three separate questions, and shipping answers none of them.

The expensive mistake is spending as though they have been answered, and it rarely looks like a mistake at the time. It looks like confidence. A founder expects fast growth, hires an agency, sponsors an event and signs a long software contract in the same quarter. The customers then arrive slowly, the way customers usually do, and the business has less money and no clearer idea of whom it should serve.

Early go-to-market has two jobs, and they are not the same job. One is to produce customers. The other is to produce evidence you can actually use. Some of what you try will fail, which is rather the point, and your spending has to leave room for it.

Separate the build from the weekly work

Some of this work is construction and some of it is operation. Confusing the two is expensive. Your listing, website, product video, screenshots, positioning, and onboarding all demand real effort upfront, and once they are usable they stop being projects and become assets you maintain.

The rest never stops. Outbound and content have to be produced week after week. Advertising needs a budget watched and results interpreted. Demos and follow-up generate more work precisely when things are going well.

The distinction matters because "we spent the week on marketing" covers both situations equally well. You might have been finishing a foundation you genuinely need. You might also have been polishing it for the fourth time because polishing is more comfortable than talking to a merchant who could say no. Give the first version a finish line and hold yourself to it. A founder rewriting the same headline in the third week running, with no conversations booked, is postponing the only thing that could tell them what it should say.

Make runway visible

PLANNING EXAMPLE

You have $48,000 allocated to the business and hold $12,000 of it back as reserve. That leaves $36,000 to operate on. At $6,000 of net monthly cash burn, you have six months.

Runway = available operating cash ÷ monthly net cash burn.

Count cash in and cash out, not MRR. MRR is a revenue measure, and nobody will accept it in payment. An annual prepayment is the clearest illustration: it helps your bank balance this month while committing you to serve that customer for the rest of the year. Appendix A carries the spending rules worksheet, three runway scenarios, and the format for a learning budget.

Through the first ten paying merchants, buy learning rather than leverage. Sit in on the objections. Watch someone attempt setup without you rescuing them. Read the cancellation reasons in full. You cannot pay an agency to scale a buying process you have never watched end to end.

The first milestone is a working transaction

What you are looking for in the first batch is a transaction that repeats: similar merchants, buying for similar reasons, reaching value, and paying through a process you could run again next month without heroics. Ten is a convenient number to review, not a threshold anyone has crossed into product-market fit.

For each one, record the problem, the reason they acted when they did, what nearly stopped the evaluation, what setup cost them, and what they actually use now. Then be honest about the mix. Five installs from friendly contacts and one heavily customized contract answer two different questions, and neither of them is whether acquisition repeats.

Plan by revenue stage

FIGURE 2

STAGE	PRIMARY QUESTION	USEFUL OPERATING EVIDENCE
$0	**Will a target merchant pay for this outcome?**	Qualified conversations, activated trials, first payments
~$1K	**Can we repeat the first wins?**	Similar buyers, repeated objections, a usable buying route
~$5K	**Can we acquire customers without losing control of costs?**	Cohort conversion, contribution, workload, early retention
~$10K	**Can the process operate consistently?**	Reliable ownership, recurring acquisition, retention and capacity

The same event means different things at different sizes. At $0, a cancellation is a lesson you can afford. At $10K a month, a pattern of cancellations quietly invalidates the acquisition budget you approved last quarter.

What these stages will not do is tell you when you arrive. "We should reach $10K in 90 days" means something only once you can show the customers, the funnel, the conversion lag, and the capacity underneath it. Without those, it is an aspiration wearing a forecast's clothes.

A first month operating brief

The first month has three jobs and a backlog. Finish the minimum buying foundation. Choose one narrow merchant segment. Run a single bounded acquisition test. Everything else that occurs to you across those four weeks goes on a list and waits there.

A reasonable brief might be ten conversations with customers or prospects, one live listing, one signup route that works, and one recorded demonstration. You can commit to the conversations. You cannot commit, on a merchant's behalf, to the payment.

At month-end, write down what you now know about fit, value, and buying friction, and then decide what to repeat. That second half is what separates learning from note-taking. Nothing you learned counts commercially until it changes what you do or what you spend next month.

COMPLETE THIS CHAPTER

Record available cash, monthly burn, runway, and a test spending ceiling. Choose one expense you will defer until you have stronger customer evidence.

MEASURE

Actual cash spent, remaining runway, merchant conversations completed, and paying customers acquired.

2

CHAPTER

Pick the Merchants You Can Win

"Shopify merchants" is a platform filter. It tells you where someone sells and nothing else, which makes it close to useless as a target market.

A merchant selling subscription food has almost nothing in common, operationally, with one selling custom furniture. A two-person team buying a simple utility app goes through a different process entirely from a larger operation replacing a system that holds years of customer data.

So the starting question has to be narrower. Which merchants have a problem we solve, already know it matters, and can buy from us without requiring that we rebuild the company around them?

Build your merchant profile from evidence

If you already have customers, start there. The pattern you are hunting sits across five things at once: who activates quickly, who pays, who stays, who actually gets a result, and who needs a level of support you can afford to give.

Write down characteristics you can observe, not ones you can only feel. “Ambitious brands” cannot be verified by anyone, including you. “Subscription merchants running a particular app and processing enough orders to justify our price” can be checked before you send a single email.

WORKSHEET 1

PROFILE FIELD	YOUR ANSWER
Merchant type or vertical	
Problem being solved	
Operational size or volume	
Existing app or workaround	
Required technical fit	
Likely buyer and user	
Expected monthly revenue to us	
Evidence that this merchant needs the product	
Reasons to exclude an account	

Defining the customer and building a list against that definition is foundational in both Predictable Revenue and Hacking Sales. Shopify adds unusually good raw material: installed technology and visible workflow make the selection concrete.

Use a fit gate before a priority score

Scoring has a specific failure mode. Strong marks on three dimensions quietly cancel out a fatal weakness on the fourth, and a bad account arrives at the top of your list looking respectable. Ask the binary question first. Can the product serve this account at all?

Three things fail an account outright: a required capability you do not have, expected revenue that cannot cover what serving them costs, and requirements that pull against where the product is going.

Only once an account clears that gate does ranking make sense. Three fields scored 0 to 2 are enough: fit, relevant evidence, and access to the buyer. The totals order your list, and the order is a working judgment you revise the moment results disagree with it.

Make the account list usable

PLANNING EXAMPLE

Research 40 accounts. Fifteen fail the fit gate and come out, leaving 25 relevant merchants. That is a starting batch, not a benchmark.

Appendix D1 has the account record. One thing is worth doing before you count anything: deduplicate companies running several stores, so the account total reflects buying entities rather than storefronts.

Separate buyer fit from store size

Pick the operating variable your app actually depends on: return volume, booking complexity, catalog size, team workload. GMV is a proxy for all of them and a direct measure of none, so it earns its place only where it approximates the variable you care about.

Where the real variable is unavailable, label the proxy as a proxy. A large merchant without the problem is a worse prospect than a small one feeling an urgent version of it, and size alone should never be what triggers a high-touch sales process.

The problem and proof interview

Early conversations are for improving the profile, not for closing. Ask how the workflow runs today, what goes wrong in it, how often that happens, who deals with it, and what fixing it would be worth.

What you are listening for is episodes. "Reporting is a pain" gives you nothing to act on. "Every Monday our operations manager combines three files before we can make a decision" tells you who feels the problem, how often, and exactly where a product could slot in.

Then ask what they have already tried. A competing app, a spreadsheet somebody maintains, a contractor, a manual process everyone tolerates: each tells you what the problem is already worth to them. An account that has tried nothing is ambiguous rather than promising, because the problem may be new, it may be unrecognized, or it may simply not matter much.

Record the words they use. When several buyers describe the same outcome in similar language, that language belongs in your listing and your outreach. Keep one distinction in view while you do it: what people say they want and what paying customers adopt are two different data sets.

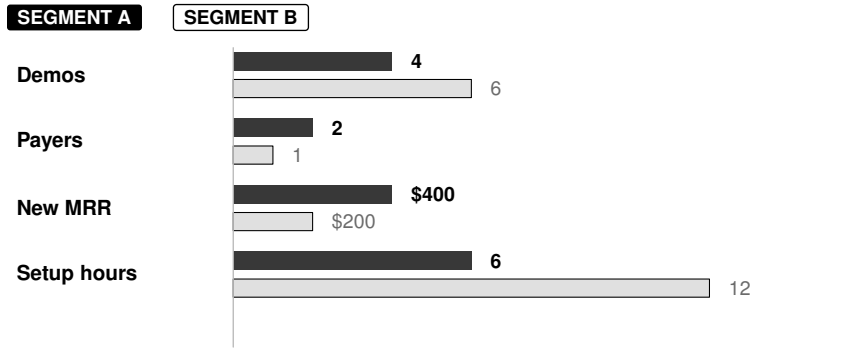
Build a segment comparison

PLANNING EXAMPLE

You test two groups of 20 accounts. Group A produces four demos, two paying merchants, and three hours of setup per payer. Group B produces six demos, one payer, and 12 hours of setup for that single account. Both pay $200 MRR.

FIGURE 3

Two segments compared on the same research effort



Group B generates more meetings and less revenue, and costs four times the delivery work per paying customer to get there. That is enough to go and find out why the buying and delivery process differs so sharply.

Choose an initial niche without becoming trapped

Start where your advantage is easiest to explain and easiest to verify. Once customers begin repeating, test an adjacent segment with a

label and a budget of its own. A workflow proven with subscription food merchants may transfer to another subscription category, and until you have run it that is a hypothesis rather than an expansion plan.

Record four things each time: the assumption, the sample, the result, and the decision you made. Keep the proven segment visible while you test the next one, because the most common way to lose a working niche is to become interested in a new one.

Map the buying group

The person who replies is not necessarily the person who will use the app, and neither of them is necessarily the person who approves it. At a small merchant all three are the same founder. At a larger one, an operator evaluates, finance approves, and an agency implements. The question that gets you there is who needs to be comfortable with this change, which works considerably better than demanding "the decision maker" before you understand the business.

Three roles are worth recording where they exist: whoever owns the outcome, whoever operates the thing daily, and whoever can authorize the spend. One name often fills two or three of them. The purpose is to avoid a surprise in week six, not to build an enterprise account map for a $50 app.

Keep a rejection log

Use a fixed set of reasons so the log can actually be counted later: technical mismatch, economics, no relevant need, missing capability, service burden, timing. Review the repeats monthly, because each cause points somewhere different. Missing capabilities are a product decision. Inaccurate inputs are a research problem. Unsuitable stores mean your filters need work.

COMPLETE THIS CHAPTER

Define the merchant profile and exclusions. Fit-check 25 accounts, explain each selection, and flag repeated custom exceptions.

MEASURE

Accounts researched, eligible accounts, disqualification reasons, and eventually paid conversion and retention by segment.

3

CHAPTER

Choose How Merchants Will Buy

Acquisition channels and buying motions get discussed as though they were the same decision. They answer different questions.

The App Store, outbound, content, ads, events and partners all describe how a merchant finds you or ends up in a conversation. Product-led, sales-led and hybrid describe something else entirely: how they evaluate, adopt and pay once they are there.

The two mix freely. An app can live on the App Store and still need a salesperson for its larger merchants. A merchant can arrive through a cold email and activate without ever seeing a demo. Treating either category as an identity you adopt is how founders end up forcing a motion their customers never asked for.

Pick the shortest credible buying path

FIGURE 4

Three buying motions



For most founders reading this, the App Store is the foundation, with a self-serve website route alongside it wherever the product supports one. The objective is not to eliminate human help. It is to remove the friction that buys you nothing and keep the help where it genuinely changes the outcome.

Match effort to customer value

PLANNING EXAMPLE

An app earns $50 MRR per merchant and produces $40 of monthly contribution before acquisition costs. If winning the account takes $300 of founder time and onboarding work, recovering that takes 7.5 months, before churn enters the picture.

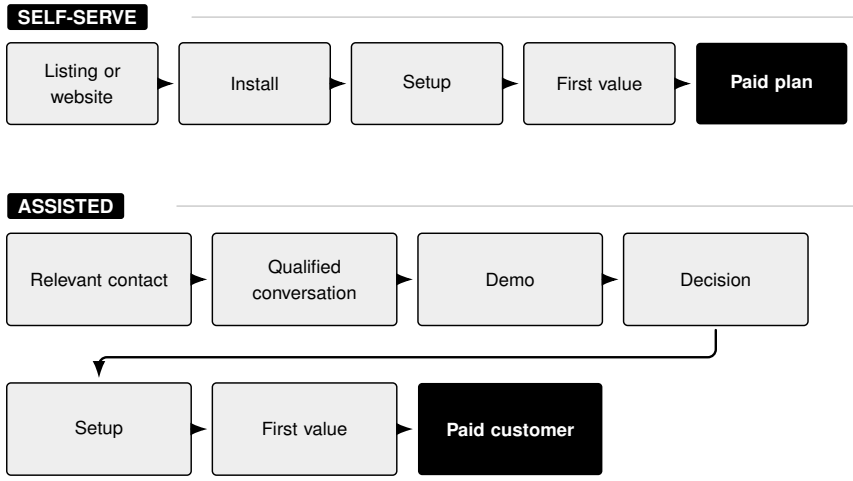
At $500 MRR and $400 of contribution, that same $300 of effort is recovered in three weeks. Nothing about the work changed. What changed is what you can afford to do to win and serve the account.

Preserve a route for each buyer

Write down the paths you actually support, rather than the ones you intend to.

FIGURE 5

The two buying routes



A trial can sit on either path. Define activation and payment as separate events, and make sure every merchant on the self-serve route can find a human without having to work for it.

A channel cannot repair the wrong buying route

Founder help is a reasonable temporary investment in learning. It stops being that when every low-priced install needs a demo. At that point the work is a permanent cost, and it should be labelled and priced before anyone describes the model as scalable.

Model contribution before choosing the service level

Build the contribution estimate from the costs your product actually creates. Hosting is often trivial while support, third-party usage, fulfillment, incentives or claims-related costs are anything but. A blanket software-margin assumption will hide exactly the difference you need to see.

PLANNING EXAMPLE

At $100 MRR, variable costs total $30 a month, leaving $70 of contribution. A $210 acquisition and setup investment gives a three-month simple payback. Let support cost rise by $20 a month and contribution falls to $50, pushing payback out to 4.2 months.

The price never moved. The service model shifted the economics further than a respectable improvement in ad performance would have. Track setup effort and ongoing support as separate lines, because they behave differently: setup is concentrated at the start, while support compounds across every month the customer stays.

Use assistance as an intentional offer

Four things need defining before you offer help: who performs setup, what the merchant supplies, how testing works, and what counts as complete. Reserve the assistance for accounts and blockers your economics can carry. “White-glove onboarding” should name a bounded service; used loosely, it is a promise of unlimited work that somebody will have to do.

Build a route selection worksheet

Five questions per target segment. Can the merchant understand the value without talking to anyone? Can they complete setup alone?

Can they tell whether it worked during an evaluation? Does the price support human help? And does the change carry operational risk serious enough to need a conversation?

The answers point at a primary route without needing to be scored. Weight them by what they are rather than by how many land each way, because two hard setup answers can outweigh five encouraging ones about the website.

FIGURE 6

The shortest credible buying route

SITUATION	PRACTICAL STARTING ROUTE
Low setup effort and clear value →	**Self-serve with accessible support**
Clear value but occasional setup blockers →	**Self-serve with optional assistance**
Material migration or negotiated economics →	**Assisted evaluation and setup**
Small and large merchants buy differently →	**Segment-specific hybrid routes**

Track handoffs in a hybrid motion

Hybrid motions produce a reporting error almost automatically, and it inflates everything downstream. The merchant installs independently, books help, then pays. Self-serve claims the win and so does sales. Both contributed. The company acquired one customer.

Keep three separate fields and the problem disappears. Source is how the merchant arrived. Assistance is what your team did. Payment is the commercial outcome. Separated, they let you answer the question that matters: whether sales assistance improves a segment enough to pay for itself.

Compare assisted and unassisted cohorts, remembering that merchants who seek help are selecting themselves. A higher conversion rate in the assisted group does not prove the call caused it.

Know when the motion should change

Two triggers are worth watching for. Add sales when larger merchants repeatedly need help and can carry its cost. Add self-serve when the setup path becomes repeatable enough to automate.

Before hiring anyone to do either, document the repeated work, the volume, and the process. Chapter 18 covers the delegation itself.

COMPLETE THIS CHAPTER

Draw your primary buying path in words and name every step that requires a human. Then map the last five real buying routes, the hours each consumed, and the first-value event in each, and compare what you observed against the motion you believe you are running.

MEASURE

Time to first value, time to payment, human hours per acquired merchant, and conversion at each step.

4

CHAPTER

Give Merchants a Reason to Choose You

Merchants do not buy your feature roadmap. They have never seen it and would not care if they had. What they buy is an expected improvement in their own business.

A feature earns its place by doing one of four things: increasing sales, reducing work, preventing mistakes, or handing the team a capability it did not have. Your offer has to make that connection clearly enough that the merchant can repeat it to a colleague without you in the room, because that is usually where the decision actually gets made.

One sentence carries most of the weight: "For [merchant type], we help [business outcome] by [specific capability], with [proof or credible demonstration]."

If the only way you can complete that sentence is with words like "powerful" and "all-in-one," it is not finished. Say what changes for the merchant.

The Switching Equation

For a merchant already running another app, price is the smallest part of the decision.

The expected benefit has to clear three things at once: the effort of migrating, the financial cost, and the operational risk of the change going wrong.

The useful move is to ask what would actually have to move: data, automated flows, points balances, staff habits, promises already made to customers, and every connected system nobody has thought about yet.

A simple utility can be swapped in an afternoon. Email marketing and loyalty drag a long tail of dependencies behind them. Even inside one category the switching costs vary enormously, and assuming otherwise is how a two-week migration becomes a two-month one.

Build an offer around the actual obstacle

FIGURE 7

Build the offer around the actual obstacle

OBSTACLE		POSSIBLE RESPONSE
Merchant lacks setup time	→	Assisted setup with clear responsibilities
Remaining contract makes switching expensive	→	A bounded financial incentive
Team fears migration failure	→	Defined migration plan, testing, and ownership
Value difference is unclear	→	Demonstration and merchant-specific economics
Needed capability is missing	→	Honest qualification and a product decision

PLANNING EXAMPLE

A merchant pays $100 a month and you offer a $200 onboarding credit. Record it as acquisition cost or as a revenue reduction, depending on how your business accounts for it. Either way it reduces what the deal is worth.

Set boundaries on custom work

Four things decide whether custom work is worth doing: expected revenue, delivery effort, the ongoing support it creates, and whether any other merchant will ever use it. When enough of those answer badly, decline. Chapter 9 has the case; Chapter 12 covers handling a feature gap while the sale is still live.

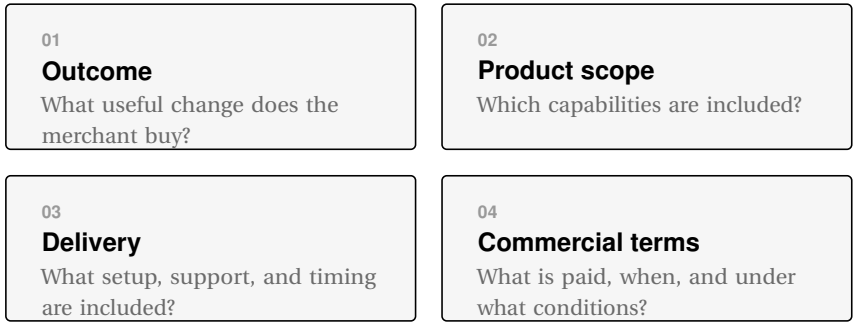
Write the outcome before writing the feature list

Take each feature and finish the sentence "This matters because the merchant can now _____." Whatever goes in the blank has to be an operating or financial consequence they would recognize.

Then name the evidence behind it: a demonstration, a customer result, observable product behavior, or an estimate built for that merchant. Match the strength of the claim to the strength of what you have. A capability that works is not evidence for a promised percentage increase in revenue, however reasonable the number sounds.

Build the offer in four fields

FIGURE 8



Write the summary so it survives being forwarded. Phrases like "setup included" and "no subscription fee" get read generously by someone who was not on the call, so define them before a buyer concludes you have promised unlimited service.

Use incentives with an economic ceiling

Every incentive should have a job. It can offset a switching cost, accelerate a decision that is already qualified, or fund a bounded pilot. Pick the smallest credible version that addresses the obstacle actually in the way.

PLANNING EXAMPLE

A $200 MRR account produces $150 of monthly contribution, and your acquisition and onboarding ceiling is $600. Research and sales effort have already consumed $250, and setup is estimated at $200. That leaves $150 of room under the ceiling for an incentive.

If the buyer asks for $500, there are two honest answers: change the model or decline. The arithmetic will not tell you how to negotiate. What it does is stop the final concession from quietly vanishing out of the business case.

Make the switching plan part of the conversation

Separate requirements from preferences as they come up. A missing data capability can stop the deal outright. A cosmetic adjustment can be scheduled for later. Both deserve a straight answer. They do not deserve the same place in the queue.

Then show a short transition plan: where they are now, where they are going, who owns which piece of work, and the point at which they can judge whether it worked. This is often more persuasive than another feature, for a simple reason. The feature addresses what they would gain. The plan addresses the cost that is keeping them where they are.

Handle common objections with useful questions

"We already use something." Acknowledge the choice rather than arguing with it, state the specific difference that made you reach out, and ask whether a demo is worth their time. None of this requires the prospect to dislike the incumbent first. Most of them do not, and the conversation can still make sense.

"Switching is too much work." Ask which part is the burden, then match the support to that answer. Reaching for a discount before you know what the work actually involves leaves the real objection exactly where it was, now with less margin attached to it.

"We need a custom feature." Get to the business requirement underneath it and what happens if they do not have it. Then be precise about three things: what works today, what would need building, and what you will not promise. A roadmap item is not an available capability, and presenting it as one buys a problem you will meet again at renewal.

"The price is too high." Four different problems hide behind that sentence: value, cash timing, comparison with something else, and genuine inability to pay. A lower price solves one or two of them. Find out which one you are dealing with before you move.

Define a credible pilot

A pilot needs four things: an outcome, a period, an owner, and a decision rule. Strip any of them out and what remains is a free period, which has a way of becoming an indefinite evaluation nobody ever formally ends.

Something like: "During the agreed evaluation period we will complete setup, observe the specified workflow, and review the result together on this date." Then define what the merchant supplies, and how continuation or payment gets decided.

And do not call a pilot successful because the app is still installed at the end of it. Installation is a stage, not a verdict. What counts is the outcome you agreed on and the commercial decision that follows from it.

COMPLETE THIS CHAPTER

Write a forwardable offer with outcome, scope, delivery, price, three reasons to choose you, and maximum incentive and setup commitments. Check its economics including acquisition and onboarding effort.

MEASURE

Offer acceptance, discount and incentive cost, onboarding effort, paid conversion, and retention by offer.

5

CHAPTER

Make Your App Store Listing Sell

The listing is a buying page, not a brochure. A merchant arrives with a job in front of them and one question about you: can this app do it?

Before they install, four things have to land. They need to recognize the use case, see the product working, understand what it costs, and conclude that installing is worth the twenty minutes.

Start with clarity about who the app is for, what it does, and what using it requires. The video should show a workflow someone recognizes. The screenshots should explain capabilities rather than parade tiny dashboard elements that mean nothing to a stranger.

The listing audit

Score the listing against the six items below, then fix the lowest-scoring one first.

WORKSHEET 2

LISTING ITEM	0	1	2
Merchant fit and use case are understandable	☐	☐	☐
Product outcome is specific	☐	☐	☐
Images make key workflows clear	☐	☐	☐
Video demonstrates the product coherently	☐	☐	☐
Pricing and requirements are understandable	☐	☐	☐
Listing accurately reflects the current product	☐	☐	☐
Total			**/ 12**

0 missing, 1 present but unclear, 2 clear and current. A working checklist, not a validated predictor of performance.

Check Shopify's current listing requirements before you publish changes. The trap worth naming: content that is perfectly acceptable on your own website is not automatically permitted in every listing field. Keep customer testimonials on the website where they belong rather than copying them into restricted listing areas. [W1]

Treat review recovery as account work

Log every negative-review issue with a date, an owner, contact status, resolution, and age. Then do the work: contact the merchant, offer a conversation, fix the problem, confirm the outcome. Appendix D5 has the record.

The public reply should acknowledge the issue and point to a route for help, without exposing anything private about the account. Keep the detailed resolution where it belongs, in the conversation.

Once it is resolved, Shopify permits a neutral, non-forceful request to consider updating the review. One rule has no exceptions: never attach an incentive to adding, changing, or removing a review. [W2]

A usable message: "Thanks for working through this with us. If your experience has changed, would you consider updating your review to reflect it?"

Pick a service target you can actually meet, such as an owner assigned by the next working day, and track both unresolved age and repeated causes. Five reviews about the same setup failure are not a support problem. They are a product fix. Resolution is your team's responsibility; whether the rating changes is the merchant's decision alone.

Establish a maintenance rhythm

While a test is running, review listing performance weekly, and keep maintaining it as pricing, features and merchant questions change. Record a source date for every asset and experiment, so the content that has aged out is easy to find later.

Write a listing brief before producing assets

Document six things before anyone produces anything: the target merchant, the primary problem, three supported benefits, the core workflow, pricing, and the objections. One brief should drive copy, screenshots and video alike.

Show the workflow plainly. A straightforward demonstration beats animation that leaves the product unexplained, and the visual assets should add up to one coherent story rather than six unrelated ones.

Give every asset an owner and a source date, and update it when the interface or the offer changes.

Research demand without treating tools as truth

Third-party tools are good for researching keywords, positions, and competitor changes. Record the query, market, language and date wherever you can. What comes back is evidence for a test, not a finding.

Judge the result by installs, activation and payment. A ranking gain on its own establishes nothing commercial.

Build a listing experiment log

PLANNING EXAMPLE

You believe merchants do not understand the setup process. You revise one screenshot and its explanation, record the date, and compare the cohort that follows against the one before it.

WORKSHEET 3

The listing experiment log

FIELD	YOUR RECORD
Hypothesis *e.g. Setup uncertainty discourages qualified installs*	
Change *e.g. A clearer view of the required setup steps*	
Primary measure *e.g. Qualified visit-to-install conversion*	
Secondary measures *e.g. Activation and paid conversion*	
Context *e.g. Traffic sources, pricing changes, ad activity*	
Decision *e.g. Keep, revise, or revert after review*	

Before-and-after comparisons are affected by traffic mix, seasonality, reviews and whatever else changed that week. Without a system that separates those effects, the honest word for your result is directional, not causal.

Do not stop measuring at installation

PLANNING EXAMPLE

One mature cohort produces 50 installs and 10 payers from 500 visits. Another produces 100 installs and five payers from 1,000 visits. Install-to-paid conversion fell from 20% to 5%. Twice the traffic, half the customers.

Three things are worth investigating in that order: audience, promise, and onboarding. Compare cohorts at equal maturity before concluding anything.

COMPLETE THIS CHAPTER

Audit the listing and give review recovery an owner. Then find the first buying question the listing leaves unanswered, fix it, log the change, and follow the next comparable cohort through activation and payment.

MEASURE

Visits, installs, visit-to-install conversion, paid conversion, and the age of unresolved merchant issues.

6

CHAPTER

Build a Website That Produces Customers

Your website is the only place you control completely. It gives you room to explain the product, answer the questions a listing cannot, and capture demand that never passes through the App Store at all.

Its job is to help a qualified merchant take one next step, whether that is installing, starting a trial, or booking a demo. Choose the primary action for each page rather than offering every visitor five equally prominent options and hoping.

Build around the buying questions

A practical first website answers six questions:

1. Is this for a business like mine?
2. What will it help us do?
3. How does it work?
4. What does it cost?

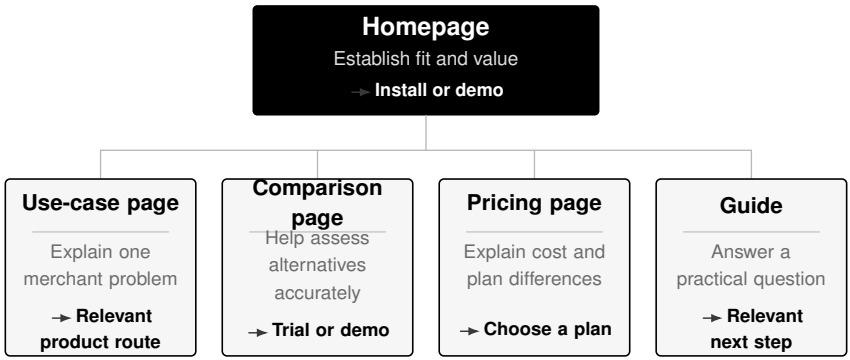
5. What evidence supports the claims?

6. What do I do next?

Answer them in merchant language, using proof you have permission to publish.

Give each page a job

FIGURE 9



Connect the measurement

Track source, landing page, install or booking, and payment dates. Preserve the cohort timing, and label every denominator you report.

PLANNING EXAMPLE

200 visits produce 10 bookings, eight held demos, and two payers. That is 5% visit-to-booking, 80% attendance, and 25% held-demo-to-paid. Payment may land in a later month, so keep the original source attached to it.

Build a minimum buying website

Start with five things: a homepage, a product explanation, pricing, support information, and an install or demo route that works. Add a segment-specific use-case page when a segment earns one. The page-role table and the six buying questions decide what belongs where.

Connect the promise across pages

Match each campaign to a destination that continues the thought. Switching outreach lands on a comparison; use-case outreach lands on that workflow. Keep product scope and pricing consistent across every page, and give temporary offers an owner and an expiry, including the references buried in older content, which is where they survive longest.

Give the demo page a clear job

A sales-led page should explain what the merchant will see and why the conversation is worth having. Booking should not feel like signing up for an undefined sales process.

Ask only for information you will use. A founder can qualify most details inside the conversation itself, and a long form creates friction without reliably preventing poor-fit meetings. After booking, confirm the time, make it easy to change, and keep any preparation proportional.

Build comparison content responsibly

A useful comparison does three things: it names the buyer's decision criteria, shows where each approach fits, and explains the differences that actually matter. It should acknowledge real tradeoffs,

because a comparison with no tradeoffs reads as marketing and gets treated as such.

Date the page and keep the evidence for anything factual. When a competitor's pricing or capabilities change, update it. Where information is not available, say so rather than filling the gap with a sales assumption that will be quoted back at you. The page does not have to declare you the winner for every reader: honest fit improves the quality of the merchants who proceed.

Use SEO as a buying question program

Start with topics your target merchants genuinely need, then organize the related questions around use cases, implementation, costs, alternatives, and operating decisions.

Technical hygiene is a short list: important pages reachable through normal site links, descriptive titles, and everything working on the devices people actually use. Keep someone responsible for broken routes and measurement failures.

Build useful assets before building volume. A large number of thin pages is a distraction from the unanswered questions sitting on the pages buyers already visit.

Establish an attribution policy

A buyer may read a guide, click an ad, and then answer an outbound email. Preserve all three interactions and count the customer once. Where tracking is missing, fix the tracking rather than inventing the attribution.

Diagnose the page before redesigning it

PLANNING EXAMPLE

A use-case page receives 300 visits, produces 15 booking starts, and five completed bookings. That is a 5% visit-to-start rate and a 33.3% start-to-completion rate. The larger loss is clearly inside booking, which still tells you nothing about why.

Check whether the form works, whether the available times suit the audience, and whether what you are asking for matches what the visitor expected. Then choose one change to test. Redesigning the whole site at this point would destroy your ability to learn which obstacle mattered.

Two cases change the priority entirely. If the page receives almost no qualified visitors, conversion optimization is not the problem and audience access is. If visitors convert and the product then fails after installation, the issue belongs in onboarding, and no amount of editing the page will reach it.

COMPLETE THIS CHAPTER

Assign one job and one primary action to each core page, and check the action works from a phone and a desktop. Then take one page and state five things about it: the audience, the buying question, the next action, the conversion event, and the obstruction in front of it today. Give that improvement an owner and a review date.

MEASURE

Qualified visits where assessable, install or booking actions, completed signups, and paying customers by source.

7

CHAPTER

Turn Installs into Paying Customers

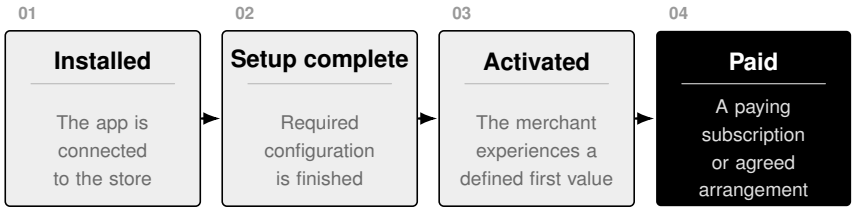
You can buy attention and still lose the merchant the moment they install.

Onboarding is where the promise meets the product. The merchant needs to know what to do, why it matters, and when they have done enough to get something back. A clunky flow makes them decide things they have no basis to decide. A guided one removes the uncertainty and moves them through the essential steps in an order that makes sense.

The Guided Value Path

Four moments are worth defining:

FIGURE 10



Activation and payment do not always happen in that order, so write the real sequence for your product. One thing to avoid: defining activation as clicking a button when no meaningful value follows the click.

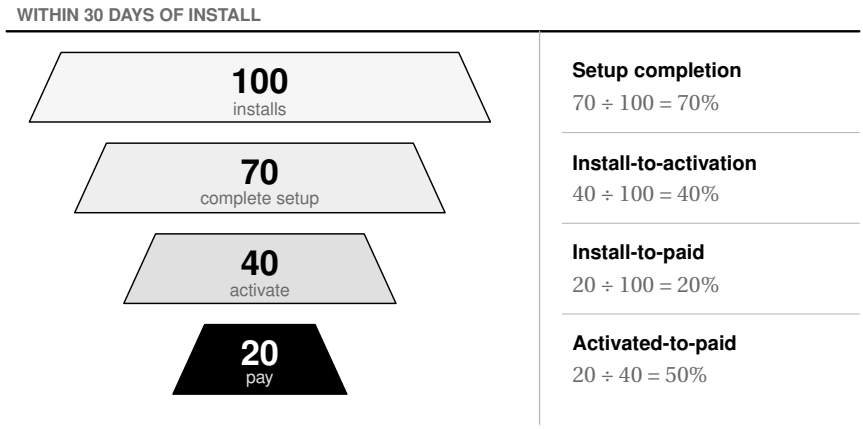
In a founder-reported dropshipping app example, replacing a clunky onboarding flow with a wizard leading toward paid signup shortened time to paid conversion by approximately 90%. Read that carefully. It is evidence about speed, and it does not by itself establish a higher conversion rate or better retention.

Measure where merchants stop

PLANNING EXAMPLE

Within 30 days, 100 installs produce 70 completed setups, 40 activations, and 20 payers.

FIGURE 11



Then go and look at the failing step through observed behavior and merchant feedback. The figures locate a loss. They never explain one.

Decide what the merchant must do first

Separate required setup from optional preferences, and at each step explain the action, its purpose, and what comes next. Preserve progress wherever the platform supports it, and show what is left.

The founder's wizard replaced an unclear route with a guided one, which is the whole mechanism. Dividing the same confusion across six screens helps nobody.

Use the four-event onboarding map

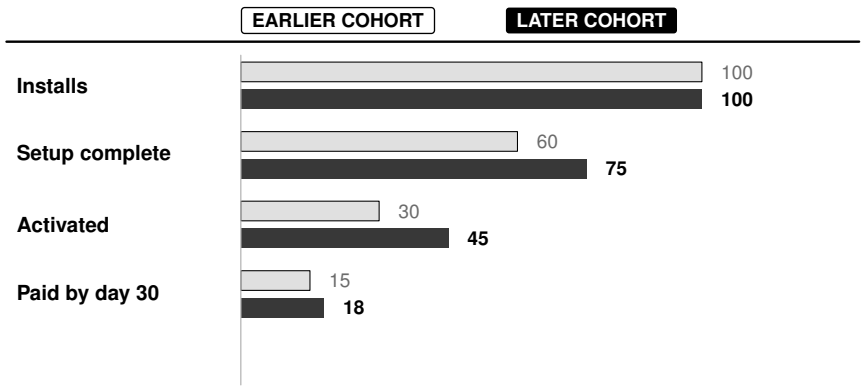
Record install, setup, activation and payment, each with a date. A fifth field for assisted intervention, such as a setup call, is worth adding when it helps explain how the merchant actually got there.

PLANNING EXAMPLE

Out of 100 installs, 60 complete setup, 30 activate, and 15 pay within 30 days. You improve the setup instructions, and the next comparable cohort produces 75 completed setups, 45 activations, and 18 payers.

FIGURE 12

Where merchants stop: two cohorts compared



Every stage improved, but the largest relative gain is activation, which is where you should look first. Before attributing it to one screen, check whether traffic quality, product changes or support also moved. Without controlling for those, the honest description is an observed improvement.

Build a follow-up that resumes the task

A generic “finish setting up” message asks the merchant to remember work they have already forgotten. A useful one names the incomplete task and hands back the route to it.

Something like: “Your store is connected. The remaining step is [required configuration]. Complete it here: [link], or reply if you need help.” Use real account state, and send it through a channel appropriate for customer communication.

A product-led flow can offer help without converting every install into a sales campaign. You are helping someone finish a task they started, not running a sequence that ignores whether they finished it three days ago.

Keep suppression rules for paid customers, completed steps, cancellations, and anyone who asked you to stop. A stale reminder makes a well-built product feel carelessly operated, and that impression is expensive to reverse.

Time to payment is only one view

Report three views together, because each one hides something the others reveal: paid conversion inside a defined window, time to payment among those who converted, and retention afterwards. A fast median across five payers can conceal 95 abandoned installs. A faster charge followed by refunds has improved nothing.

For trials, track whether merchants reach first value before the evaluation ends. Registration on its own tells you almost nothing.

Decide who receives human help

Assistance needs a trigger and a boundary. You might prioritize merchants who fit the profile, have completed part of setup, and are stuck on a known issue. Or you might help every early customer while you are still learning the flow. Both are defensible; neither should happen by accident. Label founder-assisted onboarding while you are in that stage, because if the flow needs a founder for every install you know nothing yet about how the unassisted product performs.

PLANNING EXAMPLE

Ten assisted merchants take 90 minutes each, which is 15 hours. At an internal planning value of $50 an hour, that is $750 of effort. If three become paying customers, assistance has added $250 per payer before any other acquisition cost. The hourly rate is a costing assumption you chose, not a statement about market wages.

Improve the handoff for assisted customers

Sales should hand over four things: the merchant's desired outcome, the agreed scope, known dependencies, and any commitments made. Include what was explicitly excluded, which is the part that causes arguments later. Keep it short enough that someone reads it and specific enough that nobody runs a second requirements interview.

Create one first-value checkpoint with an owner who can answer three questions: is the merchant live, is the intended workflow running, and does the merchant recognize the value. Installation and configuration are internal milestones. The customer outcome is the thing they paid for.

COMPLETE THIS CHAPTER

Define activation, map the essential onboarding steps, and find the largest observed drop-off. Then take five recent installs, map the four events for each, and ask an appropriate sample of the incomplete accounts what blocked them. Choose one change from that evidence and evaluate a comparable cohort against it.

MEASURE

Setup completion, activation rate, paid conversion, median time to activation and payment, and support effort per activated account.

8

CHAPTER

Use Advertising to Test and Grow Demand

App Store advertising has one structural advantage: the merchant is already looking for software when they see you. That does not make every click worth paying for, and it does not make every app economically suited to buying traffic.

The first thing to establish is narrow. Can a defined audience and a defined message produce paying merchants at a cost your business can carry?

Shopify's advertising documentation covers search, category and homepage placements, along with budget and bidding controls. Check placement availability and eligibility against the current version rather than this one. [W3]

Start with a hypothesis you could be wrong about. "Merchants searching for this specific use case will understand our listing and activate" can be tested. "We need to be everywhere" cannot.

Work backward from what you can afford

PLANNING EXAMPLE

Your app earns $100 of monthly revenue per customer and $80 of monthly contribution after variable service costs. You choose a three-month contribution payback target for this test.

Allowable acquisition cost = $80 × 3 = $240 per paying customer.

That is a constraint you chose, and it assumes the contribution holds and the customer stays long enough to deliver it.

If install-to-paid conversion is 20%, allowable cost per install is $240 × 20% = $48. If ad-click-to-install conversion is 10%, the corresponding allowable average cost per click is $48 × 10% = $4.80.

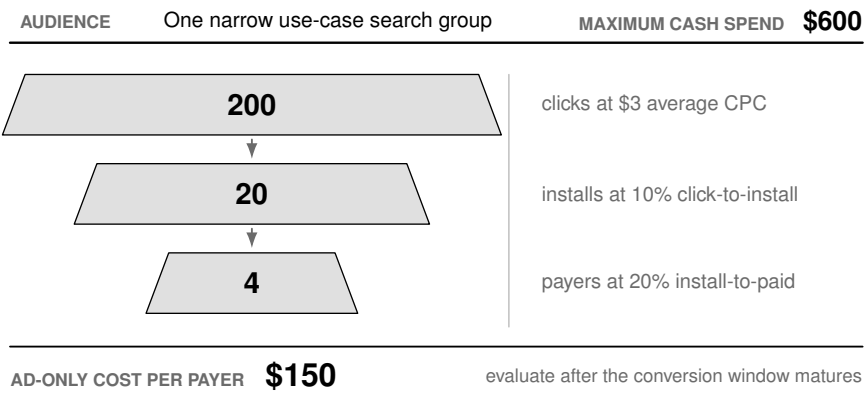
Both rates have to describe the same attributable funnel. Using an organic conversion rate as though it were a verified paid rate is the quickest way to build a confident number that means nothing.

Run a bounded test

A useful test brief contains:

FIGURE 13

A bounded advertising test



Four expected customers is a small sample. Use the result to learn something, not to declare a permanent acquisition cost. Fully loaded CAC will come out higher once creative, tools and labor are counted.

Watch spend and obvious targeting problems while the test runs. A broken destination or an irrelevant campaign should be stopped the day you notice it, not on the review date. The opposite error is just as costly: recent installs are not failures until they have had time to convert.

Add search advertising deliberately

Website search campaigns need three things working together: a landing page matched to the query, a conversion event that fires, and the same connection through to payment. Keep brand demand

separate from new demand, because a campaign collecting people already searching for your company is answering a completely different question from one reaching category buyers.

Organize campaigns around intent

Group the early tests by the merchant's likely job. A specific use-case search behaves differently from a broad category search, and a brand query differently again from someone meeting the category for the first time. Keep the groups understandable enough that the result tells you what kind of demand you reached.

The failure to avoid is mixing unrelated audiences into one small test and then concluding that "ads work" or "ads fail." A result without audience context cannot be reused, which means you paid for it twice.

Set a bid ceiling with the full funnel

The allowable-click figure is a planning ceiling. It is not advice to bid the maximum and find out.

Take a $240 allowable acquisition cost, 20% install-to-paid, and 10% click-to-install. A $4.80 average cost per click consumes the entire allowance before a single dollar of creative or labor. Leaving room for those means the media allowance has to be lower.

> **PLANNING EXAMPLE**
>
> Reserve $60 per acquired customer for creative, management and other acquisition costs, and the media allowance drops to $180. At the same funnel rates, allowable average CPC becomes $180 × 20% × 10% = $3.60.

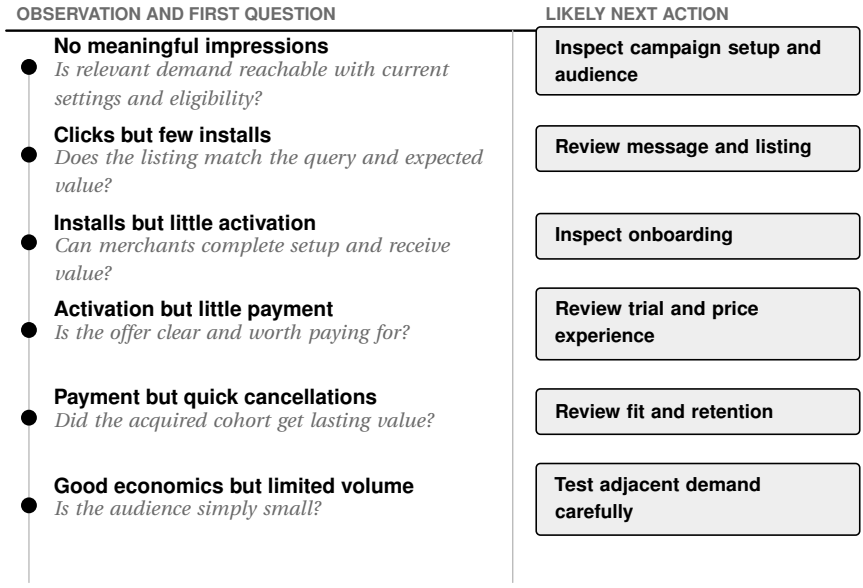
If the market wants $5 clicks, the arithmetic has told you something useful: the model does not fit at current rates. The honest responses

are to improve conversion, price, contribution or the payback constraint, with evidence. The dishonest one is to exclude costs until the number works.

Use a campaign decision table

FIGURE 14

Campaign diagnosis order



Work through it in order, remembering that one thing can imitate every entry on it: a measurement failure. Confirm the events actually fire before you rewrite the strategy.

Understand the cost of a premature conclusion

Suppose a $300 campaign produces 10 installs and no immediate payments. You stop it on day three and write it off as unprofitable. Two weeks later three of those merchants pay, and your ad-only acquisition cost is $100 per payer.

That does not mean the campaign should restart. It demonstrates why the conversion window matters at all. And if you had deleted the source records or reassigned those customers to another channel in the meantime, you would have drawn the wrong lesson with total confidence.

The reverse happens too. A customer pays quickly and cancels before the first renewal. A campaign that looked strong at first payment can look considerably worse after a retention review.

Increase budgets in reviewable steps

State four things before raising a budget: the next increment, the cash limit, the question it answers, and the cohort review date. Raise the cap from $600 to $900 while holding the audience definition steady, and avoid moving creative, price and onboarding at the same time wherever that is practical.

Coordinate paid and organic work

Paid demand shows you which merchant questions generate interest, and those questions should feed the content program. Organic content answers the objections paid visitors run into. One comparison guide might serve as a landing destination, a sales follow-up asset and an organic search page: record all three uses, and do not assign the revenue to each of them.

COMPLETE THIS CHAPTER

Set a contribution-based acquisition ceiling and write one campaign brief covering audience, cash cap, expected funnel, full acquisition cost, conversion window and review date. Then test the whole route before launch, from the destination through to the recorded payment event. An approved advertisement is not a ready campaign.

MEASURE

Spend, clicks, installs, activation, payers, ad-only acquisition cost, fully loaded CAC, and subsequent retention.

9

CHAPTER

Find Merchants with a Reason to Listen

A store directory tells you who exists. That is all it tells you. A signal tells you why this merchant and why now, which is the only question your first sentence has to answer.

A signal on its own is still not an opportunity. It has to land on a merchant you can actually serve, carrying an offer that matters to them. A furious public review from a store you have no business selling to remains a dead end, and the fact that the frustration is visible does not make it yours to solve.

Signal to Advantage to Action

1. Name the observed signal.

2. Connect a relevant, supportable advantage.

3. Ask for one action: a demo or trial.

Match your wording to what you actually know. “Our latest scan suggests” is the right register for detection data, and the hedge

costs you nothing. Check the source and the date before sending, because a scan can be weeks stale, and a review that reads like a product complaint is sometimes a complaint about the sales process instead.

Case: a negative review that became a $10,000 ARR customer

The complaint was already public. A Shopify merchant had left a negative review of a competitor, and it pointed at friction around commercial questions. An account executive found it through Alert, wrote about that specific friction, and explained how the alternative handled it.

The merchant booked a demo. Then came the unglamorous part: setup and customer-data requirements that took several conversations across sales and onboarding, existing features reconfigured, new functionality built, workarounds agreed. Two months later the merchant signed, at a founder-reported $10,000 ARR, and later left a positive review of its own.

It is worth being precise about what did the work here, because the obvious lesson is the wrong one. The signal did not win this deal. It bought one relevant email. The offer and two months of delivery earned the customer, and $10,000 is revenue rather than contribution: it says nothing about what acquisition and service cost to produce it.

Case: competitor usage that became a $5,000 ARR customer

The second case contains no complaint at all. An account executive wrote to a Shopify merchant already running a competing app, named that app plainly, and set out the additional benefits and

better economics on offer. Nothing in the message suggested the merchant was unhappy, because nothing suggested they were.

A demo followed, and the deal closed in roughly a month at a founder-reported $5,000 ARR. Switching turned out to be cheap, which mattered more than anything in the email did. This is category fit plus a reason to move, not a dissatisfaction play, and the two call for completely different messages.

Turn signals into a work queue

Each account needs seven things recorded: merchant fit, the observed event, the source date, your confidence in it, the advantage that makes it relevant, the contact, and the next action. Keep the merchant's own words in a field of their own, separate from your interpretation of them.

PLANNING EXAMPLE

Research 30 accounts, accept 20 as fit, and verify an actionable signal for 12. Those 12 form the first outreach queue. The other eight need a different reason for contact or more research. Do not manufacture urgency to fill a quota.

Build a signal record

Keep the raw material alongside the record: the review excerpt and its link, the detection observation and its date, or whatever the buyer actually said.

Four complete signal plays

Competitor currently installed. Look for merchants whose existing category use is precisely what makes your advantage relevant, and confirm the competitor serves the job you intend to discuss. Lead with the usage you observed, name one specific improvement, and ask for a demo or a trial. What you must not do is imply that using a competitor is a problem in itself. They chose it on purpose.

Competitor uninstalled. Check how fresh the detection is and whether another product has appeared in its place. Mention the change with the uncertainty it deserves, then ask whether the underlying workflow still needs solving. An uninstall is ambiguous by nature: the merchant may have stopped the activity entirely, moved systems, or simply removed something redundant.

Negative competitor review. Read the complaint rather than targeting the star rating, and match the offer to the issue you find there, because a pricing complaint and an implementation failure call for opposite responses. Two things to avoid: attacking the competitor, and repeating the merchant's frustration back at them more aggressively than they expressed it themselves.

Complementary technology. Find a shared workflow your app can genuinely support. A relevant installed product establishes useful context, but any compatibility claim needs confirming before it goes in a message. The point is to explain why the combination helps this merchant, not to demonstrate that you can read the apps listed on their site.

Prioritize by fit and usefulness

A simple queue runs on three bands. Priority A has strong fit and a recent, relevant signal. Priority B has the fit but weaker evidence on timing. Priority C needs more verification, or lacks an obvious commercial route altogether.

These are internal work priorities and nothing more. Presenting them as “intent scores”, with an implied statistical probability of purchase attached, is a claim you have not earned unless you have validated that relationship.

Let likely deal value set research depth. A complex, high-value account can justify a detailed brief. A low-priced self-serve trial may need no more than a verified use case and a short message.

Read campaign results by signal

PLANNING EXAMPLE

Over comparable mature cohorts, 40 competitor-installed accounts produce four held demos and one payer. Twenty negative-review accounts produce four demos and two payers. Twenty uninstall accounts produce one demo and no payers.

FIGURE 15

Results by signal cohort

SIGNAL COHORT	ACCOUNTS	HELD DEMOS	PAYERS
Competitor installed	40	4	1
Negative review	20	4	2
Detected uninstall	20	1	0

The review cohort looks promising. Inspect fit, offer, seller effort and timing before concluding anything from twenty accounts. A reasonable response is to fund another review cohort while checking whether the uninstall message simply assumed too much.

Service burden and retained revenue deserve the same look. A signal that produces difficult but valuable customers may suit your model perfectly well. A signal that produces a stream of low-value exceptions will not.

Preserve relevance when using automation

Tools can collect evidence, organize accounts, and draft messages. What they cannot do is take responsibility for the claim. A generated sentence that invents a merchant's problem erases the entire benefit of good targeting, and it does so in the first line the merchant reads.

Use controlled fields for the signal and the approved advantage, and review a sample before releasing anything at volume. When an input is missing, route the account back to research or fall back to a less specific message that is still true.

Alert belongs in this workflow as the way accounts and events get found and organized. Teach the interpretation alongside the tool, so that whoever uses it understands why a particular account deserves attention rather than simply trusting that it does.

COMPLETE THIS CHAPTER

Create three signal plays for your app, writing what each signal means, what remains unknown, and the truthful opening sentence. Then review 10 messages, mark the fact and the assumption in each, and rewrite every unsupported claim.

MEASURE

Verified accounts, qualified replies, held demos, paid customers, and retained MRR by signal type.

10

CHAPTER

Run Burst Prospecting

Run the outbound yourself before you hand it to anyone else.

You need to hear which problems earn attention, which claims draw skepticism, and which merchants can actually buy. No outsourced sequence will give you that, and the understanding is what makes the sequence worth outsourcing later.

Burst prospecting is the pattern used here: four outreach rounds, with coordinated channels on each outreach day where that fits. Space the rounds one to seven days apart depending on fit, likely deal value, and whatever the context suggests.

A round is not a touch, and the distinction matters for capacity. An email plus a LinkedIn message is two channel actions inside one round, so four rounds at two actions each produces eight actions for a single account. Replies, declines and unsuitable channels bring that number down.

Use coordinated channels with judgment

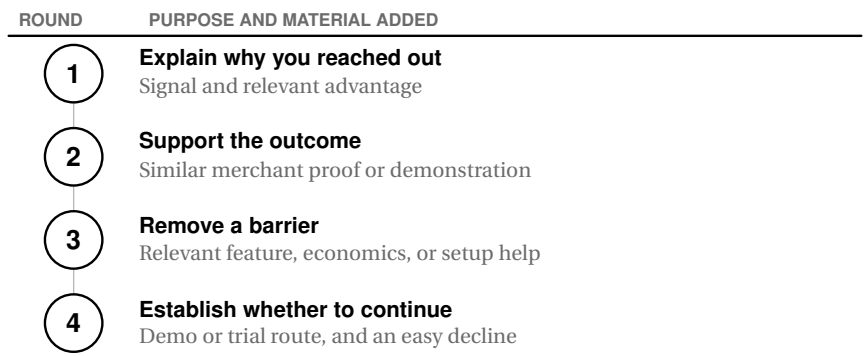
An opening round might use email, LinkedIn, a phone call and text, where each channel is appropriate and permitted. Treat that as a menu rather than an instruction to reach every person through every available route. Respect opt-outs, channel permissions, platform rules and stated preferences, and leave unavailable channels out of your capacity planning entirely.

Coordinated outreach should read as one relevant effort, not four copies of a pitch arriving with nothing added. The moment someone replies, stop the sequence and handle the conversation yourself.

The approach draws on multichannel sequencing traditions, including Justin Michael's work. The four-round implementation here is a practical operating choice rather than a finding.

Build the case across four rounds

FIGURE 16



Review-led opening:

> Hey *[Name]*, your public review mentioned difficulty getting clear answers about *[issue]*. We help *[merchant type]* handle *[relevant outcome]*, with *[specific difference]*. Open to a demo next week?

Plan accounts before counting actions

Appendix F2 has the filled 25-account, 200-action weekly plan. Cost the time before adopting it: at four minutes per action, 200 actions is over 13 hours, and that is before research, replies and demos.

Work a rolling queue. Finish the rounds already scheduled, then start new accounts. Record actions you completed rather than automation you intend to set up.

End the burst without burning the relationship

Four rounds is a default opening plan, not permission to ignore a decline. When the burst ends, move the silent accounts out of active outreach unless fit and context genuinely justify more. A new signal, a materially stronger offer, or a sensible interval can all justify coming back. What cannot is a new signal arriving for someone who opted out.

A complete competitor-installed sequence

What follows is illustrative. Replace the bracketed claims with evidence you can support, and keep each message to one question.

Round 1: relevant difference.

> Hey *[Name]*, our latest scan suggests *[Brand]* uses *[Competitor]*. We help *[merchant type]* achieve *[specific outcome]*, with *[material difference]*. Open to a demo next week?

Round 2: proof.

> Hey *[Name]*, a relevant example: *[similar merchant]* used *[capability]* to achieve *[verified result]*. Here is the workflow: *[link]*. Worth a demo for *[Brand]*?

If you cannot name the customer, use an approved anonymized example and be clear about its context. Never add a percentage just because the template reads better with one in it.

Round 3: remove the barrier.

> Hey *[Name]*, we also handle *[specific setup support]*, so your team would own *[limited merchant task]*. The commercial difference is *[accurate terms]*. Want to see it in a demo?

This round should address a barrier your category experience suggests is likely. If you already know the merchant's actual barrier, use that instead of guessing at a common one.

Round 4: close the initial attempt.

> Hey *[Name]*, last note for now on *[relevant outcome]*. If a switch is worth assessing, here is the demo link: *[link]*. Happy to leave it here if the timing is off.

The final message offers an action without manufacturing a dramatic breakup. Silence is an acceptable outcome. The job is to create an opportunity, not to force every account into a conversation it did not want.

Adapt the sequence to a review signal

The first round should reference the complaint directly and respectfully. The second can show how your process addresses it. The third explains the relevant commercial or onboarding support. The fourth closes the attempt.

Match the response to the complaint. Transparency concerns get clear terms. A missing workflow gets the available capability. A support complaint gets the support model you can actually deliver, in detail. “Better service” with no operational substance behind it is weak proof and reads as such.

Never promise that your product resolves every issue raised in a review unless you have confirmed that it does. One narrow truthful advantage is enough to start a useful conversation.

Use calls as conversations

A phone opening can be short. Identify yourself, reference the relevant event, and say why it prompted the call. Then ask permission to continue briefly, or offer the demo route.

The purpose is not to read the email aloud. Listen for three things: whether this person owns the issue, whether the timing is relevant, and what next action makes sense. Even a wrong contact will often route you correctly if the conversation is respectful.

Record call outcomes separately from calls placed. Ten calls that reached nobody and ten useful conversations are different results.

Use the prospect's response to change the route

A positive reply should stop the burst immediately. Offer specific times and a calendar link. If the buyer asked a question, answer it before sending another scheduling message.

Each response type needs something different. A "later" needs a date or a reason to return. A decline needs a reason where one is volunteered, then a clean close. A referral to a colleague needs a new contact record attached to the same account, so the company does not get counted twice.

Unsubscribe and suppression rules have to apply across the whole connected workflow. A founder who removes an email manually and keeps automated messages running elsewhere has not honored the preference; they have just made it harder to notice.

Evaluate the burst by account cohort

Track unique accounts, channel actions, positive replies, demos booked, demos held, trials, payers and MRR. When evaluating the sequence as a whole, the account is the denominator.

If 25 accounts receive 150 actions and three book, account-to-booking conversion is 12%. It is not 2% because three divided by 150 happens to be 2%. That figure measures bookings per action,

which answers a different question. Record both if they are useful, and never compare them as though they were the same measure.

Improve one part of the sequence

When a burst underperforms, there are six places to look: the list, the signal, the advantage, the action, the delivery and the channel mix. Read the actual messages and the actual replies before increasing volume on any of them.

Change one meaningful element where you can. A stronger proof asset may lift Round 2. A narrower segment may lift the whole sequence. More follow-up only helps if the underlying reason for contact was worth something to begin with.

COMPLETE THIS CHAPTER

Write four rounds for one signal and segment. Select appropriate channels, estimate delivery time, run a small batch, and record cohort results before deciding the next change.

MEASURE

New unique accounts, completed rounds, actions by channel, positive replies, held demos, and paid customers by cohort.

11

CHAPTER

Build Content That Helps Merchants Buy

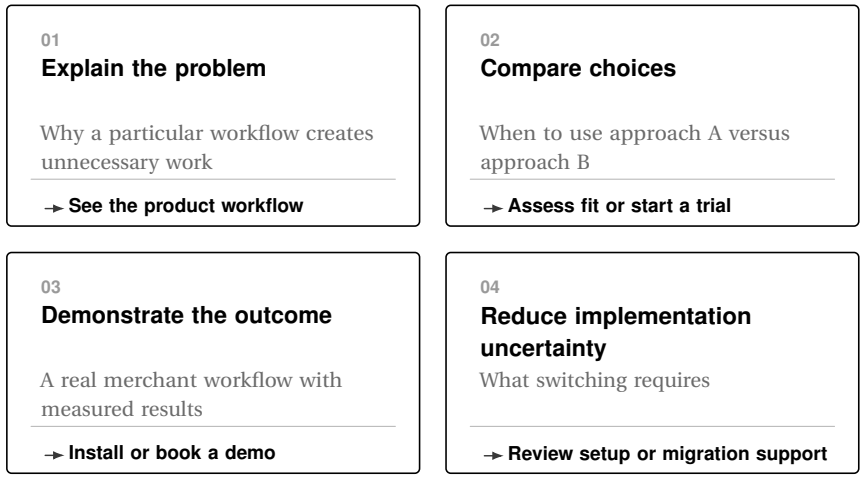
Content should answer the questions a qualified merchant asks before installing, buying, or switching. Not the questions you find interesting.

A founder can publish consistently for six months and generate almost no demand, because the topics attracted other founders, general software readers, and merchants outside the segment entirely.

Start with the questions already showing up in prospect calls, support tickets, search research and competitor comparisons. Those are the ones already attached to a buying process.

Four content jobs

FIGURE 17



Comparisons must be accurate and dated, and they should say who each approach fits. Invented competitor weaknesses cost you trust and deliver poor-fit trials, which is an expensive combination.

Build a buying question library

Every useful sales conversation can improve the content program. Record the questions that keep recurring before a decision: what the app does, how it compares, what setup involves, what it costs, and how success gets measured.

Group similar questions instead of turning each phrasing into its own article. One complete guide to assessing switching effort will answer several related concerns better than five thin posts that each answer a fraction.

Prioritize by buyer relevance and commercial usefulness rather than volume. A topic with modest search demand can be extremely valuable if it resolves an objection that appears in half your sales conversations. A popular topic outside your audience produces traffic and nothing else.

Give each asset one primary reader

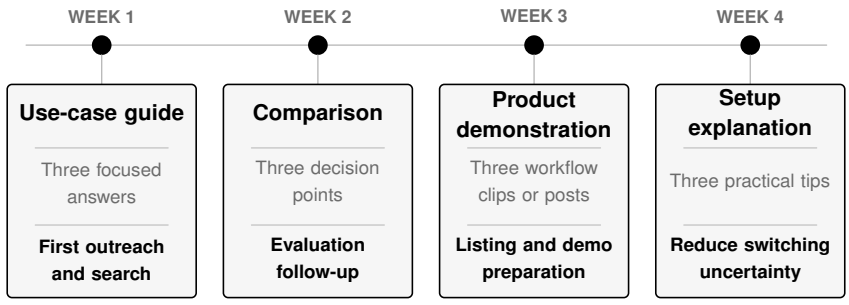
Name the merchant role and the evaluation stage before writing. Answer one primary buying question in the reader's language, with their workflow and their tradeoffs, then give them a relevant product route. A guide can serve several people without attempting to explain the whole business at once.

Build four assets before a large calendar

PLANNING EXAMPLE

Your first month contains a use-case guide, a comparison, a product demonstration, and a setup explanation. Between them they cover the problem, the alternatives, the experience, and the implementation.

FIGURE 18



Use original data carefully

Alert's observations can support genuinely useful content about technology usage, category changes and relevant merchant segments. The value is in explaining the scope of the data and what it implies for a decision.

State the observation window, the population, the detection method at an appropriate level, and the limits that matter. Three of those limits are worth naming every time: a detected install is not necessarily a new paying customer, a detected uninstall is not a verified cancellation, and a relationship between two apps does not establish that one caused the other.

A report should help the reader decide whom to research or what market question to ask next, not convert measurement uncertainty into a headline the evidence cannot carry.

Give distribution its own activity target

Publishing is one activity. Getting the asset in front of the people it was written for is a separate one, and it is the one that gets skipped.

A distribution checklist might cover company social posts, relevant sales follow-up, an appropriate customer update, and partner use where the relationship supports it. What it should not become is the same link dropped into communities with no interest in the topic.

Track distribution completion separately from the impressions it earns. Whether the asset went out through all of the planned routes is entirely within your control. How many people then see it is not.

Build a useful company voice

A founder does not have to appear in every asset. Product demonstrations, merchant workflows, practical comparisons and transparent analysis build the company's credibility perfectly well without a personality attached, and consistency matters more than sounding like a thought leader: specific observations, supported claims, one clear idea per short post.

Use first-person experience where it is real and you are permitted to share it. Do not invent a founder anecdote to make a generic point feel personal; readers who buy software for a living can tell.

Diagnose content in stages

Diagnose in stages. If the asset is not being seen, the problem is distribution and audience access. If it is seen but not read or used, the problem is relevance or presentation. If it is useful but produces no direct conversions, consider whether it plays an assisted role before removing it.

PLANNING EXAMPLE

A guide receives 150 visits, contributes to five sales follow-ups, and gets mentioned by two new customers. You do not know it caused those sales. You do know it is being used inside the buying process, which is enough to justify maintaining it while you keep measuring direct demand.

Maintain the library

Give pricing pages, comparisons, setup guides and data reports a review date. An old asset will happily keep generating traffic while telling buyers things that stopped being true a year ago.

Update the material when the product or the evidence changes, and record major revisions so acquisition comparisons stay interpretable. A content library becomes a business asset by staying accurate and useful. Size has nothing to do with it.

COMPLETE THIS CHAPTER

List 10 real buying questions and choose four for next month. Assign each a reader, asset, distribution plan, sales use, conversion route, and review date.

MEASURE

Assets shipped, distribution completed, relevant visits, assisted conversations, attributed installs, and paying customers where measurable.

12

CHAPTER

Run a Demo That Earns a Decision

A demo should help the merchant decide whether the product fits their business. A full feature tour usually achieves the opposite, adding information while making the decision harder.

Let their situation choose the story.

The meeting structure here comes from Sandler: establish rapport, agree the purpose and process, understand the problem, discuss fit, establish an outcome. Sandler's up-front contract is a mutual agreement about how the conversation will run, not a device for forcing a purchase. [W4]

Start with a clear agreement

A practical opening is: "I'd like to understand how you're handling this today, show the parts that matter, and leave time for pricing and next steps. If it isn't a fit, we can say so. Anything else you want to cover?" That gives the buyer a voice in the meeting and makes its ending considerably less awkward for both of you.

Ask about the current workflow, the problem, its impact, the reason to act now, and who else has to be involved in a decision. Learn enough to tailor what you show. What you want to avoid is turning the call into a questionnaire that postpones every piece of value to the back half.

The Three-View Demo

FIGURE 19



Some apps have no shopper-facing interface at all. Show the relevant external output or workflow instead, rather than inventing a view to satisfy the framework.

Move between the views deliberately, because each does a different job. The deck explains why a capability matters. The live example makes it concrete. The back end answers the question the operator is actually asking, which is whether their team can run this.

Make an offer with a real decision date

A time-limited incentive can help a merchant act. State the actual offer, the eligibility, the deadline, and the normal price. A $2,000 annual plan offered at $1,000 illustrates a pricing structure; it is not a recommendation that every app discount by half.

The deadline has to correspond to a real commercial offer or implementation commitment. Announcing a final deadline every week for a month teaches the buyer that your dates mean nothing, and that lesson is hard to unteach.

Where the fit is established, ask for the business. Where the merchant needs another step first, agree on what has to happen, who owns it, and when you speak next.

Prepare for the merchant rather than for the deck

Before the demo, look at the account's site, the signal that prompted contact, and whatever the buyer has already told you. Choose the workflow and the proof from that, then prepare questions about what you still do not know. The aim is enough context to avoid a generic presentation, and enough humility to let the buyer correct whatever you got wrong.

Sellers lose hours trying to estimate every number about a business from the outside. Ask the merchant for the inputs that matter when the conversation reaches the financial case. They have them, and they would rather give you the real ones.

Demonstrate the outcome in context

Start with a relevant example, then show how it gets produced. The merchant sees the result first, which is what makes the back-end settings worth their attention.

For a shopper-facing app, show the customer experience in a real or clearly labeled demonstration store. For an internal one, show the operator's task and its output. The deck frames the story and carries the commercial context; it is not there for narrating buttons.

Stop occasionally and confirm you are still on target. “This is the workflow you described. Is anything important missing?” costs ten seconds and can save the remaining forty minutes from running on a wrong assumption.

Build an honest financial estimate

Use inputs the merchant can inspect, and state the period, the baseline, the expected change, and the relevant cost. Keep estimates visibly separate from established results.

PLANNING EXAMPLE

The merchant handles 500 monthly tasks and estimates two minutes saved on each. That is 1,000 minutes, or roughly 16.7 hours. At their own $30 hourly labor assumption, the capacity value is about $500 a month. Against a $150 fee, the modeled difference is $350 before any other costs.

That is an estimate about capacity, and capacity is not cash. Nothing here guarantees that paid labor expense actually falls. Where the task volume or the time saving is uncertain, show a range instead of a point.

At one minute saved per task the same model produces about 8.3 hours and $250 of capacity value. The lower scenario is still useful to the merchant, because it gives them a floor to judge the fee against.

Handle a feature gap directly

If the app cannot do something they need, say so plainly. Then find out whether an existing configuration or an acceptable workaround meets the business need. Where new work is genuinely under

consideration, separate that decision from the demonstration you are currently in.

Do not make an engineering commitment because a prospect feels close to signing. Give the team enough detail to estimate the effort and decide whether the capability belongs in the product at all.

A confident limitation preserves trust. A vague promise can close the conversation and create a delivery problem that resurfaces later as churn and a public review.

Review the demo as an operating asset

Where recordings are available and appropriate, use them on yourself. Five questions are worth asking each time. Did the discovery questions change what you showed? Did you make a claim you cannot support? Did the merchant see the workflow that matters to them? Was the pricing clear? Was the next step explicit?

Review the losses and the wins. A won deal can contain habits that will fail you elsewhere, and a lost deal is sometimes well-qualified work that kept an unsuitable customer out of the business.

COMPLETE THIS CHAPTER

Review three demos for the merchant problem, workflow, financial case, and commitment. Use the gaps to revise your agenda, product views, and economics worksheet.

MEASURE

Held demos, qualified next steps, paid conversion, time to decision, and recurring objections.

13

CHAPTER

Keep Deals Moving or Close Them Out

A good demo can still produce a deal with nowhere to go.

The merchant liked the product. The founder sent a deck. Two weeks later neither side could tell you what is supposed to happen next.

Close that gap while the meeting is still running. Book the next meeting from this one whenever another conversation is genuinely needed, and give it a purpose: reviewing an internal decision, resolving a setup question, agreeing a launch.

Send a decision recap

A useful follow-up carries five things: the merchant's priorities, the relevant product outcome, the commercial offer, the resources, and the agreed next actions. Share the recording where it is available and appropriate for the people on it.

Recap:

> Hey *[Name]*, thanks for the conversation.
>
> Your priorities are *[priority 1]* and *[priority 2]*. We showed how *[capability]* supports *[outcome]*, with *[setup support]*.
>
> Resources: *[deck]*, *[recording]*, and *[financial estimate]*.
>
> The offer is *[terms]*, available through *[real date]*. You're confirming *[merchant action]*, and we're handling *[vendor action]*. We're scheduled for *[date]* to *[decision purpose]*.

Keep it proportional. Nobody should have to read an essay to find the action they agreed to.

Distinguish follow-up from prospecting

Follow up on the decision you share, not with a generic pitch. Two or three touches around the agreed timeline, each delivering something: requested evidence, a resolved objection, a status confirmation. When a merchant names a later decision date, respect it.

Use clear pipeline stages

FIGURE 20

Pipeline stages and the evidence each requires

STAGE	EVIDENCE REQUIRED TO BE IN THIS STAGE
Engaged	A relevant reply or conversation
Demo scheduled	A meeting accepted by the merchant
Evaluating	Demo held and evaluation next step agreed
Committed / setup	Commercial commitment established and delivery underway
Paid	Payment or a paying arrangement has begun
Closed out	Declined, disqualified, or no active next step

Adapt the stages to your own buying process. One thing should survive the adaptation: an install, a free trial, a signed agreement and a paying customer must stay distinguishable from one another, however cheerfully your CRM labels one of them a win.

Decide what advancement means

Move a deal only on buyer evidence: a meeting attended, a stakeholder introduced, information supplied, terms approved, setup agreed. Sending another resource is seller activity. It feels like progress and moves nothing.

Templates for common situations

Use the full recap above straight after the demo. For the messages that follow:

Promised action is late:

> Hey *[Name]*, checking on *[agreed action]* so we can keep *[planned next step]* on track. Can you confirm the status?

Internal discussion needed:

> Hey *[Name]*, here is a short summary you can share: *[outcome]*, *[commercial terms]*, and *[setup commitment]*. Happy to demonstrate the relevant workflow to the team. Open to booking that next step?

Opportunity has gone quiet:

> Hey *[Name]*, should I close this out for now? Happy to revisit when the timing makes sense.

Give a genuine later opportunity a place

A merchant can be a genuinely good fit and still unable to act until a contract ends, a hire starts, or a trading peak passes. That is a reason to create a dated return, not to leave the deal sitting in the current forecast.

Record the condition, the date, and who agreed to revisit. Keep that queue separate from active evaluations, or the pipeline will look healthy while the current month holds almost no buyer commitment.

When you do return, reference what was actually agreed. If circumstances have moved, qualify again from the start. A positive

conversation from three months ago is not a current purchase decision, however warm it felt.

Keep offers and deadlines coherent

A time-limited offer should read the same way in the recap and in the commercial documents. When terms change, record who approved the change and what the buyer was told. When a deadline passes, either apply the stated terms or explicitly approve a new offer for a reason you could defend. Quietly extending every deadline teaches the buyer that your dates are decorative, and they will price accordingly.

One thing not to do: send an expiry message while an unresolved question on your side is what is blocking the merchant. Fix your own commitment first. Urgency should organize a real decision, not paper over an evaluation that was never finished.

Forecast with evidence and ranges

Be careful with probability-weighted pipeline built on arbitrary percentages. A deal sitting in "evaluation" does not have a 50% chance of closing because the CRM assigned it one.

Start somewhere simpler: committed actions, likely timing, unresolved issues, and the revenue if it lands. Bring historical conversion in once you have enough consistent stage data for it to mean something.

PLANNING EXAMPLE

Three deals could add $1,500 MRR between them. One has approved terms and a scheduled launch. One is waiting on another stakeholder. One has stopped responding entirely. Reporting all three as equally

likely conceals exactly the uncertainty you need to see, so hold a scenario range and know which merchant action would firm it up.

Close the loop after a loss

Record the main reason they declined. Fit, price, timing, incumbent preference, migration effort, missing capability. Wherever you can, keep the buyer's own words rather than your summary of them.

Each loss points somewhere. A loss to a partner recommendation is information about partner strategy. A loss over setup effort belongs with the switching offer. A loss because the merchant never had the problem belongs with targeting.

Let the process improve from losses without turning every one into a product commitment. The signal is in the pattern across relevant accounts, not the last conversation you had.

COMPLETE THIS CHAPTER

Create the recap template. Audit active deals for a current buying process, next action, owner, date, and evidence needed to advance. Close out those with no active process.

MEASURE

Deals with dated next steps, stage age, held follow-up meetings, win rate, and reasons for closure.

14

CHAPTER

Make Events Worth the Cost

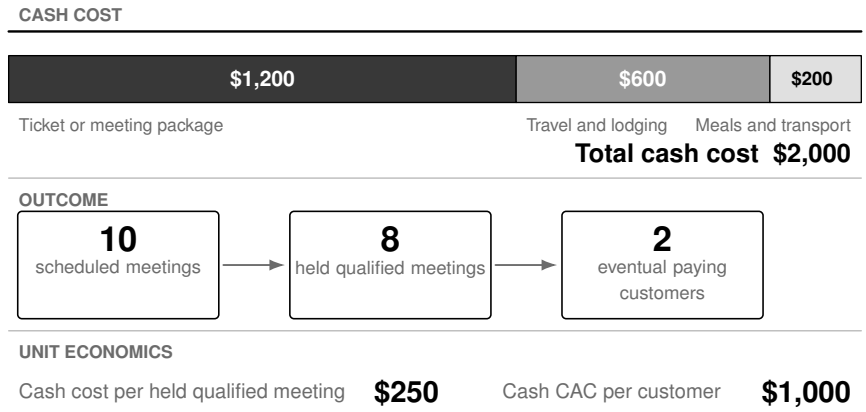
An event can put you in a room with merchants you would otherwise struggle to reach. It can also consume a large share of the acquisition budget and produce nothing you could call a qualified conversation.

The lean version is usually right at this stage: attend locally where you can, send the founder or one salesperson, and skip the booth unless you can name the economic reason for buying it. Meetups, happy hours, niche gatherings and major conferences are different formats of one question. Does this event create useful access to our target merchants?

Write the event plan before booking

FIGURE 21

The economics of one event



Run the same cash CAC against two different businesses and watch what happens. At $100 MRR and $80 of monthly contribution, $1,000 takes 12.5 months to recover before extra costs and churn. At $500 MRR and $400 of contribution, it takes 2.5. The same conference is sensible for one app and unaffordable for the other.

Choose the event by the people you need

Assess the relevant merchant roles, the verticals, the format, the access and the meeting options. A niche event or a local meetup frequently outperforms a broad, expensive show on every one of those.

Give each meeting a primary objective before you walk in: customers, partners, or learning. Counting every conversation as a sales opportunity is how event reports become fiction.

Use a pre-event pipeline

Treat the weeks before an event as a small campaign. Build a fit-checked list from the attendee information available, your existing prospects, the sponsors, and whoever has announced publicly that they are going.

Then give them a specific reason to meet. “We’ll both be there” is context. “We can show you the workflow you asked about” is a reason.

Run it through your existing CRM so the event does not quietly create a second set of duplicate accounts. Track meetings requested, accepted, held, and the next step. A scheduled-meeting count only becomes meaningful once the meetings happen with the right people in them.

A lean event preparation plan

PLANNING EXAMPLE

Four weeks out from a local event, identify 30 relevant accounts and aim for six accepted meetings.

A week out, confirm times and locations, and prepare the product view and short explanation you will actually use.

Afterwards, record the merchant’s context, send what you promised, and book the agreed next step while they still remember the conversation.

Evaluate meeting packages before buying

Four questions before you buy a meeting package. Who qualifies as a participant? Do the roles match your buyer? How are the meetings selected? And what happens when someone does not show? Ten meetings is attractive if eight are relevant and held, and poor value if most participants have no buying responsibility.

Model the whole cost. Tickets, travel, lodging, meals, preparation, attendance and follow-up all consume something. A founder attending "for free" is still spending time that could have gone into outbound or onboarding.

Paid merchant introductions are another access route

Services exist that recruit merchants for compensated conversations, with vendors paying while participants receive gift cards or attendance incentives. Running a direct program yourself is another route to the same thing.

Be clear about what the money buys. It buys participation. It does not buy purchase intent, a customer, or a favorable review. Evaluate the audience, the incentive, the qualification and the eventual economics using the process below.

Qualify incentive-based meetings

Before paying, specify the merchant type, the operating scale where it matters, the current workflow, the participant role, and the exclusions. Then ask the service two things: how those criteria get verified, and what you receive before the call.

A compensated participant can be an excellent prospect. They can also be there mainly for the incentive. The demo is what distinguishes the two, by establishing whether a genuine problem and a buying path exist.

Be explicit that compensation is for the agreed participation under the program's disclosed terms. It must never depend on buying, praising the product, or changing a review, and an attendance program and a review request should never travel together.

Model the economics of a meeting program

PLANNING EXAMPLE

A program costs $3,000 for 10 scheduled conversations. Eight are held, six meet your qualification criteria, three enter evaluation, and one becomes a $500 MRR customer.

- Cost per scheduled conversation: $300.
- Cost per held conversation: $375.
- Cost per qualified conversation: $500.
- Program-only CAC: $3,000.

At $400 of monthly contribution, simple program-only payback is 7.5 months before any additional sales and setup cost. Acquire two customers instead of one and program-only CAC halves to $1,500, with payback at 3.75 months.

That gap between one customer and two is enormous. Wait for the mature cohort before renewing a substantial package.

Give every paid introduction a disposition

Give every introduction a disposition with a reason attached: held, qualified, evaluating, declined, unsuitable. The pattern tells you where the problem is. Unqualified attendance points at recruitment. Qualified non-buyers point at your offer or your demo. Open evaluations point at nothing yet, and are not wins.

Feed the quality problems back to the service, and require visibility into the audience and the commitments they have made to you.

Decide what earns a return visit

Compare the customer value and the learning against the full cost. A promising first test can justify another bounded commitment. A weak audience fit is a reason to stop, including when the event was genuinely enjoyable, which is the case where founders most often fail to stop.

COMPLETE THIS CHAPTER

Build one event business case with cash cost, time cost, target meetings, likely customer value and a follow-up owner. Compare it against one conference package and one paid introduction program using identical cost categories, then choose on the access you need and the economics you can carry, not on which logo looks best in the calendar.

MEASURE

Scheduled and held qualified meetings, opportunities, payers, MRR, and contribution payback by event cohort.

15

CHAPTER

Build Partnerships After You Have Traction

Partnerships can move a merchant's decision because the recommendation arrives from somebody they already trust.

They also take real work. A referral page offering a percentage of revenue almost never produces an active referral engine on its own.

For an early-stage app, partnerships are usually a later investment, after the listing, website, customer proof and basic acquisition process are credible. You can start building relationships well before that, provided immediate customer acquisition stays focused somewhere else.

Give value without assuming repayment

Begin by being useful. Send relevant introductions. Spend time together at events. Help with a joint asset or a merchant problem that is actually theirs. A working relationship beats a commission link exchanged between two strangers, every time.

Buying legitimate services from an agency can build familiarity too, but buy the service because you need it. Hiring an agency does not entitle you to referrals, and treating it as though it does poisons the relationship quickly.

You may send 20 referrals and receive two. That happens even in friendly relationships. Track the commercial reality rather than the warmth.

Run a small partner test

> **PLANNING EXAMPLE**
>
> Pick three potential partners for a 90-day pilot. For each one, agree the ideal introduction, an owner on both sides, how consent-based introductions will work, and a review date.

You send six qualified introductions across the group. You receive four, three hold a demo, and one becomes a $200 MRR customer. Record the actual effort and compensation before deciding anything about continuing.

What those numbers do tell you is which relationships are active and which are dormant agreements, which is what the pilot was for.

Make the partner offer useful to the partner's customer

A partner has to put its own reputation behind your product. Commission can support that, but it cannot be the reason, because the recommendation still has to help the merchant receiving it.

For an agency, that usually means a product that implements easily, support that answers, and a result their client recognizes. For an app, it means a complementary workflow that improves the experience both of you touch.

Write the joint use case before you negotiate the referral percentage. If you cannot explain why a merchant should use both products, or why anyone should accept the recommendation, the partnership has no commercial foundation and the percentage will not supply one.

Distinguish three partner motions

Three motions hide under one word. A referral partner introduces suitable merchants. An implementation partner helps merchants adopt the product. A distribution partner provides access to an audience or a buying route. One company can play several of those roles, but the work and the economics differ sharply, so name the role you are actually testing.

Each needs something different. An agency implementing the app needs training and a dependable escalation route. An app making occasional introductions needs a concise fit guide. A newsletter placement needs audience information and tracking. Calling all three "a partnership" and giving them the same agreement and cadence is how all three underperform.

Select a manageable partner list

Start with a few relationships that show shared merchant fit, relevant access, a real product relationship, an owner, and some willingness to act. Before traction, build them without depending on them for acquisition. Three active relationships will teach you more than thirty dormant logos on a slide.

Build the first joint action

Agree one action, one owner, one date: a qualified introduction, a team demonstration, a joint guide, an implementation walkthrough.

Then give them a short fit guide covering whom you help, the problem, the exclusions, the buying route, and who responds. That makes a suitable merchant recognizable, which a general agreement to exchange leads never does.

Run the referral workflow

Use permission-based introductions where that is appropriate. A good introduction tells both sides why the conversation makes sense. The receiving team acknowledges it, schedules the next step, and keeps the referring partner informed at a level that respects everyone's time.

Settle four operating questions before a deal appears: how compensation is calculated, when it is earned, what happens after a cancellation, and how duplicate referrals are handled. The detailed terms belong in an agreement; the answers need to exist beforehand.

Never promise a commission casually on the assumption that accounting can reconstruct it later. Unclear payments will damage a relationship that took months to build, over an amount that was never worth the argument.

Measure reciprocity without making it transactional

Count the value that fits the partner's role, not just reciprocal introductions. Planning example: Partner A receives eight introductions, sends none back, and successfully implements three shared customers. Partner B sends four introductions, none of which fit. A is probably the more valuable relationship, despite producing no inbound referrals at all.

Record introductions, implementation help, joint assets, customers, revenue and effort. Then judge each partner against the role you agreed for the pilot, not against each other.

Know when to pause a partner motion

Four reasons to pause: no clear shared use case, no owner for the next action, repeated poor-fit introductions, or effort that keeps displacing more productive work.

A pause is not an ending. You can keep occasional contact while taking an inactive commercial initiative out of the operating plan.

What you should stop doing is reporting a partner as active because a contract exists. Define activity as introductions, customer work, or an agreed project completed inside a stated period.

Use partner losses as research

When a merchant picks a competitor on the recommendation of a trusted app or agency, ask what made that recommendation persuasive, if they are willing to say. It is usually one of four things: implementation familiarity, commercial terms, confidence, or simply access at the right moment.

The answer might point you toward better service, a specific relationship worth building, or an offer that is easier to recommend. It

rarely points toward paying more commission, whatever the losing instinct suggests.

COMPLETE THIS CHAPTER

Define one partner profile and one referral workflow. Choose three potential partners, define the role for each, and propose one useful joint action. Review after a bounded period and keep the relationships producing merchant value and a sustainable commercial contribution.

MEASURE

Qualified introductions given and received, held conversations, paying customers, retained MRR, commissions, and time invested.

16

CHAPTER

Keep the Customers You Worked to Win

Acquisition can look healthy while MRR barely moves.

Add $1,000 in new MRR, lose $900 from existing customers, and the business grew by $100 before anything else happened. The acquisition team hit its target. The company stayed roughly where it was.

That is why retention belongs in the GTM operating system. It decides how much acquisition you need and how much you can afford to spend getting it.

Separate cancellation reasons

A merchant who never finished setup is a different problem from one who used the app successfully for a year and then closed the store.

Keep the reason list short: setup failure, missing capability, insufficient value, technical issue, price, business closure, unknown. Attach a note with the specific circumstances.

Do not fill the unknowns with whichever reason you find most comfortable. A cancellation with no feedback attached is unknown, and it stays unknown until evidence arrives.

Review revenue by cohort

PLANNING EXAMPLE

Start the month at $5,000 MRR. Add $1,500 from new customers, add $200 from existing customers paying more, lose $100 from downgrades, and lose $600 from cancellations.

Ending MRR = $5,000 + $1,500 + $200 − $100 − $600 = $6,000.

Net new MRR is $1,000. Gross new customer MRR is $1,500. They answer different questions.

Retention starts with the customer you select

Keep the original promise and merchant profile attached to the account, then compare both against what was delivered when someone cancels. Demanding customers can be extremely valuable where the work fits your positioning and economics. Poor-fit expectations produce losses you could have avoided at the qualification stage.

Track the first renewal separately

The first renewal tells you something the first payment cannot. A merchant will often pay during setup and then stop when the app never becomes part of how they work.

For monthly subscriptions, track what share of each new cohort is still paying after its first renewal opportunity and the ones after that. For annual subscriptions, early usage and customer outcomes give you interim evidence, but the renewal itself stays unobserved until the contract period ends.

So do not compare annual and monthly retention as though the observation opportunities were equivalent. A customer nine months into an annual contract can be commercially active and operationally disengaged at the same time.

Build an account outcome record

Keep the merchant's desired result, the activation date, current product use where you can observe it, support issues, and the next review. How much detail is appropriate depends on what the account is worth.

A self-serve app can automate much of that record; an assisted account may need a named owner and an actual outcome conversation. Either way the purpose is noticing when the promise is not becoming real.

Do not report activity as value. Logging in is evidence of access. It says nothing about whether the merchant completed the workflow that was supposed to help them.

Inspect churn with the right denominator

PLANNING EXAMPLE

You start a month with 50 paying customers and $5,000 MRR. Five cancel, taking $800 of MRR with them. Customer churn is 5 ÷ 50 = 10%. Cancellation MRR churn is $800 ÷ $5,000 = 16%.

The revenue loss is larger than the customer loss because the accounts that left paid more than average, and that should change your diagnosis. A customer-count report on its own would have hidden the concentration entirely.

Now suppose the remaining customers add $300 of MRR by moving up plans. Net revenue retention for the starting group, excluding new customers, is ($5,000 − $800 + $300) ÷ $5,000 = 90%, assuming nothing else moved. State those inclusions whenever you use the metric.

Use cohort retention before broad averages

Compare acquisition sources and segments with similar time at risk. A recent cohort has had fewer opportunities to cancel than an older one, and comparing them directly flatters the newer source.

Customers acquired through compensated meetings may behave quite differently from those who arrived through category search: different motivation, different setup needs, different buying process. Compare their activation and retention before you decide which source produces better customers.

Keep the sample size visible while you do it. Two retained accounts are experience, not a retention rate you can forecast a budget against.

Recover the relationship before requesting advocacy

Fix the problem before you ask for anything. Support must never depend on a favorable review or a referral. Chapter 5 has the recovery process, the neutral message and the platform rules; Appendix D5 has the record.

Feed recurring setup or pricing friction back into the buying materials, where it can prevent the next occurrence. And treat a positive review as useful evidence rather than proof of long-term retention, because the two often disagree.

Make referrals specific and easy

Ask after an observed result or a conversation that went well. Explain the kind of merchant you can help, and offer a simple introduction route rather than a vague request.

Where you run a customer referral incentive, disclose the terms and keep it entirely separate from review requests. A referral reward is a

commercial program. An honest product review is something else, and bundling them damages both.

Measure introductions, qualified conversations, new payers and retained revenue. A customer forwarding your link has not created a qualified opportunity, however good it feels to see it happen.

Decide whether acquisition or retention is the immediate constraint

Suppose you add $1,500 of gross new MRR a month and lose $1,200. Doubling outbound will increase net growth and also feed a leaky system. Work out first whether the losses come from fit, activation, product value, or ordinary business turnover, because only the first three are yours to fix.

Some loss is unavoidable. The real question is whether an identifiable retention improvement is a better use of your time and cash than another acquisition campaign.

COMPLETE THIS CHAPTER

Define activation ownership, cancellation reasons and one recurring customer outcome review. Then compare your last five cancellations against five retained customers from similar segments, pick one preventable cause, and track the next comparable cohort.

MEASURE

New and lost MRR, customer churn, revenue churn, activation, support burden, and retention by acquisition cohort.

17

CHAPTER

Run Your GTM Every Week

A goal describes an outcome. A calendar determines whether the work happens.

The weekly operating system joins the two through Activities to Objectives to Results, or AOR. Activities are the work you do. Objectives are the intermediate outcomes you want that work to produce. Results are the business outcomes: activated customers, payment, and retained revenue.

Each level gives you a different degree of control. You control whether an account is researched and contacted. You influence whether a merchant books a demo. You measure whether the customer pays and stays. Confuse the levels and you get either vanity reporting or activity targets with no economic purpose behind them.

Start with the revenue goal

PLANNING EXAMPLE: A SELF-SERVE MONTH

You start at $4,000 MRR and want to finish at $5,000. You expect $100 of cancellations and downgrades, and you assume no increase from existing customers.

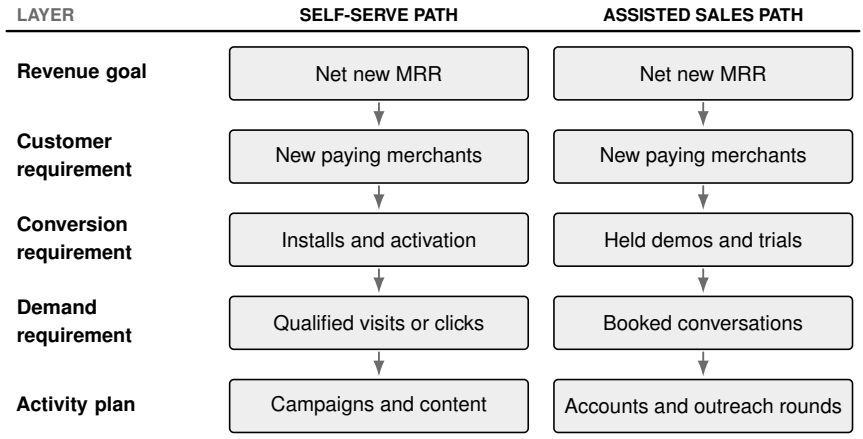
Required new customer MRR = $1,000 desired net growth + $100 expected loss = $1,100.

At $100 MRR per new customer, that is 11 paying customers. At 20% install-to-paid conversion, 55 installs. At 10% listing-visit-to-install conversion, 550 comparable visits.

Those assumptions eventually have to come from your own cohorts. And 55 installs acquired late in the month will not produce all 11 payers by month-end, so either move the acquisition schedule earlier or model the conversion lag deliberately.

The Revenue-to-Activity Ladder

FIGURE 22



Do not count the same customer in both paths. Where a merchant sees your content and later responds to outbound, keep both interactions on the record and assign the customer once in your acquisition totals.

A sales-led planning example

Suppose you want $2,000 of net new MRR, expect $500 in losses, and earn $500 MRR per new account. You need $2,500 of gross new MRR, which is five new payers.

Now work backwards. At 20% of held demos converting, you need 25 held demos. At an 80% attendance assumption, round up to 32 bookings to support 25 held. At 10% of qualified accounts booking, that means 320 accounts entering outreach.

Every assumption in that chain is yours to replace with observed cohort data as it arrives. Round operational requirements upward and plan a buffer you can afford.

Then check it against the week. Across a four-week month, 320 accounts means 80 new accounts weekly. At four rounds and two actions per round, steady-state workload peaks at 640 channel actions a week before early exits. Most founders do not have that capacity, which means the model just found a problem before the plan reached the calendar.

Five honest options follow: lower the near-term goal, improve conversion with evidence, increase customer value, add a second acquisition source, or add supported capacity. What is not an option is requiring the founder to work hours that do not exist.

Build the activity tracker

Appendix B is the master weekly sheet: targets, actuals, previous week, activities, demand, adoption, costs and MRR movements. Appendix A defines the metrics.

Carry connections, active deals, total payers and ending MRR as balances. Keep outbound prospect forms separate from inbound inquiries. Resist building a second, narrower tracker alongside it; two trackers means two versions of the truth.

Hold a short weekly review

Appendix E has the 30-minute agenda. Find the constraint using the method below, record one change, and give it an owner and a review date.

The MRR Mission

The MRR Mission puts the revenue target, the funnel requirements, the weekly activities and the resource limits on one page. A plan is ready when the work fits your calendar and your budget, and not before.

Use a range before you have a baseline

Model lower, middle and higher conversion assumptions before committing cash. For 10 payers, install-to-paid rates of 10%, 20% and 25% require 100, 50 and 40 installs, which at $20 per install means media costs of $2,000, $1,000 and $800. Replace the assumptions with cohort evidence as it matures.

Use definitions that survive a team handoff

Write down what counts as a new account, a sent email, a held demo, an install, a trial, an activated merchant and a payer. Then decide how you handle duplicates, reinstalls and multiple stores.

One buying company with three stores can generate three installs and one commercial account. Both counts are useful. Neither should quietly substitute for the other in a report.

If your CRM marks “closed-won” when a free trial begins, keep that workflow if the team needs it, but add a separate paying-customer event alongside it. An executive revenue report should never treat free evaluation as MRR.

Track activity without rewarding noise

A high action count can be real evidence of execution. It can also reward irrelevant work, if the target says nothing about quality.

Three gates prevent that: a fit gate for accounts, truthful signal requirements for messages, and clear completion definitions for content. A LinkedIn comment should be a relevant contribution, not a stock phrase posted to satisfy a number.

The other failure is a tracker so detailed that recording the work becomes a second job. Keep the source data in the systems where the work happens, and summarize only the few metrics the weekly decision actually needs.

Manage capacity with an action budget

Put a rough time estimate against research, direct messages, calls, demos, follow-up, content and analysis. Then measure a sample of real work and improve the estimate, because the first version is always optimistic.

PLANNING EXAMPLE

Twenty-five accounts need two hours of research, six of outreach, two of reply handling, and four of demos and follow-up. That is 14 hours before any content or product work. If the founder has 10 hours, the plan has to change rather than the arithmetic.

Automation reduces some of that effort. It does not make the work free. Data review, reply handling and quality control all remain, and the customers you acquire create onboarding and support demand of their own.

Use a constraint review

Start at the revenue gap and look at the step immediately before it. Too few payers often means too few evaluations. Too few evaluations often means incomplete activity or poor targeting. Keep walking back one step at a time.

Confirm the evidence, pick one obstruction, and set a test. “Reduce abandonment at required setup” is actionable. “Improve marketing” is not.

Keep the scorecard connected to decisions

Any metric on the main sheet that cannot inform an action should be redefined or removed.

WORKING EXERCISE

Complete an MRR Mission, check the cash and capacity needs under lower conversion, schedule the activities, and assign the review date.

MEASURE

The complete ladder with consistent definitions and mature cohorts.

18

CHAPTER

Build Your Next 90 Days

You do not need to activate every channel in this book at once.

What you need is a credible foundation, one focused acquisition effort, a working buying path, and enough measurement to decide what comes next. A 90-day plan gives that work a boundary. It is a management period, not a promise to reach product-market fit in three months.

Choose the current constraint

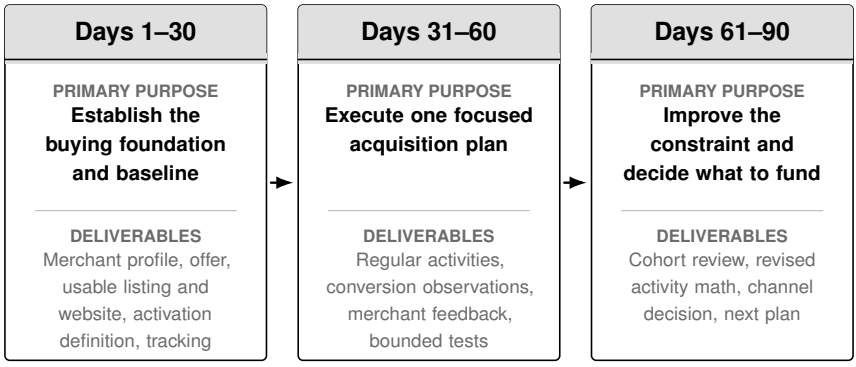
Write one sentence: "Our next growth constraint is _____, supported by _____."

The evidence might be merchants failing to finish setup, a strong conversion rate with too little qualified demand, or cancellations cancelling out new sales. Where the evidence is weak, collecting it is the first priority.

Then pick one primary acquisition channel to test or improve. Keep essential operations running in the others without treating each one as a major new initiative.

A working three-phase plan

FIGURE 23



A business that already has these foundations should start at its real constraint. Do not spend a month rebuilding assets because the table happens to begin there.

Convert the plan into a weekly commitment

PLANNING EXAMPLE

A founder has 15 hours a week for GTM. They reserve six for account research and outreach, four for demos and follow-up, three for content, one for analytics, and one for listing and website maintenance.

If demos take seven hours instead of four, the plan needs an explicit tradeoff: reduce another activity, improve efficiency, or add capacity. The calendar is a resource constraint, not a wish list.

Product-led apps usually shift more time into onboarding and advertising. Sales-led apps put more of it into conversations.

Your 90-day operating brief

Complete these eight lines:

1. Target merchant
2. Primary outcome we sell
3. Starting MRR ---- Target MRR
4. Expected losses ---- Required gross new MRR
5. Primary acquisition channel
6. Weekly activities and capacity
7. Maximum cash investment
8. Evidence required for the next spending decision

Review the brief against the decision rules below, and keep merchant feedback connected to the numbers rather than filed somewhere separate.

Use milestone gates

A milestone gate is a condition for making the next commitment. It is more useful than a date, because a date assumes the learning will be finished on schedule.

Two examples. “Before increasing the ad budget, we need verified payment attribution and an acquired cohort that has completed the observation window.” Or: “before hiring a seller, the founder needs a documented offer, a qualification process, and a recurring source of relevant accounts.”

A 90-day product-led plan

Apply the three phases above to the self-serve route. Appendix F1 or F3 has the filled funnel, the weekly commitments, the budget sensitivity and the conversion timing.

A 90-day sales-led plan

Apply the phases to founder-led signal research, bursts, evaluation and delivery. Appendix F2 and F4 have the filled account, demo, payer, workload and MRR plans.

Reserve time for the work success creates

An acquisition plan can fail operationally because it assumed new customers need no delivery effort. Five wins at four hours of setup each is 20 hours of work that nobody scheduled.

Reserve that capacity or name who will provide it. A founder who fills every available hour with prospecting closes customers into a queue, which delays their value and raises the odds they cancel.

Write decision rules before the result

Decide in advance what evidence would support continuing, changing, pausing, or increasing investment, and include both economic and operational conditions.

Both failure modes are real. A campaign can hit its acquisition-cost target and deliver customers who need unsustainable support. An event can produce no immediate sales and several qualified evaluations that simply need time. The rule has to be broad enough to recognize those and specific enough that you cannot reinterpret it conveniently later.

Record the decision even when the decision is to wait. “No change until this date because payment observations are incomplete” is a legitimate operating choice, and writing it down stops the question being reopened weekly.

Prepare the first delegation

When a task repeats successfully, document the inputs, the steps, the quality standard and the output. Start with one bounded responsibility rather than handing over the GTM system.

A researcher can build fit-checked account briefs. A producer can create content from approved buying questions. A seller can run a defined segment once the founder has demonstrated the motion works.

Finish with a reviewable commitment

Review the 90-day operating brief against the evidence. Record what changed in merchant fit, offer, acquisition, activation and retention, then set the next MRR Mission.

The test of the brief is simple: a teammate should be able to read it and understand the activities, the limits, and the conditions that would change the plan.

Appendices

The metrics, the weekly sheet, the assessment, the templates, the operating plans, and the troubleshooting guide.

The numbers founders need

Define the metrics before setting targets

New install	A newly connected store; flag reinstalls separately
Activation	A product-specific first-value event
Trial	Entry into a defined evaluation period
Paying merchant	A merchant with an active paying arrangement
MRR	Monthly-normalized recurring revenue under a consistent accounting definition
ARR	Annualized recurring run rate, commonly MRR × 12; not total sales or cash receipts
Gross new MRR	MRR first added by new paying customers
Increased existing-customer MRR	Additional MRR from customers already paying
Churned MRR	MRR lost through cancellation
Contraction MRR	MRR lost through lower plans or reduced recurring charges
Net new MRR	Gross new + increased existing-customer MRR + reactivation MRR − churn − contraction
Contribution per merchant	Revenue minus the variable costs of serving that merchant, using stated cost categories
CAC	Stated acquisition costs divided by new paying customers attributable to those costs
Contribution payback	CAC divided by monthly contribution per acquired merchant, assuming contribution remains comparable

Normalize a $1,200 annual subscription to $100 MRR, not $1,200 in the month you collect it. Keep bookings, invoices, cash, recognized revenue and recurring run rate distinct from one another. For variable or usage-based revenue, document how you normalize it and show the volatility rather than presenting it as a guaranteed subscription amount.

Revenue planning formulas

Ending MRR = starting MRR + gross new MRR + increased existing-customer MRR + reactivation MRR − churned MRR − contraction MRR.

Required gross new MRR = target net increase + expected churn and contraction − expected existing-customer increases and reactivation.

New payers required = required gross new MRR ÷ expected average new-customer MRR, rounded upward.

Installs required = new payers required ÷ install-to-paid conversion.

Qualified visits required = installs required ÷ visit-to-install conversion.

Held demos required = new payers required ÷ held-demo-to-paid conversion.

Bookings required = held demos required ÷ attendance rate.

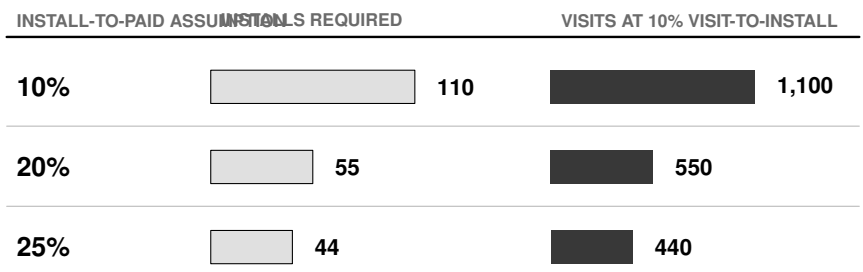
Accounts required = bookings required ÷ account-to-booking conversion.

Only divide by a rate greater than zero. When a rate is unknown, model a range and state the assumption. When it has been zero in a meaningful observed sample, diagnose the failure rather than hiding it behind an optimistic default.

Self-serve sensitivity: same goal, different conversion

Assume $1,100 gross new MRR is required and each new payer contributes $100 MRR. You need 11 payers.

FIGURE 24



The figure does not tell you which conversion rate is achievable. It shows why onboarding deserves attention. Improving the fraction of installs that become paying customers can reduce the acquisition volume required for the same goal.

Acquisition cost: label the scope

Suppose an ad test costs $600, creative costs $200, and attributable acquisition labor costs $400. Four new payers are acquired.

- Ad-only acquisition cost: $600 ÷ 4 = $150.
- Fully loaded acquisition cost under these stated categories: $1,200 ÷ 4 = $300.

Neither number should silently replace the other. Use a consistent cost policy across channels. An event reported with travel and labor should not be compared to outbound reported with only software cost.

Your first spending rules

Write these down now, while nothing is tempting you. The sponsorship proposal arrives later, and it always arrives with a deadline attached.

WORKSHEET 4

DECISION	YOUR LIMIT
Monthly acquisition cash budget	
Maximum cash at risk on one unproven test	
Founder hours available for GTM each week	
Review date for each paid experiment	
Minimum cash reserve	

Build three runway scenarios

Build three versions: a base case, a slower one, and a better one. What varies between them is acquisition timing, customer contribution, and how much you spend.

PLANNING EXAMPLE

With $36,000 available, $4,000 of fixed monthly costs, and $1,000 a month going to acquisition, you have 7.2 months before any customer contribution. Push acquisition to $2,000 and you have six.

Write a learning budget

Every acquisition expense should arrive with four things attached: the question it answers, the maximum cash and labor at risk, the date you will review it, and the observation that would change your mind.

> **EXAMPLE**
>
> No more than $500 on a single search group, then assess the install cohort once it has matured. If $500 cannot answer the question, narrow the question rather than quietly raising the budget.

Production work needs the same treatment. “Work on marketing this month” is not a budget. A completed buying asset with a plan for distributing it is one. Appendix D6 has the test brief.

The weekly operating sheet

Copy this structure into a spreadsheet or CRM report. Keep activity dates, conversion dates, and customer source records behind the summary.

WORKSHEET 5

The weekly operating sheet

WEEK STARTING _____ **OWNER** _____ **PRIMARY CONSTRAINT** _____

	KPI	TARGET	ACTUAL	PREV WK
ACTIVITY	Fit-checked new accounts			
	Outbound emails sent			
	LinkedIn requests			
	LinkedIn DMs			
	LinkedIn comments			
	LinkedIn InMails			
	Calls placed			
	Appropriate texts sent			
	Prospect contact-form submissions			
	Buying assets published			
	Short posts published			
COST	Acquisition cash spent			
	Acquisition labour hours			
DEMAND	Positive replies			
	Inbound contact forms			
	New qualified deals			
	Demos booked			
	Demos held			
ADOPTION	New installs			
	New trials			
	New activated merchants			
REVENUE	New paying merchants			
	Gross new MRR			
	Increased existing-customer MRR			
	Reactivated MRR			
	Churned MRR			
	Contraction MRR			
	Net new MRR			

BALANCES AT WEEK-END

LinkedIn connections _____ active qualified deals _____

total paying merchants _____ ending MRR _____

Conversion review: Use a mature cohort or a clearly defined rolling window. Do not divide this week's wins by this week's emails and call it the outbound conversion rate when the wins originated earlier.

Decision log: We observed _____. We believe the issue is _____. We will change _____. Owner _____. Review date _____. Maximum cash and time at risk _____.

The GTM STANDARD assessment

Use this to locate missing operating capabilities. It is an internal readiness assessment, not a validated benchmark against other apps. Score every statement:

- Missing (0). No usable process or evidence.
- Defined (1). Written down but not used consistently.
- Operating (2). Used consistently with an owner.
- Measured (3). Used, measured, and revised from evidence.

WORKSHEET 6

The GTM STANDARD assessment

	STEP AND THE THREE STATEMENTS TO SCORE	SCORES	
S	**Select** We have a specific merchant profile. We apply exclusion rules. We compare customer quality by segment.	☐ ☐ ☐	/9
T	**Tune** We state a clear outcome. We explain why merchants should choose us. We bound incentives and delivery effort.	☐ ☐ ☐	/9
A	**Assemble** Our listing and website explain fit. Our buying route works. Our onboarding leads to a defined value event.	☐ ☐ ☐	/9
N	**Name** We record signal sources and dates. We separate facts from assumptions. Our outreach reflects verified context.	☐ ☐ ☐	/9
D	**Drive** We have recurring activity targets. We can deliver the plan within capacity. We track channel costs and outcomes.	☐ ☐ ☐	/9
A	**Advance** We use a relevant demo process. Every active deal has a next action. We distinguish trial from payment.	☐ ☐ ☐	/9
R	**Retain** Activation has an owner or monitored process. We record cancellation reasons. We review retained revenue by cohort.	☐ ☐ ☐	/9
D	**Diagnose** We maintain a weekly scorecard. Revenue targets connect to activity assumptions. Spending decisions follow dated reviews.	☐ ☐ ☐	/9
Total		_____	**/72**

Use the lowest relevant step to choose the next action. A high total can still hide a serious weakness. If the onboarding path is broken, strong content and CRM discipline do not cancel that problem out.

Do not attach labels such as “elite” or “top 10%” to these scores without validation. Repeat the assessment to see whether the business is putting the practices into operation.

Templates founders can use

1. Merchant account brief

Merchant / buying entity

Store(s)

Buyer role / operator role

Fit evidence

Observed signal and source date

What we know

What we still need to verify

Relevant advantage

Likely revenue and service requirements

Owner / next action / date

2. Four-round burst worksheet

WORKSHEET 7

ROUND	DATE	CHANNELS USED	NEW VALUE ADDED	ACTUAL RESULT
1			Signal and advantage	
2			Proof or demonstration	
3			Additional benefit or offer	
4			Decision on whether to continue	

Stop conditions: Reply, decline, opt-out, disqualification, or account-specific reason to pause. Reactivation reason, if appropriate later: New evidence, materially changed offer, or agreed later timing.

3. Reply to merchant interest

> Hey *[Name]*, happy to show you. I can do *[time option 1]* or *[time option 2]*, or you can grab a time here: *[calendar link]*.

Use the merchant's timezone or state the timezone clearly. Offer times you can actually attend.

4. Demo preparation brief

Why this merchant is considering change

Known current workflow

Questions that matter

Relevant deck section

Live example

Back-end workflow

Financial assumptions to verify

Offer and real deadline, if any

Desired decision or next step

5. Review recovery record

Review date / source

Merchant’s stated issue

Owner

Contact attempt / conversation date

Resolution agreed

Merchant confirmation

Neutral update request, if appropriate

Product or process lesson

Resolution must never depend on a favorable rating or review removal.

6. Paid test decision brief

Hypothesis

Target audience and channel

Conversion event and observation window

Maximum cash spend

Maximum labor commitment

Modeled funnel

Required economic result

Reasons to stop immediately: Broken route, irrelevant targeting, or another explicit failure: _____.

Review date and decision

7. Event decision brief

Event and target audience

Cash cost / labor cost

Meeting package or pre-booking plan

Qualified held-meeting target

Customer and MRR hypothesis

Follow-up owner and dates

Cohort review date

8. Partner pilot brief

Partner and shared merchant profile

Reason the relationship helps merchants

Value we can provide

Introduction criteria and process

Owners on both sides

Compensation, if any

90-day review: introductions, customers, MRR, cost, and effort:

A 30-minute weekly review

The timing below is a suggested meeting format, not a requirement to finish a complex issue in a few minutes.

Minutes 0–5: Revenue. Starting MRR, new MRR, losses, and ending MRR. Identify material cancellations or overdue activations.

Minutes 5–10: Funnel. Compare visits, accounts, demos, installs, trials, activation, and payment against the plan. Note conversion lag and immature cohorts.

Minutes 10–15: Activities. Compare completed work with targets. Determine whether the gap was capacity, execution, or an unrealistic plan.

Minutes 15–20: Merchant evidence. Review recurring objections, onboarding friction, customer feedback, and reasons for loss.

Minutes 20–25: One priority. Choose the most consequential constraint and the next test. Avoid creating five simultaneous explanations for one result.

Minutes 25–30: Commitments. Assign owners, dates, budget, and the evidence needed for the next decision.

A solo founder can run the same review in writing. The value comes from the decisions and follow-through, not from the meeting itself.

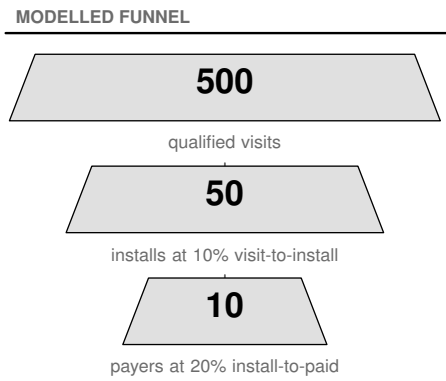
Operating plans by MRR stage

These plans are worked examples using stated assumptions, and they exclude taxes, financing and corporate accounting detail. The stages are revenue checkpoints, not validated levels of company quality. Replace the inputs with your own.

Plan 1: First 10 paying merchants through a self-serve route

FIGURE 25

Plan 1: the first 10 paying merchants



MODELLED ECONOMICS	
MRR per new merchant	**$100**
Gross new MRR objective	**$1,000**
Media budget	**$1,000**
Other acquisition cost	**$500**
Total modelled CAC	**$150**
Contribution payback	**1.9 mo**

At $80 monthly contribution, modeled payback is $150 ÷ $80 = 1.875 months, or about 1.9 months. This assumes the same measured funnel and sufficient retention; availability of 500 relevant $2 visits is unverified.

Schedule: Over four acquisition weeks, seek 125 visits weekly, implying 12.5 expected installs and 2.5 payers per mature cohort. Actual counts are whole people and will vary. If payment takes 30 days, observe the final week's cohort beyond acquisition month-end.

Weekly work: Check spend and the buying route; publish one buying asset and distribute it through three suitable uses; inspect setup abandonment and speak with an appropriate merchant sample. Verify the route on an early small cohort before consuming the full budget.

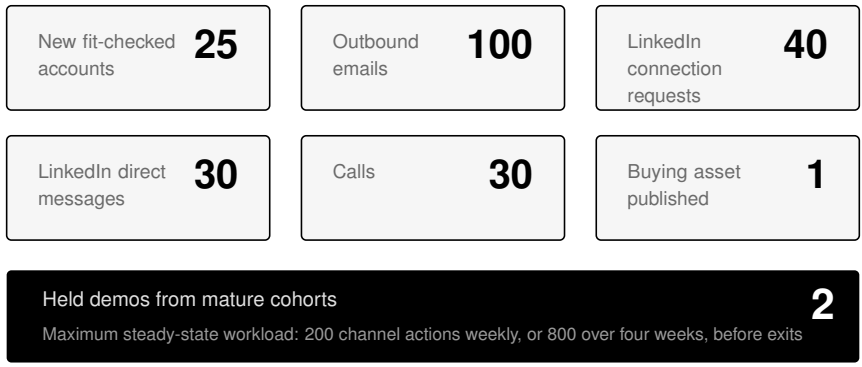
Decision: Continue on supportable audience fit and economics; revise the failing stage or pause poor-fit acquisition. Keep attribution records after stopping. Use Chapters 5–8 for execution.

Plan 2: First 1K MRR through founder-led sales

Model: Two payers × $500 MRR = $1,000 gross MRR. At 25% held-demo-to-paid conversion, require eight held demos. At 8% account-to-held-demo conversion, require 100 qualified accounts.

FIGURE 26

Plan 2: weekly work targets



Cadence: Four acquisition weeks, with up to four rounds and two appropriate actions per round. The maximum steady-state workload is 200 actions weekly, or 800 over four weeks, before exits. Early weeks contain fewer rounds. Availability, accepted connections, and replies change the channel mix; count actual work.

Capacity and timing: Reserve research, reply, demo, and delivery hours. Two wins needing five onboarding hours each add 10 hours. A one- or two-month sales cycle places payment after the original outreach period.

Decision: No replies: inspect audience, signal, delivery, and message. Demos without purchases: inspect fit, offer, and decision process. Excess custom work: revise qualification or price. Chapters 9–13 provide the procedures.

Plan 3: From 1K to 5K MRR in a planning period

Model: At $100 MRR per new payer, a $4,000 net increase plus $500 expected cancellation and downgrade losses requires $4,500 gross new MRR: 45 payers. At 20% install-to-paid and 10% visit-to-install conversion, require 225 installs and 2,250 relevant visits. At $2 per attributable paid visit, modeled media cost is $4,500. Record organic or referral contributions separately.

Schedule: In a 90-day plan, acquire during the first nine weeks at 250 visits weekly: 25 expected installs and five expected payers per cohort. Reserve the remaining period for conversion maturity and review; adjust if actual lag is longer.

Sensitivity: At 10% rather than 20% install-to-paid, the original volume yields half the payers. Reaching 45 payers would require 450 installs, 4,500 visits, and $9,000 media spend at the same price. Check early cohorts before committing the full amount.

Weekly work: Review campaigns, publish one focused buying asset, act on the largest setup obstruction, and reconcile MRR movements. Maintain listing and website against observed buying questions.

Capacity: Include support for 45 new merchants in the service model.

Decision: Buy more volume only when mature cohorts and service economics support it. Improve activation first if it is failing; test adjacent demand under a separate label and budget if the audience is limited.

Plan 4: From 5K to 10K MRR with assisted sales

Model: A $5,000 net increase plus $1,000 expected losses requires $6,000 gross new MRR. At $500 per payer, require 12 payers. At 20%

held-demo-to-paid, 80% attendance, and 10% qualified-account-to-booking conversion, require 60 held demos, 75 bookings, and 750 accounts.

Schedule: Over 12 acquisition weeks, 62.5 accounts weekly rounds to 63, or 756 total. Eight maximum actions per account produce about 504 weekly actions before exits. At three minutes each, actions alone take 25.2 hours, excluding research, demos, replies, and setup. This does not fit a 15-hour founder allowance.

Capacity options: Extend time, lower the target, add supported capacity, improve conversion with evidence, change customer value through a real offer, or fund another source. Splitting six payers to outbound halves that funnel; the other six still need their own funnel and budget. Allow time beyond the acquisition schedule if deals remain immature.

Weekly work: Maintain the queue, finish rounds, handle replies, prepare demos, keep dated next actions, and reserve delivery hours.

Decision: Add a seller only when audience, offer, materials, ownership, and sufficient volume are established. See Chapters 10, 12, 13, and 18.

Plan 5: Protect 10K MRR before adding another channel

Sensitivity: At $10K MRR, 5% monthly revenue loss removes $500; 10% removes $1,000. With $1,000 gross new MRR and no other changes, net growth is $500 or zero. These are scenarios, not acceptable churn targets.

Before adding a channel, review activation, retention, service load, and customer concentration using Chapter 16. Maintain the proven route, address one preventable loss, and document a repeated

task for delegation. Add a channel against a specific access or concentration constraint.

Adapt the plans to your price

To add $1,000 gross MRR, a $20 app needs 50 payers, a $100 app needs 10, and a $500 app needs two. Recalculate funnel, service cost, and activity capacity before copying another app's targets.

Troubleshooting the GTM funnel

Confirm measurement and cohort maturity first. Then use the symptom to choose the investigation below; each diagnosis remains a hypothesis until checked. The referenced chapters contain the full procedures.

Outreach produces no useful replies

CHECK FIRST Fit gate, buyer role, truthful signal, relevant advantage, one action, and delivery. Exclude auto-replies and declines from positive replies.

CORRECTIVE ACTION Run a small revised cohort; change targeting, evidence, message, or delivery before raising volume. Chapters 2, 9, 10.

Bookings do not attend

CHECK FIRST Invitation receipt, meeting purpose, time, participant fit, confirmation and rescheduling. Separate outbound, organic, event, and compensated sources.

CORRECTIVE ACTION Repair booking or recruitment; compare attendance by source. Chapters 6, 13, 14.

Demos do not advance

CHECK FIRST Merchant problem, relevant workflow, price, limitations, setup, and buyer next action.

CORRECTIVE ACTION Resolve the specific blocker, involve the missing stakeholder, or close out. Repeated capability gaps require product or qualification decisions. Chapters 12, 13.

Installs fail to activate

CHECK FIRST First required setup step and whether the acquisition promise matches available capability.

CORRECTIVE ACTION Observe use, ask incomplete accounts, change one obstruction, and compare cohorts including inactive merchants. Chapter 7.

Trials fail to pay

CHECK FIRST First value before expiry; pricing clarity; fit; value; missing features; timing and incomplete observation windows.

CORRECTIVE ACTION Match the fix to the cause: setup, plan explanation, qualification, or assistance. Chapters 4, 7.

Wins do not produce net MRR

CHECK FIRST Paid-versus-trial definition, annual normalization, all revenue movements, and customer value. Ten new $20 payers minus two lost $200 payers means −$200 net MRR despite eight net new customers.

CORRECTIVE ACTION Reconcile the ledger, then address acquisition or losses using the revenue ladder. Chapters 16, 17; Appendix A.

CAC looks good but cash is tight

CHECK FIRST Media versus full CAC, labor, incentives, setup, contribution, receipt timing, and annual service obligations.

CORRECTIVE ACTION Reconcile cash and MRR; use supportable payback and retention evidence rather than speculative lifetime value. Chapters 1, 3, 8.

Founder cannot keep up

CHECK FIRST Acquisition hours plus research, replies, demos, setup, and support created by wins.

CORRECTIVE ACTION Reduce scope or delegate one defined task while preserving the customer feedback loop. Chapters 17, 18.

For the chosen fix, record the action, owner, resource limit, and review date in Appendix B. Use Appendix E to review it.

Sources and attribution

The author's experience

The practical positions and anonymized examples in this book come from the author's project conversations and correspondence. They include App Store-led acquisition, founder-led outbound, burst prospecting, merchant demos, follow-up, event attendance, partnerships, review recovery, and self-serve onboarding. Case outcomes are founder-reported rather than independently audited. This book does not reproduce identifying names or emails from the restricted company or merchant examples. Alert is included as a signal research example.

Books referenced

Predictable Revenue: Turn Your Business Into a Sales Machine with the $100 Million Best Practices of Salesforce.com. Aaron Ross and Marylou Tyler. Pebblestorm, 2011. ISBN 978-0-9843802-1-3. Informs the emphasis on target selection, repeatable prospecting, and defined process. Historical response-rate claims and team structures are not imported as current Shopify expectations.

Hacking Sales: The Playbook for Building a High-Velocity Sales Machine. Max Altschuler. Wiley, 2016. ISBN 978-1-119-28164-1.

Informs the use of existing technology, competitor usage, and research workflows to select relevant accounts. Tool lists in the older source are not presented as current recommendations.

Sales Superpowers: A New Outbound Operating System to Drive Explosive Pipeline Growth. Justin Michael. Jones Media Publishing, 2023. ISBN 978-1-948382-68-7. Informs the discussion of coordinated channels and repeatable prospecting. The four-round implementation in this book is a practical operating choice rather than a finding of that source.

The book's implementation tools combine author experience with common GTM disciplines.

Official online references

[W1] Shopify, App Store requirements. Shopify developer documentation. Supports the need to check current listing rules and restrictions on listing content. https://shopify.dev/docs/apps/launch/shopify-app-store/app-store-requirements

[W2] Shopify, Always use honest and transparent review practices (App Store requirements, section 1.3). Supports neutral review requests, non-incentivized requests to update a review after resolution, and recovery practices. https://shopify.dev/docs/apps/launch/shopify-app-store/app-store-requirements#always-use-honest-and-transparent-review-practices

[W3] Shopify, Shopify App Store ads. Shopify developer documentation. Supports the general description of advertising placements and budget and bid controls. https://shopify.dev/docs/apps/launch/marketing/advertising

[W4] Sandler, The Sandler Selling System. Sandler Systems, LLC. Supports attribution of the up-front contract and structured sales conversation to Sandler. https://www.sandler.com

All four were checked in November 2026. Platform documentation changes, so recheck the requirements before implementation. This book does not attempt to provide a jurisdiction-specific outreach compliance manual.

About Alert

Alert is a sales intelligence platform built for Shopify app companies. It monitors the merchant base for the events that mean somebody has a reason to switch: a competitor installed, a competitor uninstalled, a negative review left on a category rival, a migration to Shopify Plus, a careers page that suddenly lists an ecommerce operator.

Most Shopify sales tools help you find merchants. There are a lot of merchants. The harder problem, and the one this book is about, is finding the few who have a reason to act this week, and reaching them with something true. Alert exists to make that research repeatable rather than heroic.

What to do next

The chapters in this book work without buying anything. If you have read this far, the useful next step is the one your own constraint points at: a merchant profile you can check, a listing audit, an activation definition, a week of real outreach you run yourself.

When you want the research to stop being manual, the two routes below are the fastest way to see whether Alert fits your app.

Book a demo

Thirty minutes on your app, your segment, and the signals worth watching. Bring a competitor name and we will look at live merchant data together.

https://alertcommerce.com/demo/

The GTM Standard resources

The signal-play templates, the weekly operating sheet, the assessment, and the operating plans from the appendices, in formats you can fill in and share with a team.

https://alertcommerce.com/gtm-standard

Every worksheet and template is printed at full size, in the chapter that uses it or in the appendices, so you can work entirely from the printed copy if you prefer.