CLIENT CONNECTION CONSULTING

The Compounding Customer.

The leader's guide to customer success, retention, and the economics that matter.

A strategic guide

for founders, CEOs and senior leaders

EDITION 1

CLIENT CONNECTION CONSULTING

The Compounding Customer

The leader's guide to customer success,
retention, and the economics that matter.

A strategic guide for founders, CEOs and senior leadership teams

The Compounding Customer: The Leader's Guide to Customer Success, Retention, and the Economics That Matter

First edition published 2026.

The author and publisher have made every reasonable effort to ensure the accuracy of the information in this book. The book is intended as a general guide to building and leading a Customer Success function and is not a substitute for tailored professional advice.

Published by Client Connection Consulting.

For enquiries: clientconnectionconsulting@gmail.com

CONTENTS

What's inside.

OPENING

PART ONE — THE STRATEGIC CASE

PART TWO — BUILDING AND RUNNING THE FUNCTION

APPENDIX — REFERENCE MATERIAL

INTRODUCTION

Why This Book, Why Now

If your business depends on customers coming back, the function that keeps them is no longer optional. And if you are leading the business, the function that keeps them is your problem, not someone else's.

There was a time, not very long ago, when a sale was a complete event. A customer paid, a product was handed over, and the commercial relationship was largely finished. Customer service existed, but it was a complaints desk, sized to limit the damage when something went wrong. It was a cost on the balance sheet, never a source of growth.

That world is gone for most of the companies likely to be reading this book.

In any business where customers can choose to leave every month, every quarter, or every year, the value of a customer is no longer fixed at the moment of sale. It is determined later, by how long they stay, how much they expand, and how often they recommend you. Software-as-a-service companies were the first to confront this reality, but they are no longer the only ones. Media subscriptions, fintech, infrastructure, hardware-as-a-service, marketplaces, professional services on retainer: anywhere revenue is decided after a contract is signed rather than at signature, the same economic logic applies.

And that economic logic is brutal: *in a recurring-revenue business, what happens after the sale matters more than what happened during the sale.* The customer you signed last quarter is not yet revenue. They are a hypothesis. Whether the hypothesis pays out depends on a function that, in most companies, is still under-built, under-funded, and reporting to someone who does not fully understand what it does.

That function is Customer Success.

5×

The cost of acquiring a new customer versus retaining an existing one

70%

Of mature SaaS revenue now comes from existing customers, not new logos

5pt

Lift in retention can lift profits by 25 to 95 percent

The Problem This Book Is For

Most books about Customer Success are written for the people doing the work: Customer Success Managers, team leads, the people running quarterly business reviews and tracking health scores. Useful, if that is your job.

This book is for the people one or two levels above that. The founder deciding whether to hire the first Customer Success Manager, and whom to hire. The CEO trying to understand why retention is dragging down growth despite a healthy sales pipeline. The Chief Revenue Officer working out whether Customer Success should report to them or somewhere else entirely. The Head of Customer Success who has the team but does not yet have the budget, the seat at the table, or the language to explain to a board what their function actually produces.

These are different questions from "how do I run a better onboarding call." They are strategic, organisational, and financial. They do not get answered well by a tactical playbook, and they do not get answered at all by a book that treats Customer Success as a feel-good function rather than a financial one.

This book is an attempt to answer them.

> *If your CS team disappeared tomorrow, which numbers would move? If the answer is "support ticket volume," you do not have a Customer Success function. If the answer is "net revenue retention, gross retention, and expansion bookings," you do.*

What This Book Is About

Customer Success is a young discipline. The first job titles appeared in the mid-2000s, the first formal teams not long after, and the field has been moving fast ever since. As often happens with fast-moving fields, the practice has grown faster than the consensus around it. Two CS leaders from two companies in the same industry will often describe what they do, what they measure, and who they report to in completely different terms.

This book takes a position on most of those disagreements. Where the field is genuinely split, you will see arguments for one approach over another, not a neutral summary of every option. Where there is broad agreement, that gets stated plainly and moved on from. The intent is to be useful to a leader making decisions, not to be exhaustive.

It is organised in two halves.

- ***Part One — The Strategic Case***

  For founders, CEOs, and boards working out what Customer Success should do for the business, how it differs from the functions it gets confused with, and how to think about its economic impact. Four chapters covering the definition, the customer journey as a framework, the business case for investment, and the financial metrics that matter at board level.

- ***Part Two — Building and Running the Function***

  For senior leaders translating the strategy into a working organisation. Eleven chapters covering operating models, hiring, onboarding, customer health, retention, expansion, technology, scaling, industry variation, and what artificial intelligence is genuinely changing about the work.

The two halves are independent. A CEO who only ever reads Part One will have what they need to make the central decisions: whether to invest, how much, when, and what to hold the function accountable for. A Head of CS who picks the book up to solve a specific operational problem can skip straight to the relevant chapter in Part Two without losing the thread.

Who This Is For

Four kinds of reader will find this book useful.

Founders and CEOs of recurring-revenue businesses who are either about to build a Customer Success function, or have built one and are unsure whether they built the right thing. This book answers the question of what the function is for, what it should cost, and what it should deliver.

Chief Revenue Officers and Chief Customer Officers who own retention and expansion as a number, and who need to translate that ownership into team structure, comp design, and a credible plan they can defend to a CEO or board.

Heads of Customer Success who already run a team but are working through the harder questions: how to scale beyond a generalist model, how to win the resources their function needs, and how to be taken seriously as a peer to Sales and Product rather than a service function attached to them.

Board members and investors who want a working understanding of the metrics CS produces, why they matter, and what good looks like at different stages of company maturity. The chapter on economics in particular is written with this reader in mind.

It is not written for Customer Success Managers, although a good Customer Success Manager will find a lot in it that helps them think about their work in commercial terms. It is also not written for support agents, account managers, or implementation consultants, although the boundaries between those functions and Customer Success are addressed directly in Chapter 1.

A Quick Note on Terminology

Customer Success has acquired a thicket of acronyms and jargon, much of it borrowed from sales and finance, some of it invented by software vendors. Where a term might be unfamiliar to a reader outside the field, you will see a brief inline definition the first time it appears, like this: *Net Revenue Retention (NRR — the percentage of recurring revenue you retain from existing customers year over year, including expansion and minus churn).* A full glossary is provided at the back of the book.

How to Read This Book

The book is designed to be read in three different ways, depending on what you need.

Cover to cover

If Customer Success is new to you, or new to your business, read it in order. Part One builds the case and the framework; Part Two builds the function. By the time you finish, you will have both the argument and the operating model.

Strategic only

If you are the CEO or founder and you need to make a decision rather than build an organisation, read Part One. It is roughly a quarter of the book and is self-contained. Hand the rest to whoever will be running the function.

Operational only

If you already have a CS function and are working on a specific problem, the chapters in Part Two are designed to stand on their own. The table of contents is structured around the questions a CS leader actually asks: how do I design my operating model, how do I price renewals, how do I structure expansion, how do I scale beyond ten Customer Success Managers.

Most readers will use a mix. The pages reward dipping in and out as much as they reward straight reading.

One Thing Before We Start

Building a Customer Success function is not, primarily, a matter of buying the right software or copying the right playbook. Both of those things help, and this book covers both. But the central decisions are organisational and commercial, and they have to be made by the people who set the company's direction.

If those decisions are made well, the function compounds. Retention rises, expansion follows, customer advocacy lowers your acquisition costs, and within a few years the post-sale economics of the business start to do more work than the sales engine that fills the pipeline. If those decisions are made badly, or not made at all, the function becomes an expensive and demoralised retention queue, and the business is left depending on its sales team to outrun a churn problem that is structural rather than tactical.

The difference between those two outcomes is not luck, and it is rarely the quality of the individual Customer Success Managers. It is the decisions that get made above their heads.

That is what this book is for.

THE BOTTOM LINE

In a recurring-revenue business, **the post-sale function determines whether the company compounds or stalls.** Customer Success is the name for that function. **Building it well is a leadership decision, not a tactical one**, and that is the decision this book is written to help you make.

CHAPTER 1

What Customer Success Actually Is

Three different functions get called Customer Success in three different companies, and a fourth thing that should be called Customer Success usually isn't. This chapter is an attempt to settle that.

Customer Success is a young enough discipline that you can still find two senior leaders in the same industry using the same job title to describe completely different jobs. One runs a renewals desk and reports to Sales. Another runs a strategic advisory team and reports to the CEO. A third runs what is effectively a premium support queue and reports to Operations. They are all called Head of Customer Success.

This is a problem for anyone trying to make decisions about the function. If the term means four different things, then "we should invest in Customer Success" doesn't mean anything until you specify which version. A CEO who hires a Head of CS expecting one thing and gets another has not made a hiring mistake; they have made a definitional one. The rest of this book is built on the definition this chapter is about to give you. If that definition is unstable, nothing built on top of it holds.

A Working Definition

Customer Success is the function responsible for ensuring that customers achieve measurable outcomes from your product, and that those outcomes translate into revenue your business can count on: retention, expansion, and advocacy.

That definition is deliberately narrow. It excludes a lot of what gets called Customer Success in job ads and on organisation charts: glorified support, renewal-chasing, account farming with a friendlier title. Those are activities Customer Success teams sometimes perform, but they are not what the function exists to do.

The function exists because, in any business with recurring revenue, the economics of the post-sale period are larger than the economics of the sale itself, and someone has to own them. *Recurring revenue, in this context, means any business model where the customer can choose to renew, expand, or cancel after the initial purchase: subscriptions, retainers, usage-based contracts, and so on.* If the customer can leave, the function that keeps them is no longer optional.

> *If your Customer Success team disappeared tomorrow, which numbers would move? If the answer is "support ticket volume" or "account manager workload," you don't have a Customer Success function. If the answer is "net revenue retention, gross retention, and expansion bookings," you do.*

A mid-sized B2B software company ran exactly this exercise at a leadership offsite. The CEO asked each function head what would break if a peer function vanished. When the CS question came up, the room went quiet, then the CRO said, honestly, "renewals would still go out, and the AMs would handle escalations." The CS team had thirty people. The company spent the next eighteen months rebuilding the function around outcomes the CEO could point to on a board deck.

That diagnostic is the test the rest of this chapter is built around. Most companies discover, when they ask the question honestly, that what they call Customer Success is not producing the numbers they want, because it was never structured to.

Why the Function Exists At All

Customer Success is a relatively new discipline because the business model that requires it is also relatively new. For most of commercial history, a sale was an event: money changed hands, the product was delivered, the relationship was largely complete, and customer service existed only to handle complaints. Subscription and consumption-based business models broke that arrangement. When the same customer can leave every month, every quarter, or every year, the value of the customer is no longer fixed at the moment of sale. It is a function of how long they stay and how much they expand.

Software-as-a-service companies built the discipline first, in the mid-2000s, because they felt the pain first: their entire revenue base could turn over in a year if customers chose to leave. The shift has since spread to any business with a recurring revenue component, from media and content subscriptions to fintech, infrastructure, hardware-as-a-service, marketplaces, and professional services on retainer. Wherever revenue is decided after the contract is signed rather than at signature, Customer Success has a job to do.

The shift in plain numbers

In a one-time-sale business, a customer who pays $10,000 is worth $10,000. In a subscription business, that same customer paying $10,000 a year is worth $10,000 if they leave after twelve months, $50,000 if they stay five years, and $80,000 if they stay five years and expand. The variance is large, and the function that determines which outcome you get is Customer Success.

This is also why Customer Success matters more to some businesses than others. A company selling one-off services to a fragmented market may not need a Customer Success function at all. A company whose top twenty customers represent fifty per cent of revenue almost certainly does, regardless of whether it has built one yet.

What Customer Success Is Not

Three other functions are routinely confused with Customer Success, and the confusion costs companies real money. It leads to building the wrong team, hiring the wrong people, and measuring the wrong things. Sorting out the boundaries is one of the most important things a leadership team can do early.

The differences are best shown side by side.

DIMENSION	Customer Support	Account Management	Customer Success
POSTURE	Reactive. Responds to issues raised by customers.	Commercial. Owns the renewal and upsell conversations.	Proactive. Initiates conversations about outcomes.
TRIGGER	Customer raises a ticket.	Contract date approaches.	A signal in usage, sentiment, or relationship data.
QUESTION ASKED	How do I solve this issue?	What else can we sell you?	What outcomes are we delivering, and what is blocking the next one?
PRIMARY METRIC	Resolution time, CSAT.	Revenue per account, quota attainment.	Net revenue retention, gross retention, expansion.
TIME HORIZON	The current interaction.	The current contract cycle.	The customer lifetime.

A fourth function, **Professional Services**, is paid implementation work. It is a revenue line, not an investment in retention. Professional Services and Customer Success overlap during onboarding, and the handoff between them is one of the most common failure points in the customer journey. We return to it in Chapter 7.

None of these functions is lesser than Customer Success. All four exist in mature companies, and the boundaries between them are operating-model decisions a leadership team has to make deliberately. Chapter 5 walks through how to draw those boundaries. For now, the point is that they are different jobs, and treating them as interchangeable is one of the most common and most expensive mistakes founders make in the first few years of building a recurring-revenue business. Three patterns come up repeatedly.

The strategic hire doing tactical work

A Series B SaaS company hires a senior CSM from a much larger company, expecting them to build a function. Within three months, the new hire is spending most of their time in Zendesk, working through escalations and onboarding tickets that the support team is too thin to handle. They were hired to build, but the company has staffed them as a triage resource. They leave inside a year, and the CEO concludes that "senior CS hires don't work at our stage," when what they actually proved was that they bought a strategist and put them on a queue.

The salesperson with a CS title

A mid-stage company rebrands its Account Managers as Customer Success Managers and gives them an upsell quota. Renewals look fine for two quarters because the AMs know how to close. By the third quarter, expansion stalls: the CSMs have pushed upsell conversations on customers who weren't ready, and the trust that produces real expansion has eroded. The customers who renewed at full price are now negotiating discounts at the next cycle. The company has produced a worse outcome than if they had simply kept the Account Management function as it was.

The support promotion

A growing company promotes its best support agent into a newly-created Customer Success role because they "know the customers." Six months in, the role exists but produces nothing the support function wasn't already producing. The new CSM is reactive by temperament, was trained by years in support to wait for tickets, and has neither the commercial fluency to drive expansion nor the strategic posture to drive adoption. The role looks busy. The numbers do not move.

Each of these is a definitional mistake before it becomes a hiring mistake. Get the definition right first, and the people you hire will look very different from the people who already work in adjacent functions. The discipline of hiring is downstream of the discipline of definition; getting them in the right order is one of the more durable competitive advantages a leadership team can build.

The Three Flavours of Customer Success

Even with a clean definition, Customer Success teams in different companies do meaningfully different work. The reason is not that the definition is wrong; it is that the function operates at three different levels of ambition, and most companies are unclear about which level they are at. Understanding which flavour you have is the precondition for deciding which flavour you want.

1 ***The Cost Centre***

Customer Success exists to reduce churn. The team is sized cheaply, reports into Operations or Support, and is measured on retention and CSAT. It does not own expansion. *Looks like: an early-stage SaaS company with a small CSM pool under the COO, measured on logo retention only.*

2 ***The Retention Driver***

Customer Success is a recognised revenue function. It owns the renewal motion, often parts of expansion, and is measured on gross and net retention. It reports into the CRO or CCO. *Looks like: a public mid-market B2B software company with a CCO on the executive team.*

3 ***The Growth Driver***

Customer Success is a primary engine of company growth. It owns expansion bookings, produces advocacy that materially reduces acquisition costs, and is peer to Sales and Product. *Looks like: a late-stage enterprise SaaS company whose investor materials lead with net revenue retention.*

Most companies confuse which flavour they have. A founder describes their function as a growth driver while funding it at cost-centre levels. The result is a function that disappoints everyone, because it is being held accountable for outcomes it was never designed to produce.

Choosing the Right Flavour for Your Business

Each flavour is appropriate for some businesses and inappropriate for others. The choice depends on three factors: how much of your revenue is recurring, how concentrated your revenue is in your top customers, and how strong your product's natural retention is.

A business with low recurring revenue, fragmented customer base, and high natural product retention can probably get away with a cost-centre model indefinitely. A business with high recurring revenue, concentrated revenue, and weak natural retention almost certainly needs to be at least a retention driver, and probably needs to aim for growth driver within a few years.

A simple test

If you can answer the following three questions clearly, you can usually work out which flavour you need.

What percentage of next year's revenue is decided by customers we have already signed? If the answer is more than half, you cannot afford a cost-centre model.

If our top ten customers churned tomorrow, what happens to the business? If the answer is "we are in serious trouble," you need at minimum a retention driver, and you need it before those customers start showing warning signs, not after.

What percentage of our growth, realistically, is going to come from expansion of existing customers versus new logo acquisition? If the answer is more than thirty per cent from expansion, you are not just buying retention insurance any more, you are buying growth, and the function needs to be set up to produce growth.

These questions cannot always be answered cleanly, particularly in young companies whose customer base is still forming. But the answers do not need to be precise. They need to be honest enough to make the structural decision. The cost of getting the structural decision wrong is several years of frustration with a function that was never built to do what you were asking it to do.

What Good Looks Like

Across all three flavours, a few markers separate well-functioning Customer Success teams from poorly-functioning ones. None of them are universal, but they are common enough to be worth checking.

It produces a number, not an activity

A good function reports on outcomes, not effort. The conversation with the board is about gross retention, net retention, expansion bookings, and customer health distribution. Activity metrics (how many QBRs were held, how many onboarding sessions completed) matter internally for managing the team. They are not the function's deliverable to the rest of the business.

It has an opinion about customers, not just data about them

A team that can tell you which fifteen accounts are at greatest risk this quarter, why, and what is being done about each one, is doing the job. A team that can only show you dashboards without a point of view on which numbers should worry you is not.

It is trusted by the customer, not just liked by them

There is a difference between a CSM whose customers enjoy the relationship and a CSM whose customers act on their advice. The second is what you are paying for. A team whose customers like them but ignore them is producing meetings, not outcomes.

It has a seat at the table

Product consults Customer Success before shipping changes that affect existing customers. Sales consults Customer Success before pursuing a customer profile the function thinks will not succeed. Finance consults Customer Success when building the renewal forecast. This happens because the function is trusted to have a useful point of view, not because someone wrote it into a process document.

A function with none of these markers is not a Customer Success function in the sense this book uses the term. It is an adjacent function with a misleading name, and renaming it will not change the outcomes it produces. The markers are not a checklist to satisfy, but a set of observable behaviours that emerge when the function is structured around outcomes rather than activity. Where they are absent, the structural problem is almost always upstream of the team itself.

The Working Definition, Restated

By way of summary, the definition this book is built on:

> *Customer Success is the function responsible for ensuring that customers achieve measurable outcomes from your product, and that those outcomes translate into revenue your business can count on. It is proactive rather than reactive, focused on customer lifetime rather than the current contract, and accountable for retention, expansion, and advocacy as commercial outcomes.*

It is not Customer Support (reactive and ticket-driven), Account Management (commercial and quota-driven), or Professional Services (paid implementation). It can be operated at three levels of ambition (cost centre, retention driver, growth driver), and the level you operate at should be a deliberate decision tied to how your business makes money.

A growth-stage analytics company applied this definition in a planning cycle and discovered that what it had been calling Customer Success was structurally a Cost Centre, while its board had been holding the function accountable to Growth Driver metrics. Once the gap was named, the company built a 24-month plan to move deliberately to a Retention Driver model. That kind of clarity is what a working definition produces.

Chapter 2 takes the next step. It introduces the customer journey: the framework Customer Success teams use to organise their work, and the lens through which a leadership team can see where money is made and lost across the lifecycle of a customer relationship. The journey is not an operational manual; it is a way of thinking about where to invest attention. The chapters that follow it in Part Two then convert that thinking into specific operating decisions.

THE BOTTOM LINE

Customer Success is **a proactive, outcome-focused function accountable for retention, expansion, and advocacy** in any business where customers can choose to leave. It is structurally distinct from Support, Account Management, and Professional Services, even when the boundaries look blurry. **Getting the definition right is the precondition for everything else**, including knowing whom to hire, what to measure, and what to hold the function accountable for.

CHAPTER 2

The Customer Journey as a Framework

A customer's relationship with your business is a sequence in time, not a set of departments. Most leadership teams understand this in principle and ignore it in practice, which is how money goes missing.

If you ask a finance team to draw the customer experience, they will draw a P&L. If you ask a product team to draw it, they will draw a feature roadmap. If you ask a sales team, they will draw a funnel that ends at signature. None of these are wrong, but none of them is what the customer actually experiences.

What the customer experiences is a journey: a sequence of moments that begins before they sign anything and continues for as long as they pay you money. Understanding that journey, and being honest about what happens at each step of it, is the work this chapter is about.

The journey is not an operational manual. It is a thinking framework, a way for a leadership team to see the business through the eyes of the customer rather than the org chart. Each stage of the journey involves multiple functions, has its own economics, and presents its own opportunities to either gain or lose money. The journey is most useful when it forces conversations that the functional view of the business would not produce.

The Five Stages at a Glance

For the purposes of this book, the customer journey is divided into five stages. There are longer versions of this framework (six stages, seven stages, sometimes more), and there is nothing magic about the number. Five works because it is the smallest number that still lets you talk about the journey usefully, and an executive needs something they can hold in their head.

1. **Pre-sale**
   Awareness, evaluation, purchase decision
2. **Onboarding**
   From signature to first measurable value
3. **Adoption**
   Routine, dependent use of the product
4. **Renewal & Expansion**
   Recommitment plus expanded spend
5. **Advocacy**
   Customers who actively sell for you

Each stage is a distinct economic moment. Different functions own different stages. Customer Success does not own all five, but is involved in four of them, and is the only function that thinks about all five together. That makes it the natural custodian of the journey as a whole, even when it does not own any individual step.

The next few pages walk through each stage in turn, with the same structure: what the stage is, who owns it, and where money is gained or lost. The point is not to list activities; it is to make the economic logic of each stage visible to a leadership team that may not have seen it laid out this way before.

Stage 1 — Pre-sale

Everything that happens before a contract is signed. Marketing generates awareness, sales pursues qualified opportunities, and the customer evaluates options. From the customer's perspective, this stage is about building enough confidence to part with money. From the business's perspective, it is the stage that costs the most per customer and produces no revenue.

Customer Success is usually not present in pre-sale. The most common exception is in enterprise sales, where prospective customers sometimes ask to speak to a CSM during the evaluation process to understand what they will be working with. This is more common than executives think, and the quality of those conversations affects close rates more than people expect.

1 Pre-sale

The stage that costs the most and earns nothing. Customer expectations get set here, in conversations Customer Success usually has no visibility into. The handoff at the end of this stage is the most consequential moment in the entire journey, and most companies do it badly.

WHO OWNS IT	WHERE MONEY IS LOST
Marketing and Sales. Customer Success occasionally consulted in enterprise deals.	Mis-set expectations during the sales process, signing customers who cannot succeed with the product, ignoring the cost of customers who churn within six months.

The handoff from sales to onboarding at the end of this stage is where customer expectations meet product reality for the first time. Companies that treat it as an administrative step (a CRM record being passed from one team to another) discover, six months later, that this is where most of their churn was set in motion.

Stage 2 — Onboarding

From signed contract to first measurable value. This is the stage Customer Success usually owns end to end, in collaboration with Professional Services where implementation work is involved. It is also the stage where the highest concentration of churn risk sits, and where the most leverage exists to change customer lifetime outcomes.

The defining metric for this stage is *Time to Value (TTV), the elapsed time between signature and the moment the customer demonstrably uses the product to achieve the outcome they bought it for*. Companies that obsess over TTV tend to have better retention several years later, because customers who experience value quickly form different beliefs about the product than customers who slog through a six-month onboarding before seeing anything useful.

2 Onboarding

The highest-leverage stage in the journey. Customer beliefs about the product are formed here, and those beliefs persist long after the activities of onboarding are over. A customer who experiences value in week two thinks about renewals very differently from a customer still struggling to log in three months later.

WHO OWNS IT	WHERE MONEY IS MADE
Customer Success, often with Professional Services for implementation, with Product as a stakeholder.	Customers who hit first value within target form durable habits, expand earlier, and refer others. Compounding effect throughout the lifetime.

Chapter 7 covers onboarding operationally: how to design the programme, what to measure, where to draw the line between high-touch and self-serve. For now, the strategic point is that onboarding is not a cost line. It is the most important investment in customer lifetime value that a business makes, and treating it like a cost to be optimised downward is a category error that catches up with companies eventually.

Stage 3 — Adoption

The longest stage of the journey and the one most often taken for granted. Onboarding is over; the customer is using the product. The question now is whether that use becomes routine, whether it deepens, and whether the product becomes part of how the customer's business operates.

Adoption is where the difference between "the customer uses our product" and "the customer depends on our product" gets settled. The first state produces a renewal, sometimes. The second state produces a renewal, expansion, advocacy, and a customer who is hard to dislodge by competitors. The work that turns the first into the second is largely Customer Success's, and it is the work most often under-resourced because it is the least visible.

3 Adoption

The stage where customers become either dependent or detached. The work is unglamorous: nudging usage, identifying friction, surfacing features the customer is not using, helping the customer's team build the product into their workflows. The economic payoff comes later, in renewals and expansion that look automatic but were earned here.

WHO OWNS IT

Customer Success, with Product as a critical partner for surfacing usage patterns and removing friction.

WHERE MONEY IS LOST

Customers whose usage plateaus, whose champions leave without being replaced, or who never integrate the product deeply enough to make switching costly.

A composite scenario worth keeping in mind: a customer signed eighteen months ago, completed onboarding cleanly, and uses three of the product's twelve features. They renew this year, probably at the same price. They will likely churn the year after, when a competitor offers something cheaper or when their internal champion leaves. By the time the warning signs appear, it will be too late to act. The work to prevent this outcome had to be done in adoption, twelve to eighteen months earlier.

Stage 4 — Renewal and Expansion

The stage where the work of the previous three pays off, or doesn't. Renewal is the customer's commitment to continue paying. Expansion is the customer's commitment to pay more, for additional seats, additional products, higher-tier plans, or expanded usage. They are different motions with different economics, but they happen in the same window of time and often involve the same conversations, which is why this book treats them as one stage in the journey framework.

Owners of this stage vary widely by company. In some businesses, Customer Success runs renewals and expansion together. In others, a dedicated Renewals team handles the commercial paperwork and Customer Success drives expansion. In still others, Sales takes over the moment expansion is on the table. None of these is wrong; the right answer depends on company stage and segment, and Chapter 10 walks through how to choose. What matters here is recognising that the renewal-and-expansion stage is where Customer Success's commercial impact becomes visible, even if Customer Success does not directly hold the pen.

4 Renewal & Expansion

The cash-out moment of the journey. A customer who reaches this stage in good health renews and expands; a customer who reaches it in poor health renews at best, churns at worst. The conversations look commercial but the outcomes were determined in onboarding and adoption.

WHO OWNS IT	WHERE MONEY IS MADE
Varies widely: Customer Success, Renewals teams, Sales, or some combination. Decision belongs to the leadership team.	Net revenue retention above 100 percent (existing customers paying more than they did last year) is one of the most valuable metrics in any subscription business.

Stage 5 — Advocacy

The stage most companies underestimate, and the only stage where customers actively work on the business's behalf. An advocate is a customer who refers other customers, writes public reviews, agrees to be a reference call, speaks on a webinar, or otherwise lends their credibility to the business at no cost. Advocacy is what turns a happy customer into a multiplier.

The economics of advocacy are easy to underestimate because they do not show up on the balance sheet as a separate line. They show up as a lower acquisition cost (referred customers cost a fraction of new logos), a higher close rate (prospects who have spoken to a reference convert at much higher rates), and a more durable brand position (a company whose customers speak well of it in public is harder to dislodge than one whose customers say nothing).

5 Advocacy

The stage that returns value to the top of the funnel. Advocates lower acquisition costs, raise close rates, and shorten sales cycles. Most companies neither measure nor cultivate advocacy deliberately, and as a result they get less of it than they could, despite having satisfied customers who would happily do the work if asked.

WHO OWNS IT	WHERE MONEY IS MADE
Customer Success and Marketing in partnership. Sales is the primary beneficiary but is rarely involved in cultivating it.	Lower acquisition cost on referred customers, higher close rates on deals with reference calls, durable competitive position from public credibility.

A mature advocacy programme is one of the few places in a business where the marginal cost of revenue can be near zero. It is also one of the most chronically underinvested.

The Ideal Journey vs. the Actual Journey

Every leadership team should own two journey maps, not one.

The **ideal journey** is how the customer is supposed to move through the five stages. Pre-sale produces qualified, well-fit customers. Onboarding delivers first value within target. Adoption deepens use. Renewal and expansion follow naturally. Advocacy emerges and feeds back into pre-sale. Each stage hands cleanly to the next. This is the journey the company designs and the journey marketing materials describe.

The **actual journey** is what real customers experience. A meaningful percentage of pre-sale customers are not actually qualified, but get signed anyway because the quarter is closing. Onboarding stretches twice as long as planned because Professional Services is overloaded. Adoption stalls at thirty per cent feature usage because the original champion left. Renewal happens because no competitor surfaced in time, not because the customer is enthusiastic. Advocacy is essentially nonexistent because no one asks.

Most leadership teams have a clear picture of the ideal journey and a hazy picture of the actual journey. This asymmetry is the source of an enormous amount of organisational friction. Strategy meetings discuss the ideal journey while the customer is living the actual one, and the gap between them is where the company's most expensive problems live.

The work the gap analysis does

The most useful diagnostic a leadership team can run is to map both journeys side by side and ask, at each stage, where the gap is widest. The widest gap is where to invest. Not the worst overall outcome, not the loudest complaint, but the stage where the ideal and the actual diverge most. Closing one such gap produces more value than improving any stage that is already performing close to design.

Where Most Companies Misallocate Attention

When you ask leadership teams to rank the journey stages by importance, most rank pre-sale first by a wide margin. When you ask them to rank the stages by how much management attention and budget they receive, the answer matches: pre-sale first, by a wide margin.

This is rational in early-stage companies where survival depends on signing customers at all. It becomes increasingly irrational as the company matures. Once a meaningful base of recurring revenue exists, the economic returns from incremental investment in the later stages of the journey almost always exceed the returns from incremental investment in pre-sale, because the later stages compound and pre-sale does not.

A company can spend an additional pound on advertising and acquire an additional customer. The return on that pound is the gross margin of the customer over their lifetime, minus the acquisition cost. A company can spend the same pound on improving onboarding and increase the lifetime of every customer it acquires this year, plus every customer it acquires next year, plus every customer the year after that. The arithmetic of compound returns favours the later stages even when their individual effects look smaller.

Most companies know this and do not act on it. The reasons are organisational rather than analytical. Pre-sale is loud, visible, and produces a number every quarter that everyone in the company watches. Onboarding and adoption produce numbers eighteen to thirty-six months later, by which time the connection to the investment is hard to see. Leaders who reward what they can see end up with companies that systematically under-invest in the parts of the journey that matter most.

A useful diagnostic: take the percentage of your operating budget that goes to pre-sale (sales and marketing) versus the percentage that goes to the post-sale function. Compare it to the percentage of your revenue that comes from new logos versus existing customers. If the two ratios are wildly different, you are probably over-investing in pre-sale relative to where your revenue actually comes from.

Where Customer Success Fits

Customer Success is involved in four of the five stages, but does not own all four. The distinction matters because ownership confusion is one of the more durable sources of organisational friction in recurring-revenue businesses.

In pre-sale, Customer Success is occasionally consulted, particularly in enterprise sales. The function should weigh in on customer fit, especially when sales is considering signing a customer the CS team believes will not succeed. This is unpopular advice to give, but it is one of the most valuable services a CS function can provide a CEO: telling them which deals to walk away from.

In onboarding, Customer Success usually owns the relationship, in partnership with Professional Services where implementation work is involved. The handoff from sales to CS at the start of onboarding is one of the most-discussed and most-mishandled moments in the entire journey.

In adoption, Customer Success owns the work end to end. This is the function's home stage, the one where most CSM time is spent and where most of the CS team's compounding value is created.

In renewal and expansion, ownership varies. Customer Success may run the entire motion, may partner with a Renewals team, or may hand off to Sales for expansion conversations once an opportunity is qualified. Chapter 10 covers the operating-model choices in detail.

In advocacy, Customer Success partners with Marketing. The CS team has the customer relationships; Marketing has the channels and programmes. Neither owns advocacy alone, and the partnership only works when both functions treat it as a shared deliverable.

The Handoffs Between Stages

The journey is a sequence of stages, but in operational terms it is a sequence of handoffs. The customer is passed from one function to another four times across the five stages, and the quality of those handoffs determines whether the journey holds together or breaks.

Sales to Customer Success (end of pre-sale)

The most consequential handoff in the journey. The expectations set during the sales cycle are transferred (or not) to the team that will deliver against them. When this handoff is done badly, the CS team starts onboarding having to either lower expectations or fail to meet them. When it is done well, onboarding starts from a shared and realistic understanding of what the customer signed up for.

Professional Services to Customer Success (mid-onboarding)

In businesses with implementation work, customers are handed from PS to CS at some point during onboarding. The cleanest version of this handoff is well-defined and customer-visible. The messiest version is invisible to the customer, who calls their PS contact months later about a problem that should have been raised with their CSM.

Customer Success to Sales or Renewals (start of expansion)

When expansion is owned by a function other than CS, the handoff at the point of identified opportunity is delicate. Done well, it preserves the trust the CS relationship has built and channels it into a commercial conversation. Done badly, it makes the customer feel sold to by the people they thought were on their side.

Customer Success to Marketing (advocacy)

When CS identifies a strong advocate, the customer is introduced to Marketing's advocacy programme. This handoff is the least mishandled of the four, mainly because companies handle it so rarely. Companies that systematise it find that advocacy compounds in ways the others stages do not.

How to Use the Framework

The five-stage journey is not a process to follow. It is a lens to use, and its value comes from the conversations it enables rather than the activities it prescribes. Three of those conversations matter most.

The "where is our biggest gap" conversation

Once a leadership team has both the ideal and actual journeys mapped, the question becomes which stage has the widest gap, and what it would cost to close it. This is a strategic prioritisation conversation, not an operational one. It belongs at the executive team and the board, not in a CS team meeting.

The "what does this stage cost us" conversation

Each stage has a cost. Pre-sale has acquisition cost. Onboarding has implementation cost. Adoption has CSM cost. Renewal has the cost of the team that runs the motion. Advocacy has the cost of the programme. A leadership team that can express each stage's cost in roughly comparable terms can have a real conversation about return on investment. A team that cannot is essentially making journey decisions in the dark.

The "who else needs to be in this room" conversation

Because the journey crosses functions, no single function leader sees the whole thing. The CRO sees pre-sale and parts of renewal. The CS lead sees onboarding through advocacy. Product sees adoption signals. Marketing sees the top of the funnel and the bottom of advocacy. Strategic journey conversations need at least three of these voices in the room, or the conversation defaults to the perspective of whoever is loudest.

Closing the Chapter

The customer journey is the framework the rest of this book operates inside. Chapter 3 takes the journey and asks what investing in it is worth, with a business case a CEO might use to defend the post-sale budget. Chapter 4 gets specific about the economics, including the metrics that show up on board decks. Together, those two chapters close out Part One.

Part Two then walks through each stage of the journey in detail, in the order Customer Success encounters them. Onboarding is Chapter 7. Adoption and customer health is Chapter 8. Strategic account management, which is most of the work in adoption for high-value customers, is Chapter 9. Retention and renewals is Chapter 10. Expansion is Chapter 11. Together, those five chapters are the operational core of the book.

Before then, though, the question is whether any of this is worth investing in at all. The next chapter makes the case.

THE BOTTOM LINE

The customer journey is **a sequence of stages, not a set of functions**. Customer Success is involved in four of the five stages but owns only some of them. The journey is most valuable as a leadership framework for prioritising investment, especially when the ideal journey and the actual journey are mapped side by side. **The widest gap between them is where the highest return on improvement lives.**

CHAPTER 3

The Business Case for Customer Success

A CFO does not need to be convinced that customers matter. They need to be convinced that this particular line in the budget produces a return larger than the alternatives. That is the case this chapter makes.

The previous two chapters argued that Customer Success is a real discipline, structurally distinct from the functions it gets confused with, and organised around a customer journey that crosses the whole business. Both of those arguments are necessary. Neither is sufficient.

What turns a defensible idea into a funded function is a business case. A finance team that has to choose between hiring two more Account Executives and hiring two more Customer Success Managers will pick the AEs every time, unless the case for the CSMs is made in numbers the finance team can defend back to the CEO and the board.

This chapter makes that case. It does so in three steps. First, it identifies the three revenue levers Customer Success directly controls, and what each one is worth in plain economic terms. Second, it shows why the arithmetic of retention beats the arithmetic of acquisition, with numbers that are easy to defend. Third, it lays out the cost of not having a Customer Success function: not in slogans about churn, but in the specific places where money goes missing when no one is responsible for keeping it.

By the end of the chapter, a CEO should have a coherent argument for the budget. By the end of Chapter 4, they should also have the specific metrics to support it.

The Three Revenue Levers

Customer Success directly affects three lines on a business's revenue model. Each one is real, each one is measurable, and each one compounds over time in ways that make the function's economic contribution larger than it looks on a quarterly basis.

1 ***Retention***

Customers who keep paying you. The most basic lever and the one most often discussed. A one-point increase in gross retention does not sound dramatic, but applied across a customer base over several years it compounds into substantial revenue. *The compounding is the point: retention this year affects revenue every year thereafter.*

2 ***Expansion***

Existing customers paying you more than they did last year, through additional seats, additional products, higher-tier plans, or expanded usage. Expansion is the lever most commonly under-counted because it shows up alongside new logo revenue and is hard to separate without effort. *Mature subscription businesses derive the majority of their growth from expansion, not new logos.*

3 ***Advocacy***

Existing customers actively reducing your acquisition cost on future customers. Referrals close at higher rates and at lower cost. Reference calls increase win rates on competitive deals. Public credibility shortens sales cycles. *Advocacy does not show up as CS revenue, but it shows up as a reduction in the cost of revenue everywhere else.*

A function that materially affects all three of these is not a cost centre. Whether it shows up as one in your management accounts depends on how the accounts are structured, not on what the function is actually doing.

Why Retention Math Beats Acquisition Math

A widely-cited piece of customer economics, attributed to research by Bain & Company, holds that a five percentage point increase in retention produces a 25 to 95 per cent increase in profits. The exact range is debated, but the underlying logic is not, and it is worth walking through carefully because the intuition behind it is what gives Customer Success its economic standing.

The arithmetic comes down to two facts.

Retained customers cost less than new ones

A new customer arrives carrying the full weight of the cost the business paid to acquire them: advertising, sales commissions, marketing programmes, and the operational overhead of new account setup. A retained customer arrives carrying none of that. The cost of keeping a customer is roughly one-fifth the cost of acquiring an equivalent new one, and that ratio holds up across most B2B subscription businesses with allowances for industry variation.

Retained customers earn more over time

A customer in their third year of a subscription typically pays more than they did in their first year, through some combination of expanded usage, tier upgrades, and inflation-linked price increases. They also cost less to serve, because the support burden tapers as customers become more proficient with the product. The combination of higher revenue and lower cost per customer in later years is where the compounding really lives.

A simple worked example illustrates how quickly the gap opens up. Consider two companies with identical revenue this year, identical growth rates from new logo sales, and one small difference in retention.

SCENARIO — TWO COMPANIES, FIVE YEARS APART

Starting ARR (both companies)	**$10,000,000**
New logo ARR added each year	**$3,000,000**
Company A gross retention	**85%**
Company B gross retention	**93%**
Difference in ARR after 5 years	**~$6.5M**

The two companies look almost identical on quarterly reports. The acquisition engines are running at the same speed. The only difference is the rate at which customers are leaving the back of the funnel. Five years later, Company B is substantially larger, more profitable, and more attractive to investors, despite being run by a similar team with a similar product and a similar sales motion.

This is the argument for Customer Success in its plainest form. The function is paid to influence the retention line on a model like the one above. Even modest improvements compound into outcomes the rest of the business cannot match through any amount of acquisition spending. *ARR, or Annual Recurring Revenue, is the standard way subscription businesses measure the run-rate value of their committed customer base.*

The example also explains why investors care about retention metrics so much. Two companies with the same growth rate but different retention have radically different long-term valuations, because one of them is compounding and the other is treadmilling.

The Quieter Economics of Expansion

Retention is the lever most often discussed, but expansion is the one most often underestimated. In a mature subscription business, more revenue typically comes from existing customers paying more than from new customers signing for the first time. The companies whose financial profile investors most admire are the ones whose existing customer base produces more growth than their acquisition engine does.

A useful way to see the scale of this is to look at the metric investors actually focus on: *Net Revenue Retention (NRR), the percentage of recurring revenue retained from existing customers year over year, including expansion and minus churn and contraction*. An NRR of 100 per cent means existing customers are paying the same this year as last year. An NRR above 110 per cent means existing customers, as a group, are growing the business even before any new customers are signed. NRR above 120 per cent is the territory of the most prized subscription businesses on public markets.

A note on the difference between gross and net retention

Gross Revenue Retention (GRR) excludes expansion: it measures only what you kept from the customers you started with. Net Revenue Retention (NRR) includes expansion. The two together tell a more complete story than either alone. A business with strong NRR but weak GRR is growing the customers who stay but losing too many of them outright. The metrics are covered in detail in Chapter 4.

The economics of expansion are particularly favourable because the cost to deliver expansion revenue is dramatically lower than the cost to deliver new logo revenue. Selling additional seats to an existing customer requires no marketing spend, much less sales effort, and minimal onboarding. The gross margin on expansion revenue can exceed the gross margin on new logo revenue by a wide margin, which means each dollar of expansion is worth more to the business than each dollar of equivalent new logo revenue.

The Cost of Not Having Customer Success

Most arguments for Customer Success start by describing what the function does. A better argument starts by describing what happens in its absence. The cost of not having a Customer Success function is rarely paid in a single line on the P&L. It is paid in five different places, none of which is labelled "the cost of not having CS," and which together can easily exceed the cost of building the function many times over.

Higher gross churn than necessary

Without a function whose job is to prevent churn, customers leave for reasons that could have been addressed: confusion during onboarding, lack of adoption, the departure of an internal champion, a competitor making a move. Some churn is inevitable. The portion that is preventable, and that goes unprevented, is a direct revenue loss that compounds year over year.

Lower expansion than the customer base could support

Existing customers buy more when someone shows them why and when. Without a function actively identifying expansion opportunities, those opportunities are left to the customer to discover on their own, which means most of them are never discovered. The expansion that does happen is the customer-driven kind: someone in procurement calls and asks to upgrade. The far larger expansion that comes from proactive identification simply does not occur.

Higher acquisition costs to compensate

Every customer lost has to be replaced before the business can grow. A company with weak retention runs its acquisition engine harder than necessary to maintain the same top-line trajectory as a company with strong retention. This shows up as a higher Customer Acquisition Cost (CAC) ratio, a longer payback period on each new customer, and a leadership team that feels constantly behind. The cost is real but often hidden inside the sales and marketing budget rather than identified as the consequence it actually is.

No early warning system for product or pricing problems

A well-functioning Customer Success team is the early warning system for things going wrong in the business. Patterns in churn data, repeated objections during renewal conversations, common complaints in adoption: all of these point to product gaps, pricing miscalibrations, or competitive pressure that the rest of the organisation will only learn about months later, by which time the damage has been done at scale. Without a function looking for these patterns, the business is essentially flying blind on its existing customer base.

Lower customer-generated growth

Customers who would otherwise become advocates instead become neutral or quiet. The advocacy lever from earlier in this chapter does not get pulled, because no one is responsible for pulling it. The revenue this costs is invisible because it never shows up as a referral or a reference call, but the absence is structural. A business that could be acquiring 20 to 30 per cent of its new customers through advocacy instead acquires none of them through that channel, and pays for every customer the harder way.

Across these five places, the cost of not having Customer Success is rarely smaller than the cost of having it, and is usually larger by a significant margin. The case for the function is not, in the end, that it is nice to have. It is that the business is paying for its absence whether or not anyone has labelled the cost as such.

When Customer Success Pays for Itself

A reasonable next question for a CEO is: how long does it take? The answer depends on stage and starting point, but the patterns are consistent enough to be useful as a planning input.

For an early-stage company building its first Customer Success function (one or two CSMs hired against an existing customer base), payback is typically visible within six to twelve months, mainly through preventing churn that was already in motion. The function is small, the customer base is small, and a CSM can develop a personal relationship with most accounts. The returns come quickly because the most at-risk customers benefit immediately from being given an owner.

For a growth-stage company moving from an ad-hoc Customer Success arrangement to a proper team (segmented, with playbooks and tooling), payback typically takes 12 to 18 months. The shift requires investment in operations, technology, and the time it takes to build segmented motions for different parts of the customer base. The returns come more slowly, but they tend to be larger and more durable because they are not dependent on individual heroics.

For a mature company shifting Customer Success from a retention driver to a growth driver, payback takes 18 to 36 months. This shift involves changes to comp structure, organisational reporting, hiring profile, and the company's commercial model. It is the longest of the three payback periods because it is the most structurally ambitious, but it is also the one that produces the most differentiated long-term economics.

A diagnostic for CEOs reading this

If you cannot defend the Customer Success budget on a 12 to 36 month payback horizon, depending on company stage, you are either under-investing in the function (in which case the payback would be faster than you think) or you have the wrong function entirely (in which case no amount of additional investment will produce the payback you expect).

Red Flags: You Needed CS Yesterday

There are a handful of signs that strongly suggest a business should already have built a Customer Success function and has not. The presence of any one of these is a problem. The presence of two or more is an emergency.

Churn is rising and no one can say why

When a business loses customers and the leadership team cannot explain the pattern beyond "the competition is heating up" or "customers are tightening their budgets," the business has no eyes on its own customer base. Whatever is causing the churn was visible for months before it became measurable; no one was looking.

Renewals are forecast by guessing

A renewal forecast that is not built bottom-up from actual customer health data is essentially a hope. CFOs of recurring-revenue businesses without a CS function frequently discover this fact unpleasantly, mid-quarter, when renewals that were assumed to be safe turn out not to be.

Sales is selling to anyone who will buy

Without a function that has a point of view on customer fit, sales will tend to close any deal it can. Some of those deals are customers who cannot succeed with the product, and they show up later as churn, refund requests, or expensive escalations. The cost of these customers is paid by the post-sale organisation, but the absence of pushback at the point of sale is the upstream problem.

Expansion is happening, but you can't predict it

If existing customers are expanding but the business cannot predict who will expand or when, the expansion that is happening is essentially accidental. The far larger expansion that systematic identification would produce is being left on the table.

A leadership team noticing two or more of these patterns simultaneously is looking at a structural problem, not a temporary one. The fix is the same in every case: build the function that owns the answers to those questions.

Closing the Chapter

The business case for Customer Success rests on three claims, each of which can be defended in numbers a CFO will accept.

First, the function directly affects three revenue levers (retention, expansion, advocacy) that compound over time and produce larger effects in years three and five than in year one. Second, the arithmetic of retention beats the arithmetic of acquisition in any business with a meaningful base of recurring revenue, and the gap widens the longer the comparison runs. Third, the cost of not having a Customer Success function is paid in five places that together exceed the cost of building one, even when no individual cost is labelled as such.

Building the case is necessary but not sufficient. The case also has to be expressed in the specific metrics that a board, an investor, or a CFO will recognise and trust. That is what Chapter 4 is about: the economics of Customer Success expressed in the numbers that show up on dashboards and board decks. Together, Chapters 3 and 4 close out Part One of the book and form the strategic foundation the operational chapters of Part Two then build on.

THE BOTTOM LINE

Customer Success affects **retention, expansion, and advocacy**, three revenue levers that compound over time. The arithmetic favours retention over acquisition in any subscription business. **The cost of not having the function is paid whether or not anyone identifies it as such**, and is almost always larger than the cost of building the function in the first place.

CHAPTER 4

The Economics of Customer Success

A function that cannot be measured cannot be defended, and a function that is measured by the wrong things will be defended for the wrong reasons. This chapter is about the right things.

There are roughly thirty metrics that show up in articles about Customer Success. There are about six that show up on the slides a CFO actually reviews. This chapter is about those six, plus a small number of derived numbers that explain how they connect to each other.

The reason to be selective is not that the other metrics are wrong; many of them are useful for managing a CS team day-to-day. The reason is that an executive book has to distinguish between the numbers a function uses internally and the numbers a function reports externally. Confusing the two produces the worst kind of board deck: one full of activity metrics that suggest hard work and answer none of the questions the board actually has.

The metrics that matter at the executive level fall into three groups: retention metrics (what fraction of your customer base did you keep this year), expansion metrics (what fraction of additional revenue did your existing base produce), and efficiency metrics (what is the cost of producing each unit of customer revenue). Customer Lifetime Value sits across all three and is the most important single number a recurring-revenue business produces, which is where this chapter starts.

Customer Lifetime Value

Customer Lifetime Value (CLTV), also written as LTV, is the total revenue a business can expect to receive from a customer over the entire duration of their relationship, typically expressed as gross profit rather than gross revenue. It is the metric that translates the abstract idea of a long-term customer relationship into a number a leadership team can act on.

There are several ways to calculate CLTV, each suited to different stages of business maturity. The simplest version assumes a stable customer base and uses three inputs: average annual revenue per customer, average customer lifespan, and gross margin.

SIMPLE CLTV

Average Annual Revenue per Customer* × *Average Customer Lifespan (years)* × *Gross Margin

Suitable for early-stage businesses or stable customer bases. Loses accuracy when retention or pricing changes significantly year over year.

The simple version has the virtue of being easy to defend in a board meeting. It also has the weakness of assuming things stay constant, which they rarely do. For most businesses past the early stage, a more accurate model is needed.

Two More Ways to Calculate It

The churn-based version uses the annual churn rate directly. This is more useful for businesses that have been operating long enough to have a reliable churn number, because average customer lifespan is itself derived from churn.

CHURN-BASED CLTV

Average Annual Revenue per Customer × Gross Margin ÷ Annual Customer Churn Rate

For businesses with stable churn. A 10% annual churn rate implies an average customer lifespan of 10 years.

A third version, used primarily by larger businesses with cohort data, discounts future revenue using a cost of capital to produce a present-value CLTV. The added precision matters most for businesses with very long expected lifetimes, where the time value of distant future revenue becomes significant.

For most leadership teams, the churn-based version is the right working model. It is more accurate than the simple version, less complex than the discounted version, and it expresses the core insight: lifetime value is inversely proportional to churn, which is the variable Customer Success is in the strongest position to influence.

CLTV in Practice: A Worked Example

Consider a mid-market B2B software company with the following characteristics.

WORKED EXAMPLE — SIMPLE CLTV CALCULATION

Average annual revenue per customer	**$12,000**
Annual gross retention	**90%**
Implied annual churn rate	**10%**
Average customer lifespan	**10 years**
Gross margin	**75%**
Estimated CLTV per customer	**$90,000**

The calculation: $12,000 average revenue per customer per year, multiplied by 10 years average lifespan, multiplied by 75 per cent gross margin, equals $90,000 of estimated lifetime gross profit per customer.

Now consider what happens when Customer Success improves retention by three percentage points, from 90 per cent to 93 per cent. Customer lifespan extends from 10 years to roughly 14 years (1 divided by the new churn rate of 7 per cent). With everything else held constant, CLTV rises from $90,000 to roughly $126,000 per customer. **A three-point improvement in retention produced a 40 per cent increase in customer lifetime value.**

This is the arithmetic that makes Customer Success matter. The function operates on the variable in the CLTV equation that has the largest leverage. Improving retention is not a polite optimisation; it is structurally the highest-impact variable in the model.

Net Revenue Retention

If a board asks one question about a subscription business, it asks about Net Revenue Retention. NRR is the metric that captures, in a single number, how the business is performing on its existing customer base. It is the metric that most directly distinguishes a subscription business that is compounding from one that is treadmilling.

NET REVENUE RETENTION

$$(\textit{Starting ARR} + \textit{Expansion} - \textit{Contraction} - \textit{Churn}) \div \textit{Starting ARR}$$

Measured over a defined period (typically a year). Includes only revenue from the existing customer cohort, excluding new logos.

NRR is expressed as a percentage. An NRR of 100 per cent means the existing customer base produced the same revenue at the end of the year as at the start. An NRR of 110 per cent means existing customers grew the business by 10 per cent before any new customers were signed. An NRR below 100 per cent means the existing base shrank, and new logos are working partly to fill a hole.

NRR above 120 per cent is the territory of the most prized subscription businesses on public markets. It is the metric that distinguishes a company whose customer base is, by itself, a growth engine from a company whose customer base merely retains. The difference in valuation between the two is typically large.

Gross Revenue Retention

Gross Revenue Retention is the more conservative cousin of NRR. It uses the same formula but excludes expansion, so it can never exceed 100 per cent. GRR measures pure retention: of every dollar you had a year ago, how many dollars do you still have?

GROSS REVENUE RETENTION

$$(\textit{Starting ARR} - \textit{Contraction} - \textit{Churn}) \div \textit{Starting ARR}$$

Always 100% or lower. The cleanest signal of churn risk in the customer base, untangled from expansion noise.

Both metrics matter, and they tell different stories. NRR shows whether the customer base is a source of growth. GRR shows whether the customer base is leaking. Looking at one without the other can be misleading.

- ***When NRR is strong but GRR is weak***

A business with strong NRR but weak GRR is growing the customers who stay but losing too many of them outright. The expansion from the surviving customers is masking a structural retention problem. This pattern catches up with the business eventually, because expansion compounds more slowly than churn destroys. A board that sees the two metrics together can spot this pattern; a board that sees only NRR cannot.

What Good Looks Like

The benchmarks vary by segment and stage, but the broad shape of "good" is well-established enough to be useful as a planning input. The numbers below are not promises; they are the bands that informed investors use when looking at subscription businesses.

TIER	Gross Retention	Net Retention	Typical context
BEST IN CLASS	> 95%	> 120%	Mature enterprise SaaS, deep product moats
STRONG	90–95%	110–120%	Healthy growth-stage B2B SaaS
ACCEPTABLE	85–90%	100–110%	Early-stage, SMB-heavy, or new category
WARNING ZONE	< 85%	< 100%	Structural issues with fit, product, or pricing

A few notes on how to use this table. SMB-focused businesses systematically have lower retention than enterprise-focused businesses, because smaller customers churn for reasons (going out of business, changing leadership, budget reallocation) that have nothing to do with the product. A 90 per cent retention rate in an SMB business is much harder to achieve than the same rate in an enterprise business and should not be benchmarked against the same band.

Also: a business hitting "best in class" on both metrics is rare. Most strong subscription businesses are very good at one and acceptable at the other. The combination is what investors pay the highest premiums for.

The LTV:CAC Ratio

If CLTV is the most important single number a business produces, the LTV:CAC ratio is the most important comparison. It expresses the relationship between what a business is willing to pay to acquire a customer and what that customer is ultimately worth.

LTV TO CAC RATIO

Customer Lifetime Value ÷ Customer Acquisition Cost

CAC is fully-loaded: sales salaries, marketing spend, related overhead, divided by new customers acquired in the period.

The conventional benchmark is that an LTV:CAC ratio of 3:1 or higher is healthy. A ratio of 1:1 means the business spends a dollar to acquire a dollar of lifetime value, which is sustainable only in very specific circumstances. A ratio below 1:1 is destroying value with every customer acquired.

The ratio matters to Customer Success because the function controls the L (the numerator) more than any other team in the business. Sales and Marketing influence CAC; almost everything Customer Success does influences LTV. A function that improves retention by three percentage points produces an LTV gain that an equivalent improvement in CAC cannot match, because the L is multiplicative across the customer lifetime while the C is a single transaction.

An LTV:CAC ratio that is improving over time is one of the cleanest signals of a healthy subscription business. A ratio that is deteriorating, even if both LTV and CAC are individually improving, is a sign that one of them is improving faster than the other in the wrong direction. Most often, this happens when CAC creeps up unnoticed because the most efficient acquisition channels saturate, while LTV does not improve fast enough to compensate.

Cohort Analysis for Non-Finance Leaders

A cohort, in subscription economics, is a group of customers who arrived together. The January cohort is everyone who signed in January. The Q1 cohort is everyone who signed in the first quarter. Cohort analysis tracks how each group behaves over time.

The reason this matters is that aggregate metrics can hide important structural shifts. Suppose a business reports stable 90 per cent gross retention year over year. That number looks reassuring. But if customers acquired this year are retaining at 85 per cent while customers acquired three years ago are retaining at 95 per cent, the aggregate is misleading. The business is degrading in a way that will not show up in the headline metric until enough of the newer cohorts have aged into the dataset to drag the average down.

Cohort analysis surfaces this kind of degradation early. It also lets a leadership team see whether changes the business has made (a pricing change, a product change, a shift in the ideal customer profile) are improving or worsening the customers being acquired now compared with the customers acquired previously.

- ***The minimum cohort analysis a board deck should contain***

Gross retention by acquisition quarter or year, plotted on a single chart. Each cohort gets a line. The lines should slope downward over time as customers naturally churn. The cohorts at the top of the chart should be the oldest (most mature, most likely to be best customers). If the newer cohorts are not following a similar trajectory or are diverging, that is the most important question a board can ask about the business.

Cohort analysis is the responsibility of finance or revenue operations in most companies, but Customer Success should be a consumer of it, because the patterns it reveals frequently point to gaps in onboarding or adoption that the CS function is best positioned to address.

Translating CS Activity into P&L Impact

The metrics covered so far are the language Customer Success speaks. They are not, on their own, the language a CFO or board uses. The translation step matters because Customer Success leaders who can express their function's impact in P&L terms get the budgets they ask for, and the ones who cannot do not.

There are three translations worth practising.

From retention points to revenue dollars

A one-point improvement in gross retention applied to a $50 million ARR base produces $500,000 in retained revenue this year, compounding in subsequent years. Most CFOs will accept a CS function that costs less than 30 per cent of the recurring annual value of the retention improvements it produces, which sets an upper bound on the size of the function relative to the business it supports.

From expansion bookings to growth contribution

If Customer Success is credited with influencing or owning a portion of expansion bookings, those bookings should appear on the function's scorecard. A CS function that produces $5 million of expansion against a fully-loaded cost of $2 million is producing a 2.5:1 contribution ratio that is directly comparable to the productivity ratios used to evaluate Sales.

From advocacy to acquisition cost reduction

Customers acquired through referrals typically cost a fraction of customers acquired through paid channels. If 15 per cent of new logo revenue comes from referrals, and referrals cost 30 per cent of what a paid customer costs, the savings can be calculated and credited to the function that produces them.

Each of these translations involves judgement and attribution. None of them is perfect. All of them are better than the alternative, which is presenting Customer Success metrics in their native form and hoping the audience does the translation work themselves. They rarely do.

What to Put on the Board Deck

The single Customer Success slide on a board deck has limited space and considerable expectations. The version that works has four elements, no more.

The retention numbers

Gross Revenue Retention and Net Revenue Retention for the period, with a comparison to the prior period and a comparison to the company's targets. Two numbers, two comparisons. Not a dashboard.

The cohort chart

Gross retention plotted by cohort over time. This is the chart that shows whether the customer base is structurally improving, holding steady, or degrading. A board that sees this chart once a quarter develops an intuition for the business that no other single visual produces.

The expansion line

Expansion bookings for the period, by source. If Customer Success owns a portion of expansion, that portion should be clearly identified. If the function influences but does not own expansion, the influence attribution should be explicit so the conversation is honest.

The two risks worth flagging

Of all the customers currently in the base, the two issues most likely to affect retention or expansion in the next quarter. Specific accounts can be named or anonymised; the point is to identify the risks the board should be aware of before they become surprises.

A four-element slide built this way changes the texture of board conversations about Customer Success. The function stops being a service line the board asks polite questions about and becomes a revenue function the board interrogates with the same seriousness it applies to Sales. That shift is, by itself, one of the most valuable outcomes a CS leader can produce.

Closing Part One

The four chapters that make up Part One have done a specific job. They define what Customer Success is (Chapter 1), where it operates within the customer journey (Chapter 2), what its business case looks like in plain economic terms (Chapter 3), and how to express its impact in the metrics that a board and an investor recognise (Chapter 4).

Together, these chapters are designed to be sufficient for a senior leader who needs to make a strategic decision about Customer Success without taking on the responsibility of running it day to day. A CEO who reads only Part One should be able to walk into a board meeting, defend the budget, hold the function accountable for the right outcomes, and identify when the structure they have built is mismatched with the outcomes they expect.

Part Two then turns to execution. Eleven chapters follow, covering the operating-model decisions, hiring patterns, onboarding programmes, customer health systems, strategic account management, retention motions, expansion engines, technology choices, scaling transitions, industry variations, and the way artificial intelligence is reshaping the discipline. Part Two is written for the senior leaders translating the Part One strategy into a working organisation.

A reader can stop at the end of Part One and have everything they need to make the central decision. A reader who continues into Part Two will have everything they need to execute it.

THE BOTTOM LINE OF PART ONE

Customer Success is **a defensible business function with a measurable economic case**. Retention, expansion, and advocacy compound; CLTV is the metric they all roll up into; NRR and GRR are how a board reads them. **Building the function well is a leadership decision tied to specific numbers**, not an act of faith. Part Two turns that decision into an operating plan.

PART TWO

II

Building and Running the Function

Part One made the case for Customer Success as a strategic function. Part Two builds it. The eleven chapters that follow translate the strategy into the specific operating decisions a leadership team has to make: how to design the function, who to hire, how to onboard customers, how to run the retention and expansion motions, and how to scale the whole apparatus as the business grows.

CHAPTER 5

Designing Your Operating Model

Every Customer Success function reflects a set of choices the leadership team has made, deliberately or otherwise. The deliberate ones produce a function that works. The otherwise ones produce one that doesn't.

There are roughly four decisions that determine whether a Customer Success function will succeed or fail. What service model the function uses. How it segments the customer base. What it is on the hook for, expressed as a written charter. And where it sits in the organisation. The decisions are independent of each other, but they interact in ways that produce coherent functions when handled well, and incoherent ones when handled badly.

Most companies make these decisions implicitly. The first CSM is hired and reports wherever there's bandwidth on the org chart. A segmentation emerges from whichever customers shout loudest. A charter forms over time from the work the team accumulates. By the time anyone steps back and looks at the function, it is the sum of a hundred small accommodations rather than the product of a strategy.

This chapter is about making the decisions deliberately. It is short on activities and long on choices, because choices are what compound. A function built on coherent choices will absorb growth, turnover, and changes in strategy. A function built on accumulated accommodations will not.

The Four Touch Models

A touch model is the level of human contact a Customer Success function provides to a given customer. The decision is straightforward in principle and consequential in practice, because it determines the unit economics of the function and the kind of customer experience the function can sustain.

Four models exist, and most mature organisations use some combination of them across different customer segments. The differences come down to who initiates contact, how often, and through what channel.

MODEL	Contact pattern	Typical CSM ratio	Best fit
HIGH-TOUCH	Regular scheduled meetings, named CSM, deep relationship building.	10–30 accounts per CSM	Enterprise customers, high ACV, strategic relationships.
LOW-TOUCH	Lighter scheduled contact, shared CSM coverage, exception-based engagement.	50–150 accounts per CSM	Mid-market, mid ACV, customers who need oversight but not constant attention.
TECH-TOUCH	Automated communication, in-app guidance, no named CSM by default.	500+ accounts per CSM	SMB or self-serve customers, lower ACV, scale-driven economics.
HYBRID	Combination of the above, applied to different segments of the same customer base.	Varies by tier	Any business with a meaningful spread of customer sizes.

Most businesses past the earliest stage use the hybrid model, because customer bases are rarely homogeneous enough for a single approach to fit. The question is not "which model do we use" but "how do we mix them, and where do we draw the lines."

Choosing the Right Mix

The choice between touch models is driven mostly by Average Contract Value (ACV). *ACV is the annualised value of a customer contract, used as the standard way to compare deal sizes across segments.* The reason ACV matters so much is that it sets the ceiling on what the business can spend to serve a customer profitably.

A rough working rule: the fully-loaded cost of a Customer Success Manager, including salary, benefits, tooling, and overhead, is typically between $150,000 and $250,000 per year in mature subscription businesses. A CSM who covers twenty accounts at high-touch is therefore costing the business $7,500 to $12,500 per account per year just to provide CS coverage. If those accounts are paying $30,000 a year, the unit economics work. If they are paying $5,000 a year, they do not.

Working backward from that math produces a starting heuristic for which model fits which customer.

A starting heuristic

Customers paying more than $100K ARR can usually support a high-touch model. Customers paying $25K to $100K typically suit a low-touch model. Customers paying under $25K usually need to be served through tech-touch, with human contact reserved for exceptional cases. These are starting points, not rules. Strategic value, expansion potential, and reference value can all justify deviation from the basic ACV math.

The deeper trap most companies fall into is not choosing the wrong model for a given segment, but applying high-touch to too many customers because the function feels nicer that way. Every customer the high-touch team takes on lowers the overall efficiency of the function and crowds out attention from customers who genuinely need it. The discipline is in saying no to high-touch coverage of customers whose unit economics do not support it.

Segmenting the Customer Base

Touch models are applied to segments, which means segmentation is the prior question. How a business divides its customer base into groups determines almost everything downstream: how the CS team is structured, how playbooks are written, how risk is identified, how expansion is pursued.

Most segmentation models are some combination of three dimensions.

By revenue

The simplest cut, and the most common. Customers are tiered by ARR, with thresholds set to align with the touch-model boundaries from the previous section. The top tier gets high-touch, the middle tier gets low-touch, the bottom tier gets tech-touch. Easy to operationalise; tends to miss strategic value that does not correlate with revenue.

By strategic value

A second cut layered on top of revenue. Some customers are worth more to the business than their revenue suggests, because they validate the product in a sector, because they are reference accounts for the sales team, or because they have credible expansion potential well above their current spend. These customers warrant heavier coverage than their ACV alone would justify.

By complexity or risk

A third cut for customers whose situation requires attention regardless of their revenue or strategic value. A regulated industry deployment, a contractually complex relationship, or a customer in active churn risk all warrant a different operating model from a customer of equivalent size with none of those characteristics.

A working segmentation framework typically uses revenue as the base layer and then permits exceptions for strategic value or complexity. The exceptions matter because the exceptions are where most of the political pressure comes from. Without explicit criteria for promotion to a higher tier, the team ends up making case-by-case decisions that drift toward over-coverage.

Without Alienating the Bottom of the Pyramid

Customers in the lowest tier of a segmented operating model often feel forgotten. They paid for the same product as the customers above them, in their experience, and the visible difference in attention can read as a second-class status. Handled badly, this becomes a churn driver in exactly the segment where the business can least afford it, because the bottom tier is typically the largest by customer count.

The fix is not to spend more on the bottom tier; the economics do not support it. The fix is to make the tech-touch experience genuinely good, so that the absence of a named CSM is not the absence of value.

A good tech-touch motion has three properties. It anticipates the customer's needs through in-product guidance and proactive communication, so the customer does not have to ask for things. It provides credible escalation paths so that when a real problem occurs, the customer can reach a human quickly. And it is honest with the customer about the model: customers tolerate light-touch coverage when the company tells them clearly what to expect, and bristle when they discover by accident that no one is watching their account.

A composite example worth keeping in mind: a B2B subscription business in the analytics category serves its top 50 customers with named CSMs, its next 300 through a pooled low-touch team, and its remaining 2,000 customers through tech-touch. Retention in the bottom tier is, surprisingly, the highest in the customer base. The reason is that the company invested in a high-quality automated onboarding flow, regular product-update communications, and a visible "request a CSM" button that very few customers ever use. The bottom tier feels well-served not because they get attention but because they do not need it.

The Customer Success Charter

A charter is a one-page document that defines what the Customer Success function is on the hook for and, equally importantly, what it is not. The act of writing one forces conversations that most leadership teams avoid until the function is already in trouble.

The charter answers four questions in language a peer function leader would accept.

- ***What outcomes does Customer Success own?***

  Retention is almost always included. Expansion sometimes, with the level of ownership specified (owns, drives, supports). Customer satisfaction metrics sometimes. The point is to list outcomes the function will be measured on, in language that distinguishes ownership from influence.

- ***What activities does Customer Success do?***

  Onboarding, ongoing relationship management, adoption work, renewal motion, expansion identification, advocacy cultivation. Specifying these limits scope creep into adjacent functions.

- ***What activities does it explicitly not do?***

  The most useful part of the charter. Tier 1 support, billing collections, implementation services, product feature requests, contract negotiation. Being explicit about what is out of scope prevents the function from absorbing every customer-facing task that no one else wants.

- ***Who does it work with, and how?***

  The handoff protocols with Sales, Product, Support, and Professional Services. Vague handoffs are the source of most cross-functional friction; a charter that names them explicitly removes a lot of recurring arguments.

The discipline of writing the charter matters as much as the artefact itself. Charters tend to get written in a single sitting, by the function leader, after a series of conversations with peer leaders. Each conversation surfaces a disagreement about ownership; the charter records the resolution. A leadership team that produces a charter without those disagreements has either had them already, or is going to have them later under worse conditions.

Where Customer Success Should Report

The reporting line decision shapes how the function behaves more than any other organisational choice. Four common structures exist, each with a coherent argument behind it and a predictable failure mode.

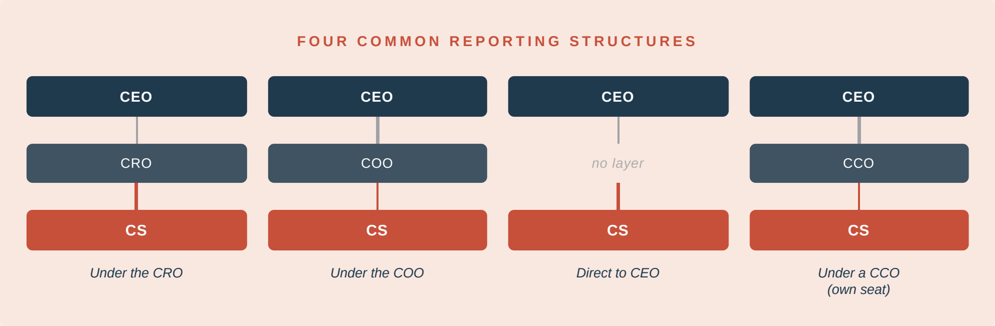

No structure is universally right; each is right for a particular combination of company stage, business model, and what the leadership team wants the function to optimise for. The next page covers each structure in turn, with the case for it and the failure mode to watch for.

1 **_Under the CRO_**

CS sits alongside Sales under a single revenue leader. *The case for:* retention and expansion are revenue outcomes; a single leader can balance new-logo and existing-customer growth. *The failure mode:* a CRO promoted from a Sales background often under-invests in CS and over-rewards new-logo growth, even when the math favours retention.

2 **_Under the COO_**

The case for: CS is operational and cross-functional, and the COO is the natural integrator of operational functions. *The failure mode:* the COO typically does not own a revenue number, and CS reporting into someone who does not have a revenue mandate tends to drift toward cost-centre behaviour.

3 **_Direct to the CEO_**

The case for: in early-stage companies, CS is too important to be filtered through anyone. The CEO needs unmediated visibility into customer outcomes. *The failure mode:* works at small scale; becomes the bottleneck once the company is past 50 or 60 people, at which point the CEO's bandwidth is the limit.

4 **_Under a Chief Customer Officer_**

The case for: CS, Support, and sometimes Professional Services consolidated under a single customer-facing executive who reports to the CEO. The strongest structure for mature businesses where the post-sale function is a primary growth lever. *The failure mode:* creating the role too early, before the function is large enough to justify a peer-level seat. Premature CCOs end up with thin teams and limited influence.

Operating-Model Mistakes in the First 18 Months

Most CS functions fail not because of the people in them but because of the operating model around them. Three mistakes account for most of the early failures, and all three are decisions made implicitly that should have been made explicitly.

Defaulting to high-touch for everyone

The first instinct of a new CS leader is to provide great service to every customer. This is admirable and unsustainable. The function ends up over-serving the bottom of the customer base, burning out its CSMs, and producing unit economics that do not support the function's continued growth. The discipline of saying "we do not provide named CSM coverage to customers below $X ARR" has to be a deliberate decision, and it has to be made early enough that the team's defaults form around it.

Skipping the charter

The team starts work without a clear scope document, and within six months has absorbed Tier 1 support, billing escalations, and implementation work that nobody else wanted. The team is busy, but the things they are busy with are not the things they were hired to do. The retention and expansion numbers, predictably, do not move. By the time anyone notices, the team's identity has shifted and is hard to reset.

Reporting line set by org-chart convenience

The first CSM is hired and reports to whichever executive has bandwidth, which is rarely the same as whoever should own CS strategically. Once the reporting line is set, it is politically expensive to change, and the function inherits the priorities of whatever function it landed under. Companies that set the reporting line deliberately, even in the early stage, save themselves a structural retrofit later.

A Pattern Behind All Three

All three mistakes share a common pattern: they are absences of decision rather than wrong decisions. The leadership team did not actively choose to default to high-touch, skip the charter, or set the reporting line by convenience. The decisions made themselves, in the absence of someone making them deliberately. That is the failure mode this chapter is written to prevent.

The fix in each case is the same. Name the decision, write it down, defend it. The act of writing forces clarity; the act of defending it surfaces the disagreements that would otherwise become organisational drift. Most CS functions that end up healthy share this pattern, and most that end up troubled share its absence.

A leadership team that finishes Chapter 5 with its four decisions made explicitly has done most of the structural work. The remaining chapters of Part Two operate inside that structure, but they do not replace it.

Closing the Chapter

The four decisions covered in this chapter are the structural ones. Together they determine whether the function the business builds is the function the business needs, or merely the function that emerged when no one was looking.

Chapter 6 takes the operating model and asks who staffs it. The hiring decisions a leadership team makes in the first eighteen months shape the function's culture for years afterward, and they involve choices (about seniority, about specialisation, about compensation structure) that are hard to undo once made. The two chapters are closely linked: an operating model without the right hires is a plan with no execution, and a hiring plan without the right operating model is execution without direction.

THE BOTTOM LINE

An operating model is **four decisions: touch model, segmentation, charter, and reporting line**. Most companies make these decisions implicitly, by default. The functions that work are the ones where the decisions were made deliberately, written down, and defended. **The cost of choosing badly is real; the cost of not choosing at all is larger**, because the function gets shaped by whoever shouts loudest rather than by what the business needs.

CHAPTER 6

Hiring and Building the Team

An operating model without the right people in it is a plan with no execution. The hiring decisions a company makes in the first eighteen months shape the function's culture for years afterward, and most of them are made with insufficient thought.

A common mistake in early Customer Success hiring is treating the role as fungible. The company decides it needs a CSM, posts a job description that could apply to any of three different kinds of person, and hires whoever interviews best. Six months later, the function is staffed by people doing the job they know how to do, which is rarely the job the business needed them to do.

This chapter is about reducing that mismatch. It walks through the four hiring decisions that matter most: what kind of profile to hire at each company stage, who the first CS hire should be, what ratios of customers and revenue to staff at, and how to structure compensation so the function does what the leadership team intends.

None of these decisions has a universal right answer. All of them have wrong answers that are easy to make and expensive to undo, which is what makes this chapter worth reading before a job description is written, not after.

The First CS Hire

The first Customer Success hire shapes the function more than any subsequent hire. Whoever takes the role sets the defaults: what gets prioritised, what good looks like, what the team values. Those defaults persist through people who join later, sometimes for years.

Three archetypes are common, and the choice between them is one of the more consequential early decisions a CEO makes.

1

The Player-Coach

A senior individual contributor who can carry a portfolio of accounts while also building the function around them. *The case for:* at small scale, the function needs both immediate output and structural design, and a player-coach delivers both. *The risk:* player-coaches often default to playing rather than coaching, because the playing is what they were good at. The function stalls structurally because the leader is too busy doing the work to design the system.

2

The Operator

A manager who has run a CS function before, hires aggressively, and builds the operating apparatus quickly. *The case for:* if the company has the customer base and the revenue to justify a small team within twelve months, an operator gets there faster than a player-coach can. *The risk:* hiring an operator into a business that does not yet have enough customers to manage produces an expensive hire with not enough work, and they leave within a year.

3 ***The Strategist***

A senior leader, often with prior VP or Chief Customer Officer experience, brought in to build the function at scale. *The case for:* for growth-stage companies that need to professionalise an existing CS team and elevate it organisationally, a strategist makes the case to the board and structures the function for the next phase. *The risk:* strategists hired too early have no team to lead, no operating problems to solve, and end up giving slide presentations about a function that doesn't exist yet.

These three archetypes are not equally common at every stage. The right choice depends almost entirely on where the company is in its lifecycle and how many customers the function needs to support within the first year. The next page works through the choice in detail.

Choosing the Right Archetype

The choice between archetypes is driven mostly by where the company is in its lifecycle and how many customers the function needs to support within the first year.

Seed and Series A stage

A player-coach is almost always the right answer. The customer base is small, the team will be one or two CSMs for the foreseeable future, and the company cannot afford a leader whose value is purely structural. The player-coach manages the customers, defines the playbooks, and builds the team as the customer base grows.

Series B and early growth stage

The choice gets harder, because the company is large enough that all three archetypes are plausible. A useful heuristic: if the CS function will be five people or fewer in twelve months, the player-coach remains the right call. If it will be ten or more, an operator becomes the better hire. The strategist remains the wrong answer at this stage in most cases.

Growth stage and beyond

The strategist becomes plausible, particularly when the company is professionalising the function for the first time after years of ad-hoc growth. The operator remains the right answer for companies whose CS function is large enough to need leadership but whose strategic architecture is already in place.

● *A diagnostic question worth asking*

What is this person going to be doing on a Tuesday morning, six months from now? If the answer is "running customer meetings," hire a player-coach. If the answer is "managing a team of CSMs and reviewing pipeline," hire an operator. If the answer is "presenting CS strategy to the board and shaping company-level decisions about retention," hire a strategist. Companies that cannot answer this question with specificity are not yet ready to make the hire.

The CSM Profile

Once the leadership question is settled, the next decision is what kind of individual contributor to hire into the CSM role itself. The temptation is to write a generic job description that asks for "customer-facing experience" and "strong communication skills" and call it sufficient. It is not.

A useful CSM job description specifies three things: the seniority required, the technical depth required, and the commercial orientation required. Each of these is a real trade-off, and getting one of them wrong produces a CSM who does the job they are wired for rather than the job the business needs.

Seniority

A senior CSM brings judgement, executive presence, and the ability to navigate hard conversations with customer leadership. A junior CSM brings energy, throughput, and a willingness to do the operational work that senior CSMs often resist. Most teams need a mix; the right mix depends on customer segment. Enterprise accounts almost always need senior CSMs. SMB accounts can be served well by junior CSMs operating within strong playbooks.

Technical depth

Technical CSMs understand the product deeply, can debug integration issues, and earn credibility with customer technical teams. They tend to be less commercially fluent and less comfortable with executive conversations. Non-technical CSMs are the inverse. The choice between them depends on what the customer needs from the relationship: a technical partner, a business advisor, or some combination. Hybrid roles exist, but they are harder to hire for and more expensive when found.

Commercial orientation

A commercially oriented CSM looks for expansion opportunities, frames conversations around business outcomes, and is comfortable with renewal pricing discussions. A non-commercial CSM focuses on adoption, customer satisfaction, and relationship health. Both are real CSM profiles. The choice depends on what the function is on the hook for: a CS team responsible for expansion needs commercially oriented CSMs, full stop. A CS team responsible only for retention can succeed with either.

A useful way to think about the three dimensions is that they are independent of each other. A junior CSM can be highly technical, or not. A senior CSM can be highly commercial, or not. The full combination of choices is what defines the role. Job descriptions that specify only one or two of these dimensions produce ambiguous hiring decisions and inconsistent teams.

● *The single most useful exercise before hiring*

Before writing the job description, write down what the CSM will be doing six months in. Three customers they have lost, what they did about it. Three expansion conversations they had, how they were structured. The technical issues they handled, the strategic conversations they led. The exercise produces clarity on which profile the role actually requires, which makes the resulting job description honest rather than aspirational.

Ratios That Work

How many customers should a single CSM cover? The question is asked often and answered confidently, despite being highly context-dependent. The honest answer is that the right ratio is whatever produces acceptable retention, adoption, and expansion outcomes for the cost involved. The ranges below are starting points for that calculation, not rules.

SEGMENT	Typical ACV	Customers per CSM	ARR per CSM
ENTERPRISE	> $250K	5–15	$2M–$5M
MID-MARKET	$50K–$250K	20–50	$2M–$5M
SMB	$10K–$50K	75–200	$2M–$4M
TECH-TOUCH	< $10K	500+	$1M–$3M

The pattern across the table is instructive. Customers-per-CSM varies enormously between segments, but ARR-per-CSM stays in a relatively narrow band ($2M to $5M for human-touch segments). This is not a coincidence: the band represents what an individual CSM can cover before the unit economics break down, regardless of how many individual customers that translates to.

Using the Ratios in Practice

The ratios are starting points, but they need to be adjusted for several factors that affect how much CSM time each customer actually requires.

Complexity of the product

A complex enterprise product with deep integrations and ongoing configuration work requires more CSM time per customer than a simple, self-explanatory product, even at the same ACV. The ratio shifts downward (fewer customers per CSM) as product complexity rises.

Stage of the customer

A customer in their first six months of a relationship requires more CSM attention than a customer in their third year. A team whose customer base skews newer than average should plan for lower ratios; a team whose base is mostly tenured can sustain higher ratios.

Expansion mandate

A CS team responsible for driving expansion has more work per account than a team responsible only for retention. The ratios in the previous table assume retention plus light expansion identification. Teams with active expansion quotas should plan for ratios 20 to 30 per cent lower.

A common error is to set ratios based on benchmark data without adjusting for these factors. The resulting team appears appropriately staffed on paper, becomes overstretched in practice, and produces results that disappoint everyone involved. The discipline is in starting with the benchmarks, then making the case for adjustments based on what is actually true about the business.

Compensation: The Question That Matters Most

More CS functions are broken by compensation design than by any other operational decision. The wrong comp plan produces the wrong behaviour, no matter what the job description says. The right comp plan aligns the function with the outcomes the business actually wants.

Three sub-decisions matter most: the split between base salary and variable compensation, what the variable compensation is tied to, and how that variable is calculated.

Base versus variable mix

Customer Success Manager compensation typically runs 70 to 85 per cent base, 15 to 30 per cent variable. The exact mix depends on what the variable is tied to, and on whether the role is considered a quota-carrying revenue role or a relationship-management role. A useful benchmark: CSMs in companies with strong CS functions typically have a higher base-to-variable ratio than Account Executives at the same companies, because the work is more relationship-oriented and less transactional.

The temptation in early-stage companies is to push variable higher to control fixed costs. This works for sales roles, where the variable links directly to revenue. It works less well for CS roles, where the variable links to outcomes that compound over time and are partly outside the CSM's direct control. Pushing variable too high produces CSMs who behave like salespeople, which is rarely what the business wanted when it built the function.

What the Variable Should Measure

The harder question is what the variable compensation is tied to. Most CS teams use some combination of three metrics, and the weighting reveals what the company actually values.

Gross retention or churn

The most defensible metric for the largest share of CS variable comp. Gross retention is directly within the CSM's influence, measurable on a stable timeframe, and tied to outcomes the business clearly cares about. The downside is that retention is influenced by factors outside the CSM's control (product issues, competitive pressure, customer-side leadership changes), and good CSMs can lose accounts that no reasonable amount of work could have saved.

Net retention or expansion

A larger share of variable compensation, tied to whether the customer base grew. Powerful when expansion is genuinely within the CSM's mandate. Dangerous when it is not: a CSM compensated on expansion they cannot meaningfully influence will become disengaged or, worse, will push for expansion at the cost of customer trust.

Customer satisfaction metrics

CSAT, NPS, or health scores can form a smaller portion of variable compensation, typically 10 to 20 per cent. These metrics work better as input signals than as primary comp drivers, because they are gameable and because the relationship between satisfaction scores and revenue outcomes is indirect at best.

• *A working comp structure*

For a CSM in a mid-market segment with expansion responsibilities, a workable comp plan might be 80 per cent base, 20 per cent variable, with the variable split 60 per cent gross retention, 30 per cent expansion bookings, 10 per cent customer health. This is a starting template, not a universal answer; the right mix depends on what the function is on the hook for.

Why Pure-Commission Models Break CS

Periodically, a CEO with a sales background decides that Customer Success should be run on a sales-style commission model. The argument is intuitive: CSMs influence revenue, so they should be compensated like the people who close revenue. The argument is also wrong, and the reasons are worth being explicit about.

CS outcomes compound; commission rewards transactions

A sales commission is tied to a discrete event: a deal closes. The CSM's work is tied to outcomes that emerge over months and years. A CSM who structures a customer for long-term success may not see the financial payoff for two years, by which time the commission structure has changed three times. The compounding nature of CS work is fundamentally incompatible with the transactional nature of commission.

Pure-commission CSMs stop being trusted

The most valuable thing a Customer Success function provides is a customer relationship based on trust. The customer believes the CSM is on their side. The moment the CSM's compensation is overwhelmingly tied to selling more product, the customer's belief erodes. The CSM's advice gets weighted differently. The relationship that took years to build is degraded within months. CSMs themselves notice this and either leave or become the kind of salespeople they were not hired to be.

Retention is not a sale

Sales commission works because closing a new logo is a discrete commercial act with a clear attribution model. Retaining an existing customer is a sustained operational act with diffuse attribution. Trying to commission-pay for retention introduces measurement problems that take more time to solve than the marginal motivation provides.

The fix is not to abandon variable compensation. It is to make the variable a meaningful but minority component of total compensation, and to tie it to outcomes that genuinely reflect the function's contribution. Pure-commission models for CS almost always need to be unwound, usually within eighteen months of being introduced.

Building the CS Leadership Bench

A CS function with one leader and a flat team of individual contributors works at small scale and breaks at scale. Most growth-stage CS organisations need a leadership bench that includes team leads or managers, a head of CS operations, and (in larger organisations) segment-specific leaders.

The question is when to start building it, and how to do so without prematurely top-heavy structure.

The first manager hire

Once the CS team is six to eight people, a first-level manager becomes necessary. The role is internal: managing CSMs, conducting reviews, running playbook execution. The right profile is someone with prior CS team management experience, ideally promoted from within if a strong candidate exists, externally hired if not.

The CS operations role

Often hired earlier than people expect. CS operations is the function that builds the tools, dashboards, processes, and analytics that the CS team relies on. In a team of ten or more CSMs, a dedicated operations person produces a multiplier effect on the rest of the team's output. Most teams that wait too long to make this hire spend the difference in CSM time wasted on operational work that should have been systematised.

Segment leaders

In larger organisations (twenty CSMs or more), splitting the team by segment becomes useful. An enterprise CS leader, a mid-market CS leader, and an SMB CS leader each run a function shaped to their segment's needs. This structure tends to emerge naturally once the team is large enough; the discipline is in resisting it before the team is large enough to need it.

Internal vs. External Hires

A recurring decision in CS hiring is whether to promote from within or hire externally. The trade-offs are real and apply differently at different levels.

For individual contributor roles

External hiring is usually the right answer for the first few CSMs, because the function has not yet generated the internal candidate pool that would justify promoting. As the function matures, internal moves from adjacent roles (support, account management, even sales engineering) become viable and often produce strong CSMs with deep product knowledge.

For first-level manager roles

A genuine toss-up. Internal promotion preserves institutional knowledge and signals career progression to the team. External hiring brings management experience that internal candidates may not yet have. The right call depends on whether a credible internal candidate exists, not on a general preference for one approach.

For senior leadership roles

External hiring is usually correct for the first VP or Chief Customer Officer role, particularly if the company is professionalising the function for the first time. Internal candidates can be strong, but the credibility a senior external hire brings to the board and the leadership team is often what the role requires. The exception is when the existing CS leader has clearly grown into the role and would be promoted into it externally anyway, in which case promotion is the cleaner outcome.

The single mistake to avoid is promoting an internal candidate because they are popular rather than because they are the right person for the job. Popular promotions that turn out badly damage the team more than external hires that turn out badly, because the unwinding is harder.

Hiring Mistakes Worth Naming

Three patterns produce most of the bad CS hires made by early-stage and growth-stage companies, and all three are avoidable with a small amount of forethought.

Hiring for the role you will need, not the role you have

A company at Series A hires a Head of CS who has built CS functions at companies ten times their size, expecting the senior hire to "level up" the team. The hire arrives, finds that the work is operational rather than strategic, and is bored within four months. The company has paid for a strategist and ended up with a frustrated senior IC. The fix is to hire for the next twelve months of work, not the next three years.

Hiring from within the same source pool

Companies that hire all their CSMs from the same background (former AEs, or former support managers, or former consultants) end up with a homogeneous team that has the same strengths and the same blind spots. The mix matters. A CS team with diverse backgrounds (sales, support, consulting, product) navigates more customer situations effectively than one with a single profile, however strong that profile is individually.

Hiring for skills that should be trained

Some skills (executive presence, judgement, commercial fluency) are hard to train and should be hired for. Other skills (product knowledge, playbook execution, internal tool proficiency) are easy to train and should not be hired for. Companies that filter candidates on trainable skills miss strong people whose ramp-up cost would have been trivial; companies that hire on hope for non-trainable skills end up with CSMs who never quite become effective.

Closing the Chapter

The four hiring decisions this chapter has covered (first hire archetype, CSM profile, ratios, compensation) determine whether the operating model from Chapter 5 produces results. A coherent operating model with the wrong people in it produces frustration. The right people in an incoherent operating model produces inconsistency. The combination of both is what produces a function that compounds.

Chapter 7 takes the function as built and asks what the team should actually do with its first ninety days with a new customer. Onboarding is the most concentrated set of activities in the customer journey and the highest-leverage period in the relationship. It is also the place where most of the structural decisions made in Chapters 5 and 6 first meet reality.

THE BOTTOM LINE

Hiring decisions **compound the same way customer relationships do**. The first hire shapes the function; the early CSM profile shapes the team; the comp structure shapes behaviour; the ratios shape what the team can sustain. **All four decisions are worth getting right before the job description is written**, because hiring corrections are slower and more expensive than the decisions that prevent the need for them.

CHAPTER 7

Onboarding and Time-to-Value

The beliefs a customer forms in their first ninety days outlast almost everything that happens afterward. Onboarding is where those beliefs are set, which makes it the highest-leverage period in the entire relationship.

Of all the stages in the customer journey, onboarding is the one where a fixed amount of effort produces the largest difference in lifetime outcome. A customer who reaches first value quickly forms a fundamentally different relationship with the product than a customer who struggles through a long, confusing setup before seeing anything useful. The first customer renews, expands, and refers. The second renews reluctantly if at all, and the work required to change their trajectory later vastly exceeds the work that would have set it right at the start.

This is why onboarding deserves a chapter of its own, and why it sits first among the operational chapters of Part Two. Everything the CS function does afterward is shaped by how well it did here.

The chapter covers the central metric of onboarding (time-to-value), how to design a programme that optimises for it, how to define and measure the moment a customer first experiences value, and what to do when onboarding goes wrong. It begins with the handoff that precedes onboarding, because onboarding that starts from a broken handoff is fighting uphill from the first day.

The Handoff That Precedes Onboarding

Onboarding does not begin with a kickoff call. It begins with the transfer of a customer from the sales team to the Customer Success team, and the quality of that transfer sets the ceiling on everything that follows.

The problem is structural. The salesperson who closed the deal knows things about the customer that never make it into the CRM: what the customer is actually trying to achieve, what they were promised in the final negotiation, which stakeholders matter, what nearly killed the deal. When that knowledge does not transfer, the CS team begins onboarding by reconstructing context the sales team already had, often by asking the customer questions the customer feels they already answered.

What a good handoff transfers

A good sales-to-CS handoff transfers four things. The customer's actual goals, in their own words, not the generic use case from the sales deck. The commitments made during the sales process, including any promises about timelines, features, or outcomes that the CS team will now be expected to deliver. The stakeholder map: who the champion is, who the budget holder is, who the skeptics are. And the deal's risk profile: was this an easy sale to an enthusiastic buyer, or a hard-fought win over internal resistance that could resurface at renewal.

The cost of a bad handoff

When the handoff fails, the customer experiences a discontinuity: the company that courted them attentively for months suddenly hands them to a stranger who does not know their situation. This is the first post-sale impression the customer forms, and it is a bad one. The damage is rarely fatal on its own, but it starts the relationship in a hole that onboarding then has to climb out of before it can make progress.

Why Time-to-Value Is the Metric That Matters

Time to Value (TTV) is the elapsed time between a customer signing and the moment they first experience the value they bought the product for. It is the single most predictive metric in onboarding, and one of the most under-monitored numbers in most subscription businesses.

The reason TTV matters so much is that it operates on customer psychology, not just customer logistics. A customer who experiences value quickly forms the belief that the product works and that the decision to buy it was correct. That belief is self-reinforcing: it increases their willingness to invest further effort, which produces more value, which deepens the belief. A customer who experiences a long delay before any value forms the opposite belief, and that belief is equally self-reinforcing in the wrong direction.

The practical consequence is that two customers who end up at the same level of product usage after a year can have completely different retention trajectories, depending on how fast they got there. The fast one believes in the product. The slow one is looking for the exit even while using it.

THE PATTERN — FAST VS. SLOW TIME-TO-VALUE

Customer A: first value in	**2 weeks**
Customer B: first value in	**3 months**
Usage after 12 months	**Identical**
Likelihood to renew and expand	**A >> B**

The two customers look the same on a usage dashboard a year in. They do not behave the same at renewal, because they hold different beliefs about the product formed during onboarding. TTV is how those beliefs get set.

Defining the First Value Moment

To optimise for time-to-value, a business first has to define what value means for its product, specifically enough to know when a customer has reached it. This sounds obvious and is routinely skipped. Many companies optimise their onboarding around completion of setup steps rather than achievement of value, which is not the same thing.

The first value moment is the point at which the customer experiences the core benefit the product exists to deliver. It is product-specific, and defining it requires honesty about what the customer actually came for.

Setup is not value

A customer who has connected their data sources, configured their settings, and invited their team has completed setup. They have not yet received value. Value comes when they get the insight, complete the workflow, close the deal, send the campaign, or whatever the product's actual purpose is. Companies that celebrate setup completion as if it were value delivery are measuring the wrong moment, and they optimise the wrong things as a result.

A test for the real first value moment

A useful test: if the customer stopped using the product the day after this moment, would they feel they had received something worth paying for? If yes, that is the first value moment. If no, the real moment is further along, and the onboarding programme should be designed to get the customer there, not merely to the point of technical readiness.

For some products the first value moment is obvious and early; for others it is subtle and requires sustained use to reach. The harder it is to reach, the more important onboarding design becomes, because the natural tendency of customers to drift away before reaching value is strongest in exactly those products.

Designing the Onboarding Programme

Onboarding design follows directly from the touch-model and segmentation decisions made in Chapter 5. Different segments warrant different onboarding intensity, for the same unit-economics reasons that drive the touch-model choice.

SEGMENT	Onboarding model	What it looks like
ENTERPRISE	High-touch, project-managed	Named onboarding specialist or CSM, kickoff, milestone plan, regular check-ins, often joint with Professional Services.
MID-MARKET	Guided, lighter touch	Structured programme with scheduled touchpoints, but more customer self-service between them.
SMB	Mostly self-serve	In-product guidance, automated email sequences, group webinars, human help on request.
TECH-TOUCH	Fully self-serve	In-product onboarding flows, contextual prompts, no scheduled human contact by default.

The common mistake is to apply a single onboarding model to all customers, usually the high-touch model, because it feels thorough. The result is an onboarding function that cannot scale and that spends disproportionate effort on customers whose contract value does not justify it.

The discipline is the same as with touch models: match the intensity of onboarding to the value of the customer, and make the self-serve experience genuinely good so that lighter-touch segments are not underserved. A good self-serve onboarding flow can produce better outcomes than a mediocre high-touch one, because consistency and thoughtful design often beat well-intentioned but improvised human attention. The segments that receive less human contact should not receive less thought.

What Good Onboarding Has in Common

Across all segments and intensity levels, well-designed onboarding programmes share a few characteristics. The touch level varies; these underlying properties do not.

A defined destination

Good onboarding knows where it is taking the customer: the first value moment, defined specifically and measured explicitly. Onboarding without a defined destination tends to default to "the customer is set up," which is the wrong target, as established earlier.

A sequence, not a checklist

Good onboarding moves the customer through a deliberate sequence designed to reach value as quickly as possible. The order matters: the fastest path to first value is rarely the same as the order in which features appear in the product menu. Onboarding designed as a feature checklist takes customers through everything; onboarding designed as a value sequence takes them to the point first, then expands.

Momentum management

Good onboarding pays attention to whether the customer is moving or stalling, and intervenes when momentum is lost. A customer who completes the first two steps enthusiastically and then goes quiet for ten days is at risk, and the onboarding system should notice and respond. Most churn that originates in onboarding comes from customers who stalled and were not noticed until it was too late.

Clear ownership

Good onboarding has a clear owner for each customer, even in self-serve models where the owner is a system rather than a person. The worst onboarding experiences happen in the gaps between functions, when the customer is technically owned by no one because Sales has moved on and CS has not yet picked up.

Onboarding Metrics That Matter

Onboarding can be measured in dozens of ways, most of which are not worth tracking at the executive level. Three metrics carry most of the signal, and a handful of commonly-tracked metrics carry almost none.

Time to first value

The headline metric, for all the reasons established earlier. Measured as the median (not the average, which is distorted by outliers) elapsed time from signature to first value moment. A business that tracks nothing else in onboarding should track this.

Activation rate

Activation rate is the percentage of customers who reach the first value moment at all, within a defined window. TTV tells you how fast the successful customers got there; activation rate tells you how many got there at all. A business with fast TTV but low activation is doing well by the customers who succeed and losing the rest. Both metrics are needed to see the whole picture.

Onboarding-specific satisfaction

A satisfaction measurement taken at the end of onboarding, distinct from overall product satisfaction. This captures whether the onboarding experience itself was good, which is a leading indicator of the customer's early relationship health. Useful as a signal, not as a primary target.

Metrics to be skeptical of

Number of onboarding calls completed, setup steps finished, training sessions attended, time spent in onboarding. These measure activity, not outcome. A customer can complete every onboarding call and attend every training session without reaching value, and a customer can reach value having skipped most of the programme. Activity metrics are useful for managing the onboarding team's workload; they are not measures of onboarding success.

When Onboarding Goes Wrong

Most onboarding failures are visible before they become terminal, if anyone is watching for the signs. Three warning patterns account for most of them.

The stall

The customer starts strongly, then stops. The most common pattern and the most recoverable, if caught early. A customer who completed initial setup and then went silent has usually hit a specific obstacle: a technical blocker, a competing priority, the departure of the person driving the implementation. The stall is a signal to reach out and remove the obstacle, not to wait and hope.

The slow drift

The customer never stops exactly, but never gains momentum either. Onboarding stretches from the planned six weeks to three months to indefinite. No single moment looks like failure, which is what makes the drift dangerous. The cumulative effect is a customer who has been onboarding for so long that they have formed the belief the product is hard to use before they ever reached value.

The premature completion

Onboarding is marked complete because the setup steps are done, but the customer has not actually reached value. The team moves on; the customer is left technically ready but not actually succeeding. This failure is often invisible in the metrics, because the onboarding system records the customer as successfully onboarded. It surfaces later as unexplained churn from customers who "completed onboarding" but never adopted.

Recovery Plays

When onboarding stalls or drifts, the response should be deliberate rather than improvised. A few recovery plays work reliably across most situations.

Re-establish the goal

Stalled onboarding often stems from the customer losing sight of why they bought the product. A conversation that returns to the customer's original goal, restates the value they were trying to achieve, and maps a short path to it can restart momentum that has been lost. This works because it reconnects the customer to their own motivation rather than to a generic onboarding plan.

Shrink the first win

When a customer is overwhelmed, the fix is usually to make the first value moment smaller and closer, not to push harder toward the original plan. A customer who is struggling to reach a large, complex first value moment can often be redirected toward a smaller one that they can reach this week. The small win rebuilds belief, which rebuilds momentum.

Find the new champion

When onboarding stalls because the person driving it has left or disengaged, no amount of process fixes the problem. The recovery play is to find and activate a new champion inside the customer organisation. This is harder than it sounds and is one of the clearest cases where a skilled human CSM outperforms any automated system.

The common thread across recovery plays is speed. A stall caught in week three is recoverable; the same stall discovered at the renewal conversation eleven months later is not. This is why momentum monitoring, covered earlier, is not an optional refinement but a core part of onboarding design. The recovery plays only work if the problem is detected while there is still time to act.

Onboarding Is an Investment, Not a Cost

A recurring tension in onboarding is the pressure to make it cheaper. Onboarding is visibly expensive: it consumes CSM and implementation time, it delays the point at which a customer becomes low-maintenance, and its costs are easy to measure while its benefits are deferred. The natural managerial instinct is to compress it.

This instinct is usually wrong, and the reason connects back to the central argument of the book. Onboarding is where customer lifetime value is set. A customer who is onboarded well is more likely to renew, more likely to expand, and more likely to refer, and all of those effects compound over the customer's lifetime. Cutting onboarding cost to save money this quarter trades a small, visible saving for a large, deferred loss.

The distinction worth drawing is between making onboarding cheaper and making it more efficient. Making it cheaper means doing less of it, which usually degrades outcomes. Making it more efficient means reaching the same first value moment faster and with less effort, which improves outcomes and reduces cost at the same time. The first is a false economy; the second is the actual goal.

A company that treats onboarding as a cost centre to be minimised will systematically underinvest in the highest-leverage stage of the customer journey. A company that treats it as the investment it is will spend deliberately on getting customers to value quickly, and will earn that spending back many times over across the customers' lifetimes.

Closing the Chapter

Onboarding is the highest-leverage stage in the customer journey because it is where the customer's beliefs about the product are formed, and those beliefs shape everything that follows. The metric that captures onboarding's effectiveness is time-to-value, and a business that optimises for nothing else in onboarding should optimise for that.

Good onboarding starts with a clean handoff from sales, aims at a clearly defined first value moment, moves the customer through a deliberate value sequence rather than a feature checklist, monitors momentum and intervenes when it stalls, and is measured by outcomes rather than activity. It is matched in intensity to the value of the customer, and it is treated as an investment rather than a cost.

Chapter 8 picks up where onboarding ends. Once a customer has reached first value, the work shifts to deepening their use and watching for the early signals of risk and opportunity. That work, customer health and adoption, is the longest stage of the journey and the one where the CS function spends most of its time. It is where the relationship either compounds or quietly begins to decay.

THE BOTTOM LINE

Onboarding sets the beliefs that govern the entire customer relationship, which makes **time-to-value the highest-leverage metric a CS function can optimise**. Good onboarding aims at a defined first value moment, not at setup completion, and is measured by whether customers reach value, not by how much activity occurred. **Cutting onboarding to save cost is a false economy**; making it more efficient is the real goal.

CHAPTER 8

Customer Health and Adoption

Most companies discover a customer is unhappy at the renewal conversation, which is the worst possible moment to find out. A health system exists to move that discovery forward by months, while there is still time to act.

Once a customer has reached first value, the relationship enters its longest phase. Onboarding is measured in weeks; adoption is measured in years. It is the stage where the customer either weaves the product into how their business operates, or keeps it at arm's length until a reason to leave eventually arrives. The difference between those two outcomes is rarely visible in a single moment. It accumulates quietly, which is exactly why it needs a system to detect it.

That system is the customer health score, and the work it enables is the shift from reacting to customer problems after they surface to anticipating them before they do. This chapter is about building that system and using it well, which turns out to be considerably harder than the dashboards most companies produce would suggest.

It covers what a health score actually is and what most companies get wrong about it, the three inputs every credible score needs, a practical plan for building a first version, the adoption work the score is meant to drive, and the organisational shift from reactive to proactive that the whole apparatus is in service of.

What a Health Score Actually Is

A customer health score is a single composite measure that estimates how likely a customer is to renew, expand, and stay. It exists to answer one question reliably: which customers need attention, and in what order. A health score that cannot answer that question is decoration.

This is where most companies go wrong. They build elaborate dashboards full of metrics, colour-code them, and call the result a health score. But a wall of green and red cells is not a health score; it is a data display. The defining property of a real health score is that it predicts something, and that someone acts differently because of it. If the score is never wrong, it is not measuring anything risky. If no one does anything different when a customer turns yellow, the score is a vanity metric.

> *The test of a health score is not how comprehensive it looks. It is whether a customer flagged at risk in March is meaningfully more likely to have churned by December than a customer flagged healthy. If the score has no predictive power, it has no value, no matter how sophisticated the dashboard.*

A good health score has three properties. It is predictive: customers it flags as at-risk do churn at higher rates than customers it flags as healthy. It is actionable: when a customer's score drops, there is a defined response, not just an alert. And it is honest: it shows uncomfortable truths rather than being tuned to look reassuring. A score engineered to keep the leadership team comfortable is worse than no score at all, because it creates false confidence.

The Three Inputs Every Health Score Needs

A credible health score draws on three distinct categories of signal. Each captures something the others miss, and a score built on only one of them is blind in predictable ways.

1 ***Usage***

What the customer actually does in the product: login frequency, feature adoption, depth of use, breadth across the team. *The strength:* objective, automatic, hard to fake. *The blind spot:* usage tells you what is happening, not why. A customer can use the product heavily and still be unhappy, or use it lightly and be perfectly satisfied because it does exactly the one thing they need.

2 ***Sentiment***

How the customer feels: survey responses, support ticket tone, the mood of recent interactions, escalation history. *The strength:* captures the why that usage data misses. *The blind spot:* sentiment is laggy and sparse. People answer surveys when they are very happy or very angry, and stay silent in between, which is where most customers actually live.

3 ***Relationship Strength***

How robust the human connection is: number of engaged contacts, seniority of the champion, whether the relationship is single-threaded or multi-threaded, recency of meaningful contact. *The strength:* predicts resilience to disruption. *The blind spot:* hardest of the three to quantify, and most dependent on the CSM honestly recording what they know.

The three inputs are complementary by design. Usage without sentiment misses the unhappy heavy user. Sentiment without relationship strength misses the satisfied customer whose only champion is about to leave. A score that combines all three sees the risks that any single signal would hide.

Combining the Inputs

Having three inputs raises the question of how to weight them. There is no universal answer, but there are sound principles and a common mistake to avoid.

The common mistake is weighting toward whichever input is easiest to measure, which is almost always usage. Usage data is automatic and abundant, so health scores tend to become elaborate usage scores with a thin layer of sentiment on top. This produces a score that looks rigorous and misses the most dangerous churn pattern: the customer using the product steadily right up until the day they announce they are leaving, because the decision was made by an executive who never logged in.

A more honest weighting reflects what actually predicts churn in a given business, which can only be learned by looking at past churn. The single most useful exercise in building a health score is to take the customers who churned in the past year and ask which signals, in hindsight, would have predicted it. Often the answer is relationship signals (the champion left, contact went quiet) rather than usage signals, which is the opposite of what most scores are weighted toward.

Start simple, then refine

A first health score should be simple enough to understand and explain: a small number of inputs, transparently weighted. Sophisticated weighting can come later, once the simple version has been tested against actual outcomes. A simple score that is used beats a sophisticated score that no one trusts, and trust comes from the score being understandable.

Building Your First Health Score

A company building its first health score does not need a data science team or specialised software. It needs a defensible first version that can be tested and improved. A practical ninety-day plan looks like this.

Days 1 to 30: learn from churn

Start with the customers who left in the past year. For each, reconstruct what happened: what their usage looked like, what their sentiment signals were, what the state of the relationship was. The goal is to identify the two or three signals that, in hindsight, most reliably preceded churn. These become the backbone of the first score. This step is skipped by most companies, which is why most first health scores are guesses dressed as analysis.

Days 31 to 60: build the simple version

Construct a score from the handful of signals identified, weighted by judgement informed by the churn analysis. Keep it simple enough to compute in a spreadsheet if necessary. Apply it to the current customer base and see which customers it flags. Then sanity-check the flags against what the CSMs already believe: if the score flags customers the team knows are fine and misses customers the team is worried about, the weighting is wrong and needs adjustment.

Days 61 to 90: operationalise and observe

Put the score in front of the team, define what happens when a customer changes colour, and start watching. The first version will be wrong in places. That is expected. The value of the ninety-day plan is that it produces a score grounded in real churn data and tested against real team knowledge, which is far ahead of where most companies start.

Adoption: From Logged In to Dependent

The health score tells you where customers stand. Adoption work is what moves them forward. The goal of adoption is to take a customer from using the product to depending on it, because dependence is what produces durable retention and expansion.

There is a meaningful difference between the two states. A customer who uses the product has a tool they could replace. A customer who depends on the product has woven it into how their business runs, such that replacing it would mean re-engineering their operations. The first customer renews if nothing better comes along. The second renews because leaving is genuinely difficult, and they know it.

What dependence looks like

Dependence shows up as breadth and depth. Breadth: more people across the customer organisation use the product, so it is embedded in more workflows and harder to remove. Depth: the product is used for more of the customer's process, so it sits closer to the core of how they operate. A customer using one feature, accessed by one person, is barely adopted no matter how often that person logs in. A customer using many features across many people is deeply adopted and structurally hard to dislodge.

The role of the CSM in driving adoption

Adoption rarely deepens on its own. Left alone, most customers settle into using the subset of the product they figured out first, and never expand beyond it. The CSM's job in the adoption phase is to notice where a customer's usage has plateaued and to show them the next layer of value: the feature they are not using that would solve a problem they have, the workflow they could move into the product, the team that could benefit from access. This is the unglamorous core of CS work, and it is where most of the function's compounding value is created.

Adoption Playbooks

Adoption work scales through playbooks: defined responses to common situations, so that the right action does not depend on a particular CSM happening to think of it. A few playbooks earn their place in almost every CS function.

The plateau playbook

Triggered when a customer's usage has flattened for a defined period. The response: identify the next feature or workflow that would deliver value given what the customer is already doing, and proactively guide them to it. The plateau is the most common adoption failure and the most addressable, because the customer is already engaged; they simply do not know what else is possible.

The single-threaded playbook

Triggered when a customer relationship depends on a single contact. The response: deliberately build relationships with additional people in the customer organisation, so the relationship survives the departure of any one person. Single-threading is one of the most common and most preventable causes of sudden churn.

The new-stakeholder playbook

Triggered when a new decision-maker enters the customer organisation, often a new executive who did not choose the product and feels no ownership of the decision. The response: proactively engage the new stakeholder, understand their priorities, and re-establish the product's value in their terms before they conduct the vendor review that new executives so often run.

Playbooks are not scripts. They are defined triggers paired with defined responses, executed with judgement. Their value is that they make the right response systematic rather than dependent on whether a particular CSM happened to notice the situation and know what to do.

The Proactive Shift

Everything in this chapter serves a single organisational shift: from reacting to customer problems after they surface to anticipating them before they do. This shift sounds obvious and is genuinely hard, because reactive work is urgent and proactive work is merely important.

A reactive CS function spends its days responding to what comes in: the support escalation, the angry email, the renewal that is suddenly at risk next week. The work is real and the team is busy, but the function is always a step behind, dealing with problems at the point where they have already become urgent and expensive to solve.

A proactive CS function uses the health score and the adoption playbooks to act before problems become urgent. It reaches the stalling customer in week three, not at the renewal. It engages the new executive before the vendor review, not after. It builds the second relationship thread before the champion leaves, not in a panic afterward. The work is less visibly urgent, which is precisely why it requires deliberate structure to sustain.

Why the shift is hard to sustain

Reactive work announces itself; proactive work does not. When a CSM spends a day on proactive outreach to healthy-looking accounts, nothing visibly breaks if they skip it, which makes proactive work the first thing sacrificed when the team is busy. Sustaining the proactive shift requires protecting the time for it structurally, because it will always lose to urgent reactive work in a fair fight.

Acting on Signals: Who Does What

A health score only creates value if a change in score reliably triggers a response. The response should be defined in advance, so that a customer turning yellow or red sets a known process in motion rather than depending on someone noticing and improvising.

STATUS	What it means	Defined response
GREEN	Healthy: using well, positive sentiment, solid relationship.	Maintain. Look for expansion opportunities. Cultivate as a potential advocate.
YELLOW	At risk: one or more signals declining, not yet critical.	CSM investigates within a defined window, identifies the cause, runs the relevant playbook.
RED	Critical: multiple negative signals, churn risk is real and near.	Escalate. CS leadership involved, save plan built, sometimes executive-to-executive contact.

The specifics matter less than the principle: every status has a defined owner and a defined action, with a defined timeframe. The most common failure is the yellow customer that everyone can see and no one owns, sitting in the dashboard turning slowly redder while each person assumes someone else is handling it. A health system without clear ownership of the response is an alarm with no one assigned to answer it.

Red-status customers in particular benefit from a defined escalation path, because the most serious churn risks usually exceed what a single CSM can resolve alone. A red status should pull in CS leadership, and sometimes the executive team, rather than leaving the CSM to carry a critical account by themselves. The cost of escalating a few accounts that turn out to recover is far lower than the cost of one major customer lost because no one above the CSM knew it was in danger.

The detailed mechanics of save plays and at-risk renewals belong to the retention discussion in Chapter 10. The point here is that the health system is what surfaces these situations early enough for those plays to have a chance of working. A save play launched in week three of a problem has options; the same play launched the week before renewal has almost none.

Closing the Chapter

A customer health score exists to answer one question reliably: which customers need attention, and in what order. It earns its place only if it predicts something and changes what someone does. Built on the three inputs of usage, sentiment, and relationship strength, grounded in an honest analysis of past churn, and paired with defined responses to changes in status, it becomes the instrument that makes proactive Customer Success possible.

The adoption work the score enables is what moves customers from using the product to depending on it, and dependence is what produces the durable retention and expansion that the rest of the book is concerned with. The whole apparatus serves the shift from reactive to proactive: catching the stall in week three rather than the renewal, engaging the new executive before the review rather than after.

Chapter 9 takes the highest-value customers, the ones where the relationship justifies the most deliberate attention, and examines the practices that strategic account management is built on: quarterly business reviews, executive sponsorship, multi-threading, and the hard conversations that the health system will, inevitably, surface.

THE BOTTOM LINE

A health score is only real if it **predicts churn and changes what someone does**. Built on usage, sentiment, and relationship strength, and grounded in honest analysis of past churn, it makes the proactive shift possible. **The goal of adoption is dependence, not just usage**, because a customer who depends on the product renews because leaving is hard, not merely because nothing better came along.

CHAPTER 9

Strategic Account Management

The handful of customers who represent most of your revenue do not want to be managed. They want a partner who understands their business well enough to be worth listening to. That is a higher bar, and it is the subject of this chapter.

Not all customers warrant the same depth of attention, a point established in the operating-model and segmentation discussions of Chapter 5. The customers at the top of that segmentation, the ones whose revenue and strategic value justify the most deliberate investment of CS time, require a distinct set of practices. Those practices are what this chapter is about.

Strategic account management is the discipline of running the most important customer relationships deliberately rather than reactively. It is not simply high-touch Customer Success with more meetings. It is a different posture: the CSM operates less as a service provider and more as an advisor whose value comes from understanding the customer's business and helping them succeed in it, of which the product is one part.

The chapter covers the quarterly business review and how to run one that is worth the time it consumes, executive sponsorship and when it earns its cost, the practice of multi-threading relationships so they survive the loss of a single champion, the hard conversations that strategic accounts inevitably require, and the trap that catches CS teams who confuse being helpful with being valuable.

The Quarterly Business Review

A Quarterly Business Review (QBR) is a periodic structured meeting between the CS team and a customer's stakeholders, intended to review progress, align on goals, and plan ahead. It is the central ritual of strategic account management, and it is done badly more often than it is done well.

The QBR done badly is a status update dressed as a strategic meeting. The CSM presents usage statistics, recaps support tickets, and walks through a deck the customer did not ask for. The customer sits politely through it, learns nothing, and leaves wondering why an hour was required. After a few of these, the customer's senior people stop attending, sending a junior delegate instead, and the QBR quietly loses its purpose.

The QBR done well is a different meeting entirely. It is about the customer's business, not the vendor's product. It opens with the customer's goals and asks how they are progressing against them. It surfaces problems the customer may not have noticed and opportunities they have not considered. It earns the attention of senior people because they leave it with something they did not have when they walked in.

ELEMENT	The QBR that fails	The QBR that works
SUBJECT	The vendor's product and usage.	The customer's business and goals.
DIRECTION	Vendor presents, customer listens.	Two-way; customer talks as much as the vendor.
OUTPUT	A deck the customer forgets.	Decisions, actions, and a plan both sides own.
ATTENDANCE	Declines over time as value fades.	Holds or rises as value becomes clear.

What a Good QBR Requires

Running a QBR that customers value is harder than running one that merely happens. Three things separate the two.

Preparation that the customer can feel

A good QBR is built on genuine understanding of the customer's situation, which takes preparation. The CSM arrives knowing what the customer is trying to achieve this quarter, what has changed in their business, and where the product is and is not helping. This preparation is visible to the customer; it is the difference between a meeting that feels tailored and one that feels templated. Customers can tell within five minutes which kind of meeting they are in.

A point of view, not just a report

The most valuable thing a CSM brings to a QBR is a point of view: a recommendation about what the customer should do next, grounded in what the CSM has seen work for similar customers. A CSM who only reports what the customer already knows adds little. A CSM who says "here is what I would do if I were you, and here is why" earns the standing that makes the rest of the relationship work.

The right people in the room

A QBR with only operational contacts is a check-in. A QBR that includes the customer's decision-makers is a strategic meeting. Getting senior people to attend requires giving them a reason, which loops back to the first two points: senior people attend meetings that are about their business and that produce a useful point of view. They skip meetings that recap usage statistics.

The frequency matters less than the name suggests. Quarterly is a convention, not a requirement. Some accounts warrant monthly strategic contact; others are well served by twice a year. The cadence should match the pace of the customer's business and the complexity of the relationship, not a calendar rule applied uniformly.

Executive Sponsorship

An executive sponsorship programme pairs a senior leader at the vendor with a senior leader at the customer, creating a relationship that sits above the day-to-day CSM connection. Done well, it is one of the most effective retention and expansion tools available for the largest accounts. Done indiscriminately, it is an expensive use of executive time that produces little.

The value of executive sponsorship comes from two things. It creates a peer relationship: a customer CFO takes a meeting with a vendor CFO that they would not take with a CSM, and the conversation operates at a level the CSM cannot reach. And it creates resilience: when something goes wrong, an existing executive relationship is worth more than any escalation process, because the trust was built before it was needed.

When it is worth the cost

Executive time is the scarcest resource in any company, so executive sponsorship should be reserved for the accounts where it genuinely moves the needle: the largest accounts by revenue, the most strategic by reference or expansion value, and the ones where a peer-level relationship would materially change the trajectory. A common mistake is to extend executive sponsorship too broadly, which dilutes the executive's attention across too many accounts and reduces the programme to a series of obligatory calls that neither side values.

● *The test for executive sponsorship*

Would a relationship between our executive and theirs change what happens to this account? For the top handful of accounts, the answer is usually yes, and the investment pays for itself in a single saved renewal or unlocked expansion. For the next tier down, the honest answer is usually no, and the executive's time is better spent elsewhere. The discipline is in restricting the programme to where it actually matters.

Multi-Threading: Surviving the Champion's Departure

The single most common cause of sudden, surprising churn in a strategic account is the departure of the champion: the person inside the customer organisation who chose the product, advocates for it, and holds the relationship together. When that person leaves, a relationship that looked healthy can collapse within a quarter, because it turned out to rest on one set of shoulders.

Multi-threading is the deliberate practice of building relationships with multiple people across the customer organisation, so that no single departure can unravel the account. It is one of the highest-return activities in strategic account management, and one of the most neglected, because it feels unnecessary precisely when the relationship is healthy and the champion is firmly in place.

Why single-threading happens

Single-threading is the path of least resistance. The champion is responsive, helpful, and easy to work with, so the CSM works through them for everything. The relationship feels strong because the one relationship that exists is strong. The fragility is invisible until the champion announces they are leaving, at which point it is too late to build the relationships that would have provided resilience.

What multi-threading requires

Building relationships beyond the champion takes deliberate effort against mild resistance. The champion may not see why the CSM needs to know their colleagues; the colleagues may not see why they should spend time with a vendor. The work involves finding legitimate reasons to engage other stakeholders: a QBR that includes the champion's boss, a use-case expansion that brings in a different team, an executive sponsorship relationship at a higher level. The goal is for the account to have several independent reasons to keep the product, held by several different people.

The Hard Conversations

Strategic account management inevitably involves conversations that are uncomfortable. How a CSM handles them often determines whether the relationship deepens or frays. Four come up repeatedly.

The price increase

Telling a valued customer their price is going up is difficult precisely because the relationship is good. The conversation goes better when it is framed in terms of value rather than cost: what the customer is getting, how that has grown, and why the new price reflects it. It goes worse when the increase arrives as a line item in a renewal quote with no conversation, which customers experience as a betrayal of the relationship they thought they had.

The contract reduction

When a customer wants to reduce their spend, the instinct is to resist. Often the better move is to understand why first. A reduction driven by genuine over-purchasing is different from one driven by dissatisfaction, and the responses differ accordingly. A customer allowed to right-size their contract sometimes stays for years; a customer forced to keep paying for what they do not use churns entirely at the next opportunity.

The escalation

When something has gone badly wrong, the customer is angry, and the CSM is the face of the vendor. The conversation goes better when the CSM owns the problem rather than deflecting it, communicates clearly about what is being done, and follows through visibly. Escalations handled well can paradoxically strengthen relationships, because the customer learns that the vendor performs under pressure. Escalations handled badly confirm the customer's worst fears.

The churn-risk conversation

The hardest conversation of all is with a customer who is seriously considering leaving. The instinct is to defend the product, list its features, and argue the customer out of their decision. This rarely works, because a customer who has reached the point of considering departure has usually stopped listening to product arguments.

The conversation that has a chance is one that starts by understanding, genuinely, why the customer is considering leaving. Sometimes the reason is fixable and no one knew about it. Sometimes the reason is a misunderstanding that can be corrected. Sometimes the reason is real and unfixable, in which case the honest move is to acknowledge it rather than argue. A CSM who listens first, and who is willing to admit when the customer has a legitimate grievance, retains more accounts than one who launches into a defence.

Across all of these conversations runs a common thread: the relationship is strengthened by honesty under pressure and damaged by evasion. The CSM who delivers hard news directly, owns problems rather than deflecting them, and listens before responding builds the kind of trust that survives the next difficulty. The one who avoids hard conversations until they become crises finds that the avoidance itself becomes the reason the relationship fails.

> *The customers who trust you most are not the ones who never heard bad news from you. They are the ones who heard it from you first, delivered honestly, before they heard it from anyone else or discovered it themselves.*

The Trusted-Advisor Trap

Customer Success teams are told to become trusted advisors, and the advice is sound. But there is a trap inside it, and strategic account management is where teams most often fall into it. The trap is confusing being helpful with being valuable, and over-serving accounts to the point where the service destroys the economics of the relationship.

A trusted-advisor relationship is supposed to be efficient: the customer trusts the CSM's judgement, so decisions happen faster and with less friction. The trap is the opposite, where the CSM becomes so deeply embedded in the customer's operations that they are effectively providing free consulting, absorbing hours of work that the contract does not pay for, in the name of being a good partner.

How the trap forms

It forms gradually and from good intentions. The CSM helps with something outside the product's scope because the customer asked and it seemed churlish to refuse. The help is appreciated, so more is requested. Over months, the CSM is spending a third of their time on a single account, doing work that is not really CS work, for a customer whose contract value does not remotely justify the investment. The customer is delighted; the unit economics are quietly underwater.

Why over-serving destroys margin

The cost is rarely visible in any single interaction, which is what makes it dangerous. Each individual act of help looks reasonable. The cumulative effect is a CSM whose effective cost per account has ballooned, who is unavailable to other accounts that need attention, and who has trained one customer to expect a level of service that cannot be sustained or scaled. When the service is eventually pulled back, the customer experiences it as a downgrade, and the relationship that the over-service was meant to protect is damaged anyway.

Staying Valuable Without Over-Serving

The way out of the trap is to distinguish clearly between value that deepens the relationship and service that merely consumes it. Helping a customer get more from the product is value: it increases their dependence and their likelihood to renew and expand. Doing the customer's job for them, outside the product's scope, is service that consumes margin without building durable value.

The discipline is to be generous with the first and disciplined about the second. A strong CSM gives freely of judgement, insight, and product expertise, because those things compound the relationship and cost little to provide. The same CSM is careful about giving away hours of operational labour that the contract does not cover, because that does not compound; it simply drains. Knowing the difference, and holding the line on it even when a valued customer pushes, is part of what separates a strategic account manager from an over-extended one.

Chapter 10 turns to the moment all of this work is ultimately measured against: the renewal. Retention, the renewal motion, churn forensics, and the save plays that the health system surfaces are where the relationship either converts into continued revenue or does not. Strategic account management is, in large part, the work that makes those renewals a formality rather than a fight.

THE BOTTOM LINE

Strategic account management runs the most important relationships **deliberately rather than reactively**. The QBR is about the customer's business, not the vendor's product; executive sponsorship is reserved for where it moves the needle; multi-threading protects against the champion's departure. **The trap to avoid is confusing being helpful with being valuable**, and over-serving accounts until the service destroys the economics it was meant to protect.

CHAPTER 10

Retention, Renewals, and Churn

The renewal is where every previous decision is settled. Everything the function has done in onboarding, adoption, and account management either holds together at this moment, or doesn't.

Renewals are the most consequential events in a subscription business. They are also the most often mishandled, because most companies treat them as administrative rather than commercial. The contract comes up for renewal; the customer either signs again or doesn't; the result shows up in the retention metrics next quarter. The framing makes the outcome look passive, but it is anything but. A renewal is a decision the customer makes about whether to keep paying, influenced by a year of accumulated experience, and the work that shapes that decision happens long before the contract is signed.

This chapter is about that work, and about the renewal motion itself. It covers the three structurally different kinds of churn and why each requires a different response, the question of who owns the renewal motion (with the trade-offs that come with each answer), the forensic work of understanding why customers actually leave, the save plays that have a chance of working and the ones that merely delay the inevitable, and finally how to build a renewal forecast that a CFO will trust.

The thread running through all of it is that retention is not preserved by trying harder at the renewal. It is preserved by what happens in the eleven months that precede it.

The Three Types of Churn

Churn is often discussed as a single number, but the number aggregates three structurally different things, each with its own causes and its own responses. Conflating them produces churn-reduction efforts that are misdirected at best and counterproductive at worst.

1

Involuntary Churn

Customers who leave for reasons unrelated to the product: the customer's business closed, the customer was acquired and consolidated to a different vendor, leadership changed and rationalised the stack. *The response:* minimise where possible, but recognise that some involuntary churn is genuinely outside the function's control. The mistake is to spend Customer Success effort fighting churn that no amount of work could have prevented.

2

Voluntary Churn

Customers who actively choose to leave: they did not get value, they found a better alternative, they decided the product was not worth the cost. *The response:* this is the churn the function exists to prevent. Most of the work in the previous chapters was about reducing voluntary churn before it becomes inevitable.

3

Downgrade Churn

Customers who stay but reduce their spend: fewer seats, lower tier, less usage. Not visible in logo retention but a direct hit to net revenue retention. *The response:* often the result of poor adoption work earlier in the year, where the customer never expanded into the product enough to justify the original contract size at renewal.

The distinction matters because a function that does not separate the three ends up trying to apply voluntary-churn tactics to involuntary-churn situations and vice versa. The leadership team should know what the headline churn number is and what its composition is, because the composition is what indicates where to focus the work.

A useful exercise: take last year's churn rate and break it down by these three categories. The resulting picture often reveals that the headline churn number, while accurate, was telling a much less useful story than the composition would have. A function focused on the wrong category of churn does work that does not move the headline number, then concludes that the work was ineffective, when the real issue was that the effort was aimed in the wrong direction.

Who Owns the Renewal

Three models exist for who runs the renewal motion. Each has a coherent argument behind it and a characteristic failure mode. The decision matters because the function that owns the renewal optimises for it, and the optimisations are not all the same.

MODEL	The case for it	The failure mode
CS-OWNED	CS has the relationship and the context; the renewal is a natural extension of the year's work.	CSMs default to relationship preservation over commercial outcomes; price increases and difficult terms get softened.
SALES-OWNED	Renewals are revenue events; sales is built to close them and to push pricing.	A salesperson treating a renewal like a new deal damages relationships and surfaces objections that did not need to be surfaced.
RENEWALS TEAM	A dedicated team optimises for the commercial motion without the conflicts of either CS or Sales.	Adds a handoff and a third party to a relationship the customer already understands; expensive to build.

The right answer depends on stage and segment. Early-stage companies almost always have CS own renewals. Growth-stage companies often experiment with the second or third model as their CS team scales. Mature enterprise businesses frequently end up with a hybrid: a dedicated renewals team for transactional renewals, with CS-led renewals for strategic accounts.

What Changes With Each Model

Beyond the structural choice, the model dictates what the renewal motion actually looks like in practice. A CS team that owns renewals runs them very differently from a renewals team that owns them, and the differences shape what the customer experiences.

When CS owns it

The renewal is woven into the ongoing relationship rather than treated as a discrete event. The conversation about next year's contract begins months before the contract date, embedded in QBRs and strategic discussions. By the time the renewal arrives, both sides have already discussed terms. The risk is that CSMs avoid the commercial conversations until too late, then accept whatever terms the customer asks for in order to preserve the relationship.

When Sales owns it

The renewal is treated as a commercial deal, with the same rigor as a new opportunity: pricing analysis, negotiation strategy, executive support. The motion is sharper commercially but can feel cold to the customer, who expected the renewal to reflect a year of working together. The risk is that the salesperson, optimising for the immediate deal, makes pricing or term decisions that win this renewal at the cost of the next one.

When a renewals team owns it

A specialist team focuses on the commercial mechanics: timing, pricing, paperwork, negotiation. CS continues the relationship work in parallel, and the two functions coordinate. Done well, this combines the strengths of both. Done badly, the customer experiences a confusing handoff at the worst possible moment, when a new face appears asking about contracts while the familiar face appears asking about value.

No model is universally right. What matters is choosing one deliberately, defining the handoffs clearly, and accepting the failure mode that comes with the choice. A leadership team that ducks the decision ends up with all three models partially in place, which produces the worst version of each.

Churn Forensics: Why Customers Actually Leave

A company that does not know why its customers churn cannot prevent the next ones from churning for the same reasons. This sounds obvious, but most companies' answer to the question of why their customers leave is wrong, because the data they rely on is the wrong data.

The exit survey is the most common source of churn data and the least reliable. Customers who fill it out tend to give the easiest answer: "budget," "consolidation," "no longer needed." These answers are sometimes true and frequently a polite cover for the real reasons, which are harder to say to the vendor: "we did not get value," "we never figured out how to use it," "we lost the person who cared about it." The exit survey produces a clean dataset that systematically understates voluntary churn.

Where the real reasons are

The honest signals are usually visible months before the customer announces departure. Declining usage by the people who originally championed the product. Engagement dropping off in the months before the renewal date. Support tickets that hint at unmet expectations. A CSM's gut sense that something has shifted, recorded if the CSM remembered to record it. Going back through these signals after the fact produces a far more accurate picture of why the customer left than any survey.

The five-customer exercise

A useful discipline: each quarter, take the five most painful churns and reconstruct what happened to each, in detail. Not a generic root-cause analysis, but a specific narrative for each customer. What did their first six months look like? When did the relationship change, and what changed? Who was involved, who left, who came in? This exercise produces actionable patterns that aggregated survey data never will.

Common Patterns the Forensics Reveal

When companies do this work honestly, a handful of patterns appear over and over. Recognising them in the live customer base is one of the highest-leverage uses of CS time.

The customer who never adopted

Reached the end of the first year having barely used the product. The reasons trace back to onboarding: a long time-to-value, an unsupported transition from the original buyer to the day-to-day users, a champion who lost interest. The renewal looks bad, but the loss happened in months one through six, not at the renewal itself.

The customer who lost their champion

Used the product well for two years and then collapsed in the third. The trigger is almost always a personnel change: the executive who chose the product left, and the replacement either ran a vendor review or simply decided to consolidate. Multi-threading work done in years one and two is what prevents this; the absence of it is what causes the loss.

The customer who outgrew you

A real outcome, sometimes. The customer's needs evolved beyond what the product can do, and they moved to something more capable. Not preventable in most cases, but worth understanding because it informs product strategy: a recurring pattern of customers outgrowing the product is a signal to the company, not just to the CS team.

The customer who consolidated

Real consolidation does happen, particularly during economic downturns. The honest version of this pattern is that the product was useful but not useful enough to defend against a CFO mandate to cut vendor count. The protection here is being one of the products the customer cannot do without, which loops back to the dependence argument from the adoption chapter.

Save Plays That Work

When a customer signals they may not renew, the response should be deliberate. A few save plays work reliably across most situations, and an equal number of common responses do not work and should be avoided.

The honest reset

When a customer is unhappy because something went wrong, the play that has the best chance is genuine acknowledgement: the vendor knows what went wrong, takes responsibility for it, and presents a specific plan to do better. This works when the issue is real and the vendor can credibly address it. The customer wants to know they have been heard and that something will change.

The right-size

A customer considering leaving because they over-bought can sometimes be saved by letting them right-size their contract. The reduction is painful in the short term but often preserves a customer who stays for years at a lower tier rather than churning entirely. This requires courage from the CS leader, because the immediate revenue hit is visible while the saved long-term revenue is not.

The executive-level conversation

For strategic accounts at risk, a conversation between senior people on both sides can shift dynamics that operational-level conversations cannot. The executive sponsor relationship discussed in the previous chapter pays off here. The play works because it signals to the customer that they matter, and it gives both sides a forum to discuss the relationship at a level above the immediate friction.

Save Plays That Don't Work

Just as important as knowing which plays to run is knowing which to avoid. Three responses look like saves but are not, and recognising them prevents the function from wasting effort on customers who will churn anyway.

The desperate discount

A large last-minute price cut to retain a customer who has already decided to leave. The customer accepts the discount, stays one more cycle, and then leaves anyway, having extracted twelve more months of usage at half price. The discount did not save the customer; it merely delayed the loss and reduced the value of the time bought. A customer who needs a major discount to renew is usually a customer whose underlying issues were never about price.

The feature promise

Promising the customer that a missing feature will be available "soon" if they renew. This sometimes works if the feature is real and near; it usually works against the vendor if the feature ships late or different, at which point the customer feels deceived and is harder to retain than before. Feature promises made to save renewals tend to be the ones the product team most regrets.

The extra service

Offering additional CSM time, training, or services to save the renewal. Sometimes this addresses a real underlying problem and works. More often it papers over the actual issue (the product did not deliver value) with a layer of service that consumes margin and merely delays the outcome. The trusted-advisor trap from the previous chapter often originates in saves like these.

The common thread across these failed plays is that they address symptoms rather than causes. A customer leaving because they never got value will not be saved by a discount, a feature promise, or extra service. They will be saved only by a credible path to actual value, and that path usually had to be laid down months earlier. By the time the renewal is in question, the options have narrowed considerably.

Building a Renewal Forecast a CFO Will Trust

Most companies forecast renewals badly, which is a problem because a recurring-revenue business's forward financial picture depends on the accuracy of the renewal forecast. A CFO who cannot trust the forecast either pads it with conservatism or ignores it; either response degrades the quality of company-level planning.

A renewal forecast that earns trust has three properties.

It is built bottom-up from actual data

A trustworthy forecast starts from the customer base and works upward: for each customer with a renewal in the period, what is the probability they renew, and at what value. Aggregate this across customers to produce the forecast. The probabilities come from the health scoring system from Chapter 8, calibrated against historical accuracy. A top-down forecast that simply applies a percentage to the prior period's renewals is not a forecast; it is a guess in spreadsheet form.

It separates the probable from the at-risk

A useful forecast distinguishes customers whose renewal is essentially certain from those whose renewal is in question. The first group can be forecast accurately; the second requires explicit assumptions about which way each will go. Combining them obscures the actual risk and produces a number that looks precise while being misleading.

It is held to account

A forecast that is never compared to actuals does not improve. Each quarter, the previous quarter's forecast should be reviewed against what actually happened. Where it missed, why. Where it was accurate, the methods that produced the accuracy. This discipline is what turns forecasting from a ritual into a capability.

The Forecast in Practice

A working renewal forecast typically classifies upcoming renewals into a small number of buckets, each with a defined probability and a defined response. The structure matters more than the specific numbers.

RENEWAL FORECAST — A PRACTICAL STRUCTURE	
Committed (high confidence)	**95% probability**
Likely (healthy signals)	**80% probability**
At risk (mixed signals)	**50% probability**
Critical (negative signals)	**20% probability**
Weighted forecast	**Sum across categories**

The categories should be tied to the health scoring system, not invented separately. A customer with a red health score in October cannot be classified as "committed" in the November renewal forecast without an explicit reason for the divergence. The two systems should agree, or the disagreement should be examined.

Calibration matters. If the "committed" bucket renews at 85 per cent rather than the 95 per cent assumed, the probabilities are wrong and need adjustment. This is the kind of disciplined feedback loop that distinguishes a forecast a CFO can rely on from one they will routinely override with their own conservative estimate. The forecast does not need to be perfect; it needs to be honest about its accuracy and improving over time.

Closing the Chapter

Retention is the metric the Customer Success function is most directly accountable for, and the renewal is where retention is decided. But the work of retention does not happen at the renewal. It happens in the eleven months before, in onboarding that gets customers to value, in adoption that turns use into dependence, in account management that builds multi-threaded relationships and runs the strategic conversations early enough for them to matter.

When the renewal arrives, the options available to the function are largely set by everything that came before. A customer with a strong relationship, deep adoption, and a senior champion does not require a save play; the renewal is a formality. A customer who is none of those things is in a position no save play reliably rescues. The renewal motion is the harvest of a year's work, not a separate activity from it.

Chapter 11 turns from retention to its quieter cousin, expansion. Where retention measures keeping the revenue you have, expansion measures growing it from the customer base you have already won. In mature subscription businesses, expansion is the larger source of growth, and Customer Success is increasingly the function that owns it.

THE BOTTOM LINE

Retention is set by what happens in the eleven months before the renewal, not at the renewal itself. The three types of churn require different responses; the choice of who owns the renewal motion shapes how it is run; **save plays only work when they address causes rather than symptoms**. **A renewal forecast that earns CFO trust is built bottom-up from health data, separates committed from at-risk, and is calibrated against actual outcomes**.

CHAPTER 11

Expansion as a Growth Engine

In mature subscription businesses, more new revenue comes from existing customers paying more than from new customers signing for the first time. This is not a footnote to the growth story. It is the growth story.

If retention is the floor, expansion is the ceiling. Retention preserves the revenue a business has already earned; expansion grows it from the same customer base, often more efficiently than new logo acquisition can. The economics are favourable enough that the most prized subscription businesses on public markets are the ones whose existing customers, by themselves, would grow the business even if no new logos were ever signed.

For Customer Success, expansion represents both an opportunity and a question. The opportunity is to move the function from a defensive role (preventing churn) to an offensive one (driving net new revenue). The question is whether CS is structurally the right function to own that work, or whether it belongs elsewhere, and whether asking CSMs to drive expansion damages the trust that made the function valuable in the first place.

This chapter takes that question seriously and provides a working answer. It covers the economic argument for expansion, the structural debate about who should own it, how to identify expansion opportunities in the data the business already has, how to build playbooks that produce expansion without turning CSMs into junior salespeople, and the compensation structures that make the whole apparatus actually work.

Why Expansion Has Become the Growth Engine

A subscription business with strong retention has a customer base that does not need to be replaced each year. The cost saved on acquisition can be redeployed into deepening relationships with the customers already won. When done well, this produces a flywheel: the customer base itself becomes a source of growth, and each year's growth begins from a larger base than the year before.

The arithmetic is the reason mature subscription businesses look so different from their younger selves. An early-stage company growing at fifty per cent does so primarily through new logos. A mature company growing at twenty per cent often does so with more than half of that growth coming from existing customers expanding their spend. The two companies look similar on a top-line growth chart and structurally are not similar at all.

The economic gap

A dollar of new logo revenue typically costs the business between 80 cents and a dollar to acquire, depending on industry. A dollar of expansion revenue from an existing customer typically costs between 10 and 30 cents. The gross margin difference is the entire reason expansion is the growth engine of mature subscription businesses; it is several times more efficient than the alternative, and the gap widens as the customer base matures.

The other reason expansion matters is that it compounds. A customer who expanded last year is now a larger account, which means they are now a candidate for further expansion at higher absolute dollar amounts. New logos do not compound this way; each one is a fresh start. Expansion produces customers who get more valuable to the business every year they stay, which is the financial profile every investor in a subscription business is hoping to see.

The Four Types of Expansion

Expansion comes in four structurally different forms. Each has its own dynamics and its own implications for who should drive it and how.

1

Seat expansion

More users from the customer organisation using the product. *The motion:* identify teams or roles within the customer that would benefit, then make the case for adding them. The simplest form of expansion and often the largest in aggregate, because it draws on the customer's organic growth as well as deliberate expansion work.

2

Tier upgrades

The same customers moving to a higher product tier with more features or higher limits. *The motion:* demonstrate value the customer would gain by moving up, typically by showing what they cannot do at their current tier. Often driven by changes in the customer's needs as their use of the product deepens.

3 ***Cross-sell***

Selling additional products to the same customer. *The motion:* identify the customer's adjacent needs and position the next product as the solution. The most commercially complex of the four because it involves a new buying decision rather than the expansion of an existing one.

4 ***Usage growth***

More consumption of the product within the existing contract, particularly in consumption-priced products. *The motion:* drive deeper adoption that naturally produces higher usage. Less negotiated than the other three; emerges from the adoption work covered in Chapter 8.

The four are not equally common across business models. Per-seat SaaS products see most of their expansion in seats and tier upgrades. Consumption-priced products see most of it in usage growth. Multi-product portfolios see significant cross-sell. The composition of expansion in a given business is itself diagnostic: it reveals what kind of growth the customer base is producing, and where the most attention should be focused.

Who Should Own Expansion

This is the question with the most disagreement in the field, and the answer matters because it determines how the work gets done. Three models exist, with the same trade-off pattern as the renewal-ownership question from Chapter 10.

MODEL	The case for it	The failure mode
CS-LED	CSMs have the relationship and see expansion signals first; the customer prefers buying from someone they trust.	CSMs without commercial training under-pursue expansion and avoid the harder conversations; trust gets used as an excuse for inaction.
SALES-LED	Sales has the skills and incentive structure to close expansion deals; CSMs stay focused on relationship and adoption.	Sales arrives as a stranger, the conversation feels commercial rather than consultative, and the trust CS built gets damaged in the process.
HYBRID	CS identifies and qualifies; Sales closes. Combines the strengths of both functions if the handoff is clean.	The handoff is rarely clean. CS and Sales each blame the other when deals stall, and the customer experiences the friction.

The right answer depends, again, on segment and stage. SMB and lower mid-market expansion is usually best handled by CS. Enterprise expansion is often best handled by Sales or a hybrid model, because the deals are large enough to warrant the commercial machinery.

But the structural answer is not the most important answer. The real source of disagreement about who should own expansion is not structural.

The Question Underneath the Question

The real question is whether asking CSMs to drive expansion damages the relationship they have built with the customer, and whether the customer experiences the CSM differently once their compensation depends on selling more product.

The honest answer is that it depends on how it is done. A CSM whose compensation is overwhelmingly tied to expansion does, in fact, become a salesperson, and the customer notices. The trust degrades. The CSM begins to be heard differently. The transformation is gradual and often invisible to the CSM themselves, who continues to believe they are advising the customer while the customer begins to see them as a vendor pushing more product.

A CSM whose compensation is mostly tied to retention and adoption, with a meaningful but minority component tied to expansion, can drive expansion without losing the trust that makes the role valuable. The compensation discussion from Chapter 6 (mostly base, modest variable, expansion as one component of variable rather than the dominant one) is exactly what the structural integrity of the CS function depends on when expansion enters the picture.

> *The CSM who is paid mostly to keep customers and partially to grow them is still a CSM. The CSM who is paid mostly to grow customers is a salesperson with a different job title, and the customer can tell.*

This is the conceptual answer to the structural question. CS-led expansion works if the compensation structure preserves the relationship orientation. CS-led expansion fails if the compensation structure converts CSMs into salespeople. The answer to "should CS own expansion" depends almost entirely on how the function is paid.

Identifying Expansion Signals

Most expansion that could be happening in a customer base is not happening because no one has noticed the signals that indicate readiness. The data is usually present; the discipline of looking at it is not. Three categories of signal are particularly useful.

Usage approaching the ceiling

A customer using 80 per cent of their seats, hitting their usage limit, or approaching the boundaries of their tier is signalling that the product is working and that they may need more of it. These signals are mechanical and easy to detect automatically. They are also the easiest expansion conversations to have, because the customer already understands the need.

Adoption spreading beyond the original use case

A customer originally using the product for one workflow who has begun applying it to a second or third workflow is demonstrating organic deepening of the relationship. The expansion conversation here is about helping them do that more effectively, which often means adding seats, tiers, or adjacent products. Easier to spot in product usage data than most teams realise.

Strategic signals from the customer's business

Funding announcements, leadership changes, new initiatives, geographic or product-line expansion. The customer's growth often drives their need for more of the product, and the CSM in a strategic relationship hears about these signals before they become public. A customer entering a growth phase is a candidate for expansion conversations the CSM is uniquely positioned to start.

A useful discipline is to review the customer base systematically for these signals on a regular cadence, rather than waiting for them to surface in the course of normal CSM work. Most expansion opportunities are missed not because the customer was not ready, but because no one looked.

Expansion Playbooks

Once a signal is identified, the response should be deliberate rather than improvised. A few playbooks work reliably across most situations and are worth codifying.

The natural-fit conversation

Triggered when usage data suggests the customer is approaching a constraint. The CSM raises the constraint in the next strategic conversation, frames the expansion as a natural next step given what the customer is already doing, and lets the customer drive the timing. This works because the customer reaches the same conclusion the CSM has reached, rather than feeling sold to.

The new-use-case conversation

Triggered when the CSM sees adoption spreading into adjacent workflows. The conversation is about helping the customer succeed in the new workflow, which often requires adding seats, features, or another product. The framing is consultative rather than commercial: the question is how to make the customer more successful in the work they are already doing, of which expansion is one possible answer among others.

The executive-level conversation

Triggered by strategic signals from the customer's business. The conversation operates at a higher level, often involving an executive sponsor, and frames the expansion as a strategic alignment of the product with the customer's direction. The deal sizes are larger and the conversations take longer, but the close rates are higher because the alignment is real.

Playbooks like these scale expansion work without scaling the CSM headcount linearly with it. A team that has codified its playbooks identifies more opportunities, runs them more consistently, and produces more expansion per CSM than a team that improvises each conversation. The expansion outcomes follow from the system, not from the individual heroics of particular CSMs.

Compensation That Makes Expansion Work

The compensation question was covered in Chapter 6, but it bears revisiting in the context of expansion specifically, because this is where the compensation structure most directly affects behaviour. A CSM whose variable compensation is heavily weighted toward expansion will pursue expansion aggressively, sometimes at the cost of relationships. A CSM whose variable compensation is weighted toward retention with a smaller expansion component will pursue expansion when the opportunity is right and decline when it is not, which is the behaviour the business actually wants.

A workable structure for a CS team that owns expansion looks roughly like this: 75 to 80 per cent base salary, 20 to 25 per cent variable, with the variable split between retention (the largest share), expansion bookings (a meaningful share), and customer health signals (a smaller share). The exact percentages depend on the business, but the relative ordering matters: retention first, expansion second, satisfaction signals third.

Why this structure works

A CSM paid this way will not walk away from a legitimate expansion opportunity, because the variable component is meaningful. They will also not push an expansion that does not make sense for the customer, because the bulk of their compensation depends on the customer continuing to renew at all, which means pushing too hard would risk a larger payout for a smaller one. The compensation aligns the CSM's incentives with the customer's interests, and expansion happens as a natural consequence.

Why pure-commission structures fail here too

Pushing variable compensation higher, particularly the expansion component, breaks the structure. The CSM begins to optimise for the variable rather than the customer relationship. Expansion conversations get pushed when they should not be. Renewals get sacrificed to chase expansion bookings. The pattern from Chapter 6 reasserts itself: pure-commission CSMs stop being trusted, and the trust loss costs more than the commission earned.

Expansion as the Path to High Net Revenue Retention

Net Revenue Retention, introduced in Chapter 4, is the metric that captures the combined effect of retention and expansion against churn and contraction. A business with high NRR is one where existing customers, in aggregate, are paying more this year than last. Above 100 per cent means the customer base is growing on its own. Above 120 per cent is the territory of the most prized subscription businesses on public markets.

There is no path to high NRR that does not run through expansion. Pure retention, even at 100 per cent of revenue, only holds the line. Growth from the existing base requires that some customers pay more than they did last year, by more than the customers who pay less. The shape of high-NRR businesses is one where a meaningful percentage of customers expand each year, by enough to overcome the inevitable contraction and churn elsewhere.

THE COMPOSITION OF 120% NRR — A WORKED BREAKDOWN

Starting ARR base	**$100M**
Gross churn during the year	**− $8M (8%)**
Contraction (downgrades)	**− $3M (3%)**
Expansion from existing customers	**+ $31M (31%)**
Ending ARR from existing base	**$120M (120% NRR)**

The expansion line carries most of the weight, and the gross churn and contraction lines have to be controlled tightly for the NRR to land above 120. A business achieving this is doing the work covered in every previous chapter of Part Two, with expansion as the additive layer on top. The number is not produced by expansion alone; it is produced by retention, adoption, and account management creating customers ready to expand, then by the expansion work itself converting that readiness into bookings.

Closing the Chapter

Expansion has become the growth engine of mature subscription businesses because the unit economics favour it heavily over new logo acquisition. For Customer Success, expansion is both an opportunity and a structural question: it offers the chance to move from a defensive function to a growth driver, but only if the work is structured in a way that preserves the trust that made the function valuable.

The structural answers are the ones the chapter has worked through. Expansion can be owned by CS, by Sales, or by a hybrid, and the right answer varies by segment and stage. The compensation structure determines whether CS-led expansion strengthens or undermines the function. Playbooks turn opportunistic expansion into systematic expansion. And the four types of expansion (seats, tiers, cross-sell, usage) each have their own dynamics that the operating model has to accommodate.

Chapter 12 turns from the work the function does to the systems it does the work with. Customer Success technology has matured rapidly in the past several years, and the choices a leadership team makes about tooling, data infrastructure, and the analytical foundation of the function increasingly determine what is operationally possible. The right tools enable everything in this chapter and the previous ones. The wrong ones make it slower and more expensive at every step.

THE BOTTOM LINE

Expansion is **the growth engine of mature subscription businesses**, and the most efficient source of revenue a business has. Whether CS owns it, Sales owns it, or a hybrid handles it depends on segment and stage; **whether the structure works depends overwhelmingly on the compensation design**. A CSM paid mostly to retain customers, with a meaningful but minority expansion component, drives expansion without losing the trust that made the role valuable. A CSM paid mostly to expand is a salesperson, and the customer can tell.

CHAPTER 12

The Tech Stack and Data Foundation

Buying software does not produce a Customer Success function. It produces a Customer Success function with software, which is a different thing, and only better than the version without if the software was the right one for the work.

Customer Success technology has matured rapidly in the past several years. A category that did not really exist in 2010 has become a multi-billion-dollar software market, with platforms claiming to solve everything from health scoring to renewal forecasting to in-product onboarding. The maturation is real and the tools are genuinely useful. The trap is buying the software before deciding what the function is for, which produces an expensive system that automates the wrong work efficiently.

This chapter is about the technology decisions a Customer Success leader makes and the data foundation those decisions rest on. It covers the four categories of CS tooling and what each is genuinely for, a stage-by-stage view of what to buy when, the data infrastructure that makes any of the tools work, the build-versus-buy question that every CS team eventually faces, and the over-spends and under-spends that show up most often.

The thread running through it all is that tools amplify whatever operating model and process the function already has. A coherent function with the right tools moves faster. A confused function with the right tools moves in the wrong direction faster. The technology decisions in this chapter only produce value if the decisions in Chapters 5 through 11 have been made first.

The Four Categories of CS Tooling

CS technology spans four broad categories. Each addresses a different problem, and confusion about which tool solves which problem is one of the most common sources of overbuying. A team that already has the CRM solving a particular need does not need to buy a CSP to solve the same need; a team buying a product analytics platform should be clear about what it is doing that the CSP does not.

1

Customer Success Platforms (CSPs)

The category built explicitly for CS work: customer 360 views, health scores, playbook execution, renewal management, alert workflows. *What it is genuinely for:* giving CSMs a single workspace that surfaces what to do, for which customer, when. Examples include Gainsight, Totango, ChurnZero, Catalyst. Useful only once the function is large enough to need shared workflows.

2

CRM Systems

The system of record for customer relationships, originally built for Sales but doing increasing duty for CS in many companies. *What it is genuinely for:* maintaining the source of truth for accounts, contacts, contracts, and pipeline. Most CS work in early-stage companies happens here, often with light extensions for CS-specific fields.

3 *Support Systems*

Ticketing, escalation, and knowledge base systems, including Zendesk, Intercom, Freshdesk and others. *What it is genuinely for:* handling reactive customer issues, capturing the volume and type of issues for analysis, and providing self-service resources. Distinct from CS work but often the source of important signals about customer health.

4 *Product Analytics*

Tools that capture what customers actually do in the product: usage data, feature adoption, behavioural patterns. Examples include Pendo, Mixpanel, Amplitude, Heap. *What it is genuinely for:* providing the usage signals that feed health scores and adoption playbooks. The foundation underneath any data-driven CS function.

The four categories are best understood as complementary rather than substitutable. The CRM holds the relationship; the support system handles reactive issues; product analytics captures behaviour; the CSP, when it makes sense to buy one, integrates the others into a workflow CSMs can act on. A team that confuses the categories ends up either buying tools that solve problems they already had solved, or assuming a single tool will do work that genuinely requires more than one.

How the Categories Fit Together

A working CS technology stack uses all four categories, with each performing the role it is suited to and clear handoffs between them. The CRM holds the customer record. The product analytics tool captures behavioural data. The support system handles reactive issues. The Customer Success platform pulls these together into actionable views and workflows.

The most common failure mode is buying tools that overlap heavily and then arguing about which one is the source of truth. A team that uses a CSP for customer records, a CRM for the same records, and a spreadsheet for the records that did not fit either, has three versions of the truth and no actual one. The discipline is in deciding, deliberately, which system owns which data, and making the other systems read from it rather than maintaining their own copies.

A useful question to ask before buying

If we already have a CRM, a support system, and a product analytics tool, what specifically does the CSP add that the other three cannot? If the answer is "a single dashboard for CSMs," the CSP is potentially worth it. If the answer is "health scores" or "renewal forecasting," the team may be paying a six-figure platform for something a competent operations person could build in the existing CRM. The honest answer determines whether the spend is justified or merely tempting.

The right stack is not necessarily the largest. A team with a well-configured CRM, a product analytics tool, a support system, and disciplined operations can produce most of what a CSP provides, at a fraction of the cost, until they hit a scale where the operational overhead exceeds the platform license. The discipline is to buy when buying produces leverage, not when the marketing makes the case look compelling.

What to Buy When

The right tools at each stage are different, and the most expensive mistakes come from buying tools the company is not yet ready to operate. The progression below is a starting point, not a rule, but it reflects the patterns that hold up across most subscription businesses.

STAGE	What to buy	What not to buy yet
SEED / SERIES A	CRM (likely already in place), basic support system, lightweight product analytics if the data is critical.	CSP. Not enough customers to justify it; spreadsheets and CRM extensions are sufficient.
SERIES B	Better product analytics, mature support system, beginning to evaluate CSP options.	Premium CSP tier with every module. Buy the simplest version that meets current needs, expand later.
GROWTH	CSP becomes worth the investment; deeper product analytics; integration work between systems.	A second CSP "to evaluate alongside the first." Pick one, commit, get good at it.
SCALE	Mature CSP, dedicated CS operations function, data warehouse feeding multiple systems.	Custom-built replacements for off-the-shelf tools, unless the differentiation is real and defensible.

The pattern is consistent across all four stages: buy the next tool when the operational pain of not having it exceeds the cost of having it, not before. A CSP bought too early sits underused and underperforms its potential. A CSP bought too late means the team has been doing in spreadsheets what a platform could have done for them, at a cost in time that vastly exceeds the platform licence.

The Data Foundation Underneath

No CS tool works better than the data flowing into it. A CSP with poor data is a beautiful dashboard of meaningless numbers; a CRM without consistent CS fields is a customer record nobody trusts. The data foundation is what makes the tools work, and it deserves more attention than it typically gets.

A working CS data foundation has three properties.

A clear source of truth for each kind of data

Customer records belong in one system. Usage data belongs in another. Support ticket history belongs in a third. The systems read from each other through defined integrations, not by maintaining independent copies. When two systems disagree about a customer's status, there should be a defined rule for which one is authoritative.

Consistent data entry discipline

Even the best tooling fails if the underlying data is patchy or inconsistent. A CSM who fills in health scores honestly produces useful data; one who marks everything green to avoid difficult conversations produces noise. The discipline of accurate, honest data entry is more important than the tool capturing it, and it has to be modelled by the leadership team before it propagates to the CSMs.

Someone responsible for the data itself

Past a certain scale, a dedicated CS operations role earns its cost almost immediately by maintaining the data quality the rest of the function depends on. Without that role, every CSM ends up partly responsible for data hygiene, which means no one is fully responsible, which means the data drifts over time and the tools that depend on it become less reliable.

Build versus Buy

A recurring question in CS technology is whether to build internally or buy from a vendor. The honest answer depends on what is being built and what the company's broader engineering capacity looks like, but a few principles hold up reliably.

Buy for the commoditised work

CRM, support ticketing, basic product analytics, email automation: these are mature, well-served markets where vendors compete on quality and the price-per-feature is reasonable. Building these internally is almost always a mistake. It diverts engineering resources from work that actually differentiates the business and produces an inferior version of something that can be bought working.

Build for the differentiated work

Where the company has unique requirements that no off-the-shelf tool addresses well, building can make sense. Custom health scoring models that reflect the specific way the business's customers actually behave. Specialised dashboards that combine signals no vendor offers. Integration glue between systems that vendors do not provide. These are legitimate build candidates when the differentiation is real.

Beware the in-between

The dangerous territory is building partial replacements for tools that mostly exist. A team that decides to "build our own CSP because Gainsight is too expensive" usually underestimates the engineering investment, ends up with a partial version, and then maintains it forever as technical debt. The lower licence cost gets eaten many times over by the maintenance burden.

A working rule: if a credible vendor solves the problem reasonably well, buy. Build only for the problems where no vendor solves it well enough, and where the company has the engineering bandwidth to maintain what it builds for years. The total cost of ownership of internal tools is consistently larger than companies estimate at the time of the decision.

Common Over-Spends

A few spending patterns appear in CS technology decisions repeatedly, where the money goes out the door and the operational value never quite materialises.

The premium CSP for a small team

A Series A company buys a top-tier CSP with every module activated, on the theory that the platform will allow the team to scale. Eighteen months later, the team uses fifteen per cent of the platform's capability and the licence cost is materially larger than its operational return. The platform was not wrong; the timing was. The same platform purchased two years later, against a function that had actually grown into needing it, would have been a sound investment.

Multiple overlapping tools

The CSP overlaps the CRM, which overlaps the support system, which overlaps the product analytics platform. Each was bought to solve a problem the others did not solve at the time, but the overlaps grew as each tool added features. The team now maintains four tools where two would do, and the integration work to keep them synchronised consumes more time than the tools themselves save.

Implementation consultants who outstay the implementation

A platform implementation that was supposed to take three months becomes a permanent consulting engagement, with the external team owning the platform's configuration indefinitely. The internal team never learns to operate it independently; every change requires a consultant's billable hours. The platform's recurring cost is now the licence plus the consultant retainer, often a multiple of the original budget.

Common Under-Spends

The opposite pattern is just as common and often more expensive over time: the spending that did not happen and should have, where the operational cost of not having the tool quietly exceeds the cost of buying it.

Product analytics that no one is collecting

A subscription business operating without serious product analytics is flying blind on adoption. The CS team can guess at how customers are using the product; they cannot know. Investing in product analytics is one of the cheapest, highest-return technology decisions a CS function can make, and it is consistently delayed because the cost is visible and the return is not.

CS operations headcount

A team large enough to need a dedicated operations person, but operating without one. The CSMs spend hours each week on data hygiene, dashboard building, and process documentation that an operations specialist would handle in a fraction of the time. The unbudgeted operational drag is often larger than the cost of the hire would have been.

Integration work between systems

The systems exist but do not talk to each other. The CRM has account information, the product analytics tool has usage data, the support system has ticket history, and none of them appear together in a useful view. CSMs alt-tab between five tools to assemble a picture of a single customer. The integration work to bring them together would pay for itself in CSM time within months, and yet is consistently deprioritised.

A Technology Strategy, Not a Shopping List

The most useful frame for these decisions is to treat CS technology as a strategy rather than a shopping list. The strategy answers a small number of questions, in order, before any individual purchase is considered.

What is the function on the hook for, expressed in metrics? What work do CSMs need to be doing every week to produce those metrics? What data do they need to do that work? Where does that data live, and how does it reach them? Once those questions are answered, the tooling decisions follow naturally. The platform that solves the actual problem becomes obvious; the platforms that solve adjacent problems can be deferred.

A CS leader who can articulate the strategy in this form can defend the technology budget to a CFO, sequence investments coherently, and avoid the pattern of accumulating tools that each made sense when bought but together produce less than their cost. A CS leader who cannot articulate the strategy in this form will end up with whatever stack the vendor sales teams successfully sold, which is rarely what the function actually needed.

Chapter 13 turns from the tools the function uses to the organisation around it. The question of how to scale Customer Success from one to ten to fifty CSMs, when to introduce specialisation, and how the leader's role changes at each phase is the subject of the next chapter.

THE BOTTOM LINE

CS technology **amplifies whatever operating model the function already has**, for better or worse. The four categories (CSP, CRM, support, product analytics) each have a distinct role; over-buying happens when tools overlap and under-buying happens when the operational cost of not having the tool is invisible. **Buy for the commoditised work, build only for the genuinely differentiated work**, and treat the technology stack as a strategy that follows from the function's metrics, not a shopping list that precedes them.

CHAPTER 13

Org Design and Scaling

The CS function that works at ten people does not work at fifty. The mistake is to grow the team by adding more of the same and discover, two years later, that the shape of the organisation has stopped being capable of doing the job.

Most CS organisations grow by accretion. The first CSM is hired, the second follows, then a third, and the team scales by simply having more of them. For a while this works. The CSMs are generalists, each owning a portfolio of accounts and doing all the work each account requires. Hiring is the only growth lever, and the function grows linearly with the customer base.

At some point this pattern stops working. The team is too large to manage flat; the customer segments demand different motions; the work itself is differentiated enough that asking everyone to do all of it produces no one doing any of it well. The function has to change shape, not just size, and the transitions involved are the most consequential organisational decisions a CS leader makes.

This chapter is about those transitions. It covers what breaks at each scaling phase, when to introduce specialisation, how the structure of the function changes from generalist to specialised, the additional complexity of building CS across global timezones, and the way the leader's own role shifts as the function grows.

The Three Scaling Phases

CS organisations move through three broad phases as they grow. The transitions between them are not gradual; each represents a structural change in how the function operates. Recognising which phase the function is in, and which transition is approaching, is most of what a CS leader needs to navigate the scaling problem well.

1

The Generalist Phase (1 to 10 CSMs)

Every CSM is a generalist, owning their accounts end to end. The leader is a player-coach, often carrying accounts personally while managing the team. Process is informal; tools are light. *What works:* speed, individual ownership, low overhead. *What breaks first:* inconsistency between CSMs, the leader's bandwidth, the inability to specialise where the work demands it.

2

The Specialised Phase (10 to 50 CSMs)

The function differentiates. Onboarding specialists, renewals specialists, and segment-specific teams emerge. The leader stops carrying accounts and runs the organisation. CS operations becomes a real function. *What works:* consistency, scale, ability to optimise each part of the motion separately. *What breaks first:* coordination between specialised teams, the customer experiencing handoffs they did not before, the cost of the operating apparatus.

3

The Multi-Function Phase (50+ CSMs)

CS becomes a multi-function organisation in its own right, often under a Chief Customer Officer. Separate sub-organisations handle different segments, regions, or product lines. The leader's role is closer to a CEO of a smaller business than to a frontline manager. *What works:* ability to operate at enterprise scale across segments and geographies. *What breaks first:* the leader's distance from the customer, organisational politics, the slow accumulation of bureaucracy.

What Breaks at the First Transition

The shift from generalist to specialised, somewhere between five and ten CSMs depending on segment, is the harder of the two transitions. Most leaders see it coming late, after the function has already started showing the symptoms.

Inconsistency between CSMs

In the generalist phase, the best CSM produces noticeably different outcomes from the average CSM. Some of this is talent; much of it is that each CSM has developed their own approach with no shared system. As the team grows, the inconsistency becomes harder to ignore. Customers in different portfolios have meaningfully different experiences; renewal rates vary by CSM more than they should. The fix is specialisation and shared playbooks, which is what the next phase introduces.

The leader's bandwidth

A player-coach leader can stay close to maybe six or seven CSMs while still carrying their own accounts. At ten, the leader is permanently behind. Either the leader stops carrying accounts (and the team feels the loss of capacity) or the leader keeps carrying them (and the team feels the loss of management). Either way, the structure that produced the leader-as-IC stops working, and the role has to formally change.

Work that everyone does, badly

Onboarding requires one set of skills; renewals require another; strategic account work requires a third. In the generalist phase, every CSM does all of these, with the predictable result that they are good at some and weaker at others. The variance is masked when the team is small. It becomes the dominant source of organisational pain as the team grows, because the customer experience of any one of these motions depends on which CSM happened to draw their account.

When and How to Introduce Specialisation

Specialisation is not a single decision. It is a sequence of decisions, each made when the function reaches the scale to justify it. The order matters, because introducing a specialisation prematurely is its own failure mode.

Onboarding specialists, usually first

The first specialisation most CS functions introduce is onboarding. The work is distinct enough (project management, technical configuration, intensive first-ninety-days engagement) that it benefits from people dedicated to it. The CSMs who hand off to onboarding specialists initially resist, then come to appreciate the focus on their renewal and adoption work. The threshold is usually around eight to ten CSMs, and the return on the investment is visible within a year.

Segment specialisation, often second

As the customer base diversifies, splitting the team by customer segment becomes useful. Enterprise CSMs, mid-market CSMs, and SMB CSMs each operate on different motions, with different tools, ratios, and playbooks. This specialisation typically follows onboarding specialisation by twelve to eighteen months, once the team is large enough to support distinct sub-teams for each segment.

Renewals and expansion specialists, sometimes

Whether to introduce dedicated renewals or expansion teams depends on the operating-model choices from Chapter 10 and Chapter 11. Not every CS function should do so. The threshold is high: enough customers to support specialised teams that do not dilute the CSM relationship, and a deliberate choice that the commercial motions are best handled by people whose entire focus is on them.

The progression is rarely linear and the right sequence depends on segment composition. What matters is that each specialisation is introduced because the function has reached the scale to need it, not because the leader has read about it in a book.

What Generalists vs. Specialists Own

The transition from generalist to specialised teams changes what each role is responsible for. The shift is worth making explicit, because the change is often felt as a loss by the CSMs whose scope is narrowing, even when the overall outcome for the customer improves.

STAGE OF WORK	Generalist phase	Specialised phase
SALES HANDOFF	CSM receives the customer directly from Sales.	Onboarding specialist receives the handoff; CSM joins later.
ONBOARDING	CSM owns end to end.	Onboarding specialist owns; CSM coordinates and observes.
ONGOING RELATIONSHIP	Same CSM through full lifecycle.	CSM takes over after onboarding handoff.
RENEWALS	CSM runs them.	Renewals specialist, with CSM context; varies by segment.

The visible cost of specialisation is the handoffs it introduces. A customer who had one relationship now has two or three, and the handoffs between them have to be handled deliberately or the customer experiences a confusing series of transitions. Done well, the customer gains specialised expertise at each phase and loses very little. Done badly, the customer feels passed around, and the specialised model produces worse outcomes than the generalist one it replaced.

The Second Transition: Becoming a Multi-Function Org

The shift from a specialised CS team to a multi-function CS organisation, somewhere around fifty CSMs, is structurally different from the first transition. It is less about specialisation within the function and more about the function itself becoming complex enough to require its own internal management layers.

A few changes characterise this phase.

A Chief Customer Officer (or equivalent) becomes the role

The CS leader is no longer running a function with a handful of direct reports. They are running an organisation with managers underneath them, and the work is closer to running a smaller business than to managing a team. The skills the role requires shift accordingly: less about being the best CSM in the room, more about strategy, executive presence, and cross-functional politics.

CS operations becomes a peer function

In the specialised phase, CS operations is a small group inside the CS team. In the multi-function phase, it is a recognisable peer function with its own leader, often a Head of CS Operations, who runs the data, tooling, and process infrastructure for the whole organisation. The role's seniority signals that operations is treated as a strategic capability, not a clerical service.

Regional and product-line splits emerge

At sufficient scale, the function splits by geography (Americas, EMEA, APAC) or by product line, with separate sub-organisations under the CCO. Each operates with significant autonomy, on the model that worked for the unified function at the smaller scale, but now adapted to its specific context. Coordination across these sub-organisations becomes a real management problem in its own right.

Building CS Across Timezones

A company selling globally needs CS support globally, which sounds straightforward and is genuinely difficult. The challenges of running CS across timezones and cultures compound the structural problems of scaling, and the failures are common enough to be worth naming.

Coverage versus consistency

A customer in Singapore wants their CSM to be reachable during Singapore business hours, not at 11pm because the CSM is in San Francisco. Solving this requires hiring CSMs in the customer's region, which solves coverage but introduces consistency problems: regional teams develop their own approaches, and the global function fragments. The discipline is in supporting regional autonomy on customer-facing work while maintaining shared playbooks, tooling, and standards across regions.

Cultural differences in CS expectations

Customer expectations of CS vary by region. North American customers often expect proactive engagement and frequent contact. European customers may prefer more independence and lighter touch. Some Asian markets expect a level of relationship-building and ceremony that other regions consider excessive. A global CS function that applies a single approach across regions produces worse outcomes in most of them. The motions have to be calibrated to regional expectations while remaining recognisably the same function.

The follow-the-sun trap

Some companies try to provide round-the-clock CS coverage by handing off accounts between regional teams as the sun moves. This usually fails. Customer relationships do not transfer cleanly mid-conversation, and the operational overhead of the handoffs exceeds their value. Most successful global CS functions assign each customer to a single region and team, with that team responsible for the relationship regardless of where the customer is. Coverage gaps are accepted as the cost of the simpler model.

How the Leader's Role Changes

The CS leader's job is structurally different at each phase, and one of the more common failure modes is a leader who scales the team without scaling themselves. The leader who was excellent at running ten CSMs personally is often the wrong person to lead an organisation of fifty, and recognising this honestly is one of the harder discussions in a growing company.

In the generalist phase

The leader is a player-coach. They carry their own accounts, manage the team, and act as the function's external face. The role's value comes from individual contribution as much as from management. The team is small enough to be managed by walking around, and the leader knows every customer's situation in detail.

In the specialised phase

The leader stops carrying accounts. They run the organisation: hiring, structure, process, performance. They manage managers rather than individual contributors. The role's value comes from organisational design, from the quality of the leadership bench they build, and from the political work of representing CS to the rest of the company.

In the multi-function phase

The leader is functionally a CEO of a smaller business. They set strategy, manage senior leaders, allocate resources between sub-organisations, and represent the function to the board and the executive team. They are several layers removed from the customer, and the discipline of staying connected to the customer reality through their senior reports is its own management challenge.

A leader who transitions cleanly through these phases is rare. More commonly, the company has to change leaders at one of the transitions, either through promotion, lateral move, or external hire. This is uncomfortable but often correct; the same skills do not produce success at all three scales.

Closing the Chapter

A Customer Success function does not scale by becoming a bigger version of itself. It scales by changing shape: introducing specialisations, building management layers, splitting by region or segment, and shifting the leader's role at each phase. The transitions are predictable, in the sense that the patterns are the same across most companies, but they are also genuinely difficult, because the shape of the function that worked yesterday is often the wrong shape for tomorrow.

Recognising the transition before it becomes acute is most of the work. A leader who sees the inconsistency, the bandwidth problem, and the work-everyone-does-badly pattern, and who acts on them by introducing specialisation before they become crises, manages the first transition well. A leader who waits until the function is visibly failing has a much harder problem to solve and less time to solve it.

Chapter 14 turns from the structural questions that hold across most subscription businesses to the variations that matter by industry and business model. Enterprise SaaS, SMB SaaS, product-led growth, fintech, healthcare, marketplaces, and hardware-as-a-service each shape the CS function in different ways, and the next chapter walks through what changes and what stays the same.

THE BOTTOM LINE

CS organisations scale through **three structural phases**: generalist, specialised, and multi-function. Each transition involves changing the shape of the function rather than just its size. **The leader's role changes at each phase as much as the team's structure does**, and the same skills that produced success at one scale often do not produce it at the next.

CHAPTER 14

How CS Differs by Industry

Almost everything written about Customer Success assumes a B2B SaaS company selling to mid-market businesses. Most of it transfers to other contexts, but the parts that do not transfer are exactly the parts that matter most.

The principles in this book are general, but their application is not. A Customer Success function at an enterprise software company looks different from one at a self-serve product company, which looks different again from one at a regulated fintech or a hardware-as-a-service business. The differences are not cosmetic. They change the touch model, the metrics, the hiring profile, and sometimes the fundamental purpose of the function.

This chapter maps the major variations. It is organised around the dimensions that produce the biggest differences: the contrast between enterprise, SMB, and product-led motions, and then the specific demands of regulated industries, platform businesses, and physical-product subscriptions. For each, the question is the same: what changes, and what stays the same.

The thread throughout is that the strategic logic of Customer Success is constant (retention, expansion, and advocacy compound; the function exists to drive them) while the execution varies enormously. A leader moving from one context to another should expect the principles to transfer and the playbooks not to.

The B2B SaaS Default

Most Customer Success thinking assumes a particular context, even when it does not say so. The default is a B2B software company selling subscription products to other businesses, with contracts in the tens of thousands of dollars per year, a sales-led motion, and customers who need ongoing help to get value from a moderately complex product.

This default is worth naming because it is invisible. When an article says CSMs should run quarterly business reviews, it assumes the customer is large enough to warrant one. When it says the function should drive expansion, it assumes a product with expansion potential. When it recommends a health score built on usage data, it assumes the product generates usage data worth scoring. All of these assumptions hold for the default context and break, partially or completely, outside it.

The rest of this chapter is organised around where and how the default breaks. The first and largest variation is the spectrum from enterprise to self-serve, which changes the function more than any industry vertical does.

A reader's caution

If your business does not match the default, the right response to general CS advice is not to discard it but to translate it. Ask what the advice is actually trying to achieve, then ask how that goal is best achieved in your context. The principle usually transfers even when the specific recommendation does not.

Enterprise, SMB, and Product-Led: Three Motions

The single biggest variation in Customer Success is not the industry vertical but the motion: how the product is sold, how much each customer is worth, and how customers expect to be served. Three broad motions dominate, and they are different enough to be almost different professions.

1 ***Enterprise***

High contract values, long sales cycles, complex deployments, many stakeholders. *What the CS function looks like:* high-touch, senior CSMs, named relationships, executive sponsorship, quarterly business reviews. The work is relationship-heavy and strategic. The whole apparatus this book describes applies most directly here, because this is the context most CS practice was developed for.

2 ***SMB***

Lower contract values, faster sales, simpler deployments, fewer stakeholders. *What the CS function looks like:* lighter touch, pooled or tech-touch coverage, heavy reliance on playbooks and automation, metrics-driven rather than relationship-driven. The economics do not support named CSMs for most accounts, so the function scales through systems rather than people.

3 ***Product-Led (PLG)***

Users adopt the product directly, often free, and convert to paid through usage. *What the CS function looks like:* mostly in-product, data-driven, focused on activation and conversion rather than relationship management. The "customer" may be thousands of individual users rather than a single buying organisation. CS here is closer to growth and product than to traditional account management.

A company often runs more than one of these motions at once: an enterprise motion for large accounts, a self-serve motion for small ones, and a product-led motion for the free tier that feeds both. The CS function then has to operate all three, which is one of the harder organisational problems in the field, because the three motions require different people, metrics, and tools.

How the Three Motions Differ in Practice

The differences between the three motions show up in almost every operational decision the function makes. Laying them side by side makes the contrast concrete.

DIMENSION	Enterprise	SMB	Product-Led
TOUCH MODEL	High-touch, named CSM	Pooled / low-touch	Tech-touch, in-product
PRIMARY METRIC	Net revenue retention	Gross retention	Activation & conversion
CSM PROFILE	Senior, strategic	Efficient, playbook-driven	Data & growth-oriented
EXPANSION	Negotiated, relationship-led	Semi-automated	Self-serve, in-product

The table understates how different these motions feel in practice. An enterprise CSM and a product-led growth specialist do jobs with the same title and almost nothing else in common. Moving a person from one motion to the other, on the assumption that CS is CS, is one of the more common and more costly hiring mistakes a scaling company makes.

Regulated Industries: Fintech, Healthcare, and Beyond

Selling into regulated industries changes Customer Success in ways that go beyond the touch model. The regulation itself becomes part of the relationship, and the CS function has to accommodate constraints that simply do not exist in unregulated contexts.

Compliance becomes part of the value

In a regulated industry, helping the customer stay compliant is often as important as helping them get value from the product. A CSM serving a healthcare or financial-services customer needs to understand the regulatory environment well enough to be credible, and the product's role in the customer's compliance posture becomes a central part of the relationship. This raises the bar on CSM expertise considerably.

The sales cycle and the relationship are longer and heavier

Regulated customers move slowly, involve more stakeholders (including legal, compliance, and security), and require more documentation. The CS relationship inherits this weight. Onboarding takes longer, changes require more approval, and the customer's tolerance for product issues is lower because the issues can have regulatory consequences. The function has to be staffed and paced accordingly.

Trust is the entire game

In regulated industries, a vendor's reliability and security posture are not features; they are preconditions. A single serious incident can end a relationship that took years to build, regardless of how good the day-to-day CS work was. The CS function's role in maintaining trust, communicating transparently during incidents, and demonstrating ongoing reliability is heightened to a degree that unregulated businesses rarely experience.

Marketplaces and Platforms

Marketplace and platform businesses introduce a structural complication that single-sided businesses do not have: there is more than one kind of customer, and their interests do not always align. A marketplace has buyers and sellers; a platform has developers and end users. Customer Success has to serve both sides, and what is good for one is sometimes bad for the other.

Two customers, sometimes in tension

A marketplace CS function serving sellers wants those sellers to succeed, which often means helping them capture more value from buyers. The same function serving buyers wants buyers to get good outcomes, which sometimes means protecting them from sellers. Balancing these is a structural challenge with no clean solution; the function has to hold both sides' interests simultaneously and make deliberate choices about where they conflict.

Network effects change retention dynamics

In a platform with strong network effects, retention is driven partly by the network itself rather than by the CS relationship. A seller stays on the marketplace because that is where the buyers are, not because their CSM is excellent. This changes the function's role: less about preventing churn through relationship, more about helping participants succeed within a network whose gravity does most of the retention work. The CSM's leverage is real but operates differently.

Scale is extreme

Marketplaces often have enormous numbers of participants, most of them small. The economics rarely support human CS for the long tail, which means tech-touch and self-serve models dominate, with human CS reserved for the largest participants on each side. The function looks more like the SMB or product-led motion than the enterprise one, even when individual large participants warrant high-touch attention.

Hardware-as-a-Service and Physical Subscriptions

When the subscription involves a physical product, Customer Success acquires dimensions that pure software businesses never deal with. The product can break in the physical world; logistics enters the relationship; and the cost structure is fundamentally different because there are real marginal costs per customer.

The product exists in the physical world

Hardware fails in ways software does not, and the failures are visible and disruptive. A CS function for a hardware-as-a-service business has to account for device health, replacement logistics, and field service in a way no software CS function does. The health score includes signals about the physical device, not just usage of the software layered on top of it.

Marginal cost changes the economics

Pure software has near-zero marginal cost per customer, which is why software CS can spend freely on retention; the customer kept is nearly pure margin. Physical-product subscriptions have real per-unit costs, which compress the margin available to fund CS work. The function has to be leaner, and the calculations about how much to spend retaining a customer are tighter because the customer is less profitable to begin with.

Logistics becomes part of the experience

Delivery, installation, maintenance, and returns are part of the customer experience in a way they never are for software. A CS function for a physical-product subscription has to coordinate with logistics and field operations functions that software CS does not even have. The relationship is only as good as the weakest link in a longer operational chain.

What Stays Constant Across All of Them

Having spent the chapter on variation, it is worth returning to what does not vary. The execution of Customer Success changes enormously across these contexts. The strategic logic does not.

In every context, retention is more valuable than acquisition over the lifetime of a customer base, and the gap compounds. In every context, customers who reach value quickly behave differently from those who do not. In every context, expansion from existing customers is more efficient than new logo acquisition. In every context, the function's job is to drive retention, expansion, and advocacy, even when the means of doing so look completely different.

This is why the principles in Part One of this book are general while the playbooks in Part Two require translation. A leader who understands the strategic logic can adapt the execution to any context. A leader who has only memorised the playbooks of one context will struggle the moment they move to another, because they learned the what without the why.

> *The strategic logic of Customer Success is universal. The execution is local. A leader who confuses the two, treating the playbooks of their last context as laws rather than adaptations, will misapply them in the next.*

The variations in this chapter are not exhaustive. Every business has idiosyncrasies that shape its CS function in ways no general treatment can anticipate. But the method for handling them is consistent: hold the strategic logic constant, and adapt the execution to the context in front of you.

Closing the Chapter

Customer Success varies more by motion than by industry. The spectrum from enterprise to SMB to product-led changes the function more than any vertical does, because it changes the fundamental economics of how much can be spent serving each customer and how customers expect to be served. Regulated industries add the weight of compliance and the heightened importance of trust. Marketplaces add a second customer whose interests sometimes conflict with the first. Physical-product subscriptions add logistics and real marginal costs that compress the economics.

Through all of it, the strategic logic holds constant. Retention compounds, value-speed matters, expansion is efficient, and the function exists to drive the outcomes that make a recurring-revenue business work. The execution is what changes, and a leader who keeps the distinction clear can move between contexts without losing their footing.

Chapter 15, the final chapter of Part Two, turns to the question on every CS leader's mind: how artificial intelligence is reshaping the discipline. It separates what AI genuinely changes from what the vendors claim it changes, and offers a framework for thinking about the function over the next decade.

THE BOTTOM LINE

Customer Success varies **more by motion than by industry**. Enterprise, SMB, and product-led are almost different professions; regulated industries, marketplaces, and physical-product subscriptions each add their own constraints. **The strategic logic is universal; the execution is local**, and a leader who keeps the distinction clear can adapt to any context while one who memorised only the playbooks cannot.

CHAPTER 15

AI and the Next Decade of CS

Every vendor now claims their product is AI-powered, and most of the claims are noise. Underneath the noise, a smaller set of real changes is reshaping what Customer Success teams do. Separating the two is the work of this chapter.

Artificial intelligence is reshaping Customer Success, and most of what is said about how is wrong. The vendor claims run far ahead of the reality, the predictions of wholesale automation misunderstand what the function actually does, and the genuine changes are quieter and more specific than the hype suggests. A CS leader trying to make sense of it needs a way to separate the real from the marketed, which is what this chapter aims to provide.

The approach here is deliberately conservative, because a chapter about technology in a book meant to last has to distinguish durable changes from passing excitement. The specific tools will change; some that seem essential today will be forgotten in three years, and capabilities that seem like science fiction will become routine. What will not change as quickly is the underlying logic of where AI helps, where it does not, and what it asks of the people running the function.

The chapter covers what AI genuinely changes today versus what vendors claim, the three layers of CS work that AI is reshaping, what disappears and what gets created, how the CS leader's skillset shifts, and a practical framework for cutting through vendor pitches to the substance underneath.

What AI Actually Changes Today

The honest starting point is that AI changes less, today, than the marketing suggests, and changes it in narrower ways. This is not a reason for complacency; the real changes are significant. It is a reason for precision about which changes are real.

What AI genuinely does well today, in the context of Customer Success, falls into a few categories. It is good at finding patterns in large volumes of data that a human would miss, which makes it useful for signal detection. It is good at drafting and summarising text, which makes it useful for the administrative load CSMs carry. It is good at handling routine, well-defined interactions, which makes it useful for the high-volume, low-complexity end of customer contact. These are real capabilities and they are already changing the work.

What AI does not do well today is the core of high-value Customer Success: building trust with a customer's leadership, navigating a complex political situation inside the customer organisation, exercising judgement about a non-obvious situation, or holding a difficult conversation with a customer considering departure. These are the activities that produce most of the function's value in the enterprise context, and they remain stubbornly human.

> *The pattern is consistent across most functions AI touches: it absorbs the routine and the data-intensive while leaving the relational and the judgement-heavy to people. Customer Success happens to be a function with a great deal of both, which is why AI changes it substantially without replacing it.*

The Three Layers AI Is Reshaping

The changes AI brings to Customer Success operate at three distinct layers. Separating them is useful because they mature at different rates, demand different responses, and carry different risks.

1

Signal Detection

AI analysing usage data, support interactions, and communication patterns to surface risks and opportunities earlier than a human could. *Maturity:* the most developed layer today. *What it changes:* health scoring becomes more predictive, at-risk customers surface earlier, expansion signals get caught that would otherwise be missed. The most immediately useful application, and the lowest risk.

2

Automation

AI handling routine work that previously consumed CSM time: drafting follow-ups, summarising calls, preparing QBR materials, answering common questions. *Maturity:* developing rapidly. *What it changes:* CSMs spend less time on administrative work and more on the relational and strategic work AI cannot do. The promise is leverage, not replacement, if the freed time is redirected well.

3

Customer-Facing Intelligence

AI interacting directly with customers: in-product assistants, automated onboarding guidance, conversational support. *Maturity:* least mature for complex work, increasingly capable for routine interactions. *What it changes:* the tech-touch end of the customer base can be served better and at larger scale, freeing human CS for the accounts that warrant it. Highest potential and highest risk of getting the customer experience wrong.

The three layers should be adopted in roughly this order. Signal detection is low-risk and immediately useful. Automation requires more care but pays off quickly. Customer-facing intelligence demands the most caution, because a poor automated customer experience can damage the relationships the function exists to protect.

Where AI Helps Most: Signal Detection

The clearest near-term value of AI in Customer Success is in signal detection, because it plays directly to what the technology does best and avoids what it does worst. A health score built on AI-analysed signals can be meaningfully more predictive than one built on hand-weighted rules, because the AI can find patterns across far more data than a human could reasonably weight.

The concrete improvements are real. AI can detect subtle shifts in how a customer communicates that precede churn, patterns in usage that indicate an account is about to expand or contract, and combinations of signals that no rule-based system would have flagged. The customer who is using the product steadily but whose email sentiment has cooled, whose champion has gone quiet, and whose support tickets have taken on a particular tone is exactly the kind of multi-signal pattern AI catches and rules miss.

The caution that still applies

A more predictive signal is only valuable if someone acts on it, which returns to the argument from Chapter 8: a health score, however sophisticated, is decoration unless a change in score reliably triggers a response. AI makes the detection better; it does not make the response happen. A team that buys AI-powered signal detection and does not build the response system around it has bought a more accurate alarm that still nobody answers.

The other caution is that AI signal detection can be confidently wrong in ways that are hard to audit. A rule-based score is transparent: you can see why a customer was flagged. An AI score can be opaque, flagging a customer for reasons no one can fully explain. This is acceptable when the stakes are low and the track record is good, but it requires the same calibration discipline that Chapter 8 recommended for any health score: check the predictions against outcomes, and trust the system only as far as its track record justifies.

What Disappears, What Changes, What Gets Created

The honest way to think about AI's effect on the Customer Success workforce is in three categories: work that disappears, work that changes, and work that gets created. The category that gets the most attention (work that disappears) is the smallest of the three.

What disappears

The administrative and routine work that consumed a surprising fraction of CSM time: manually summarising calls, drafting routine follow-ups, compiling data for QBRs, answering the same common questions repeatedly, assembling reports. None of this was the high-value core of the role, and most CSMs will not miss it. The disappearance frees time rather than eliminating jobs, provided the freed time is redirected to work that matters.

What changes

The core CSM work changes in character without disappearing. Health monitoring shifts from manual review to acting on AI-surfaced signals. Preparation shifts from assembling materials to reviewing AI-prepared materials and adding judgement. The CSM becomes an editor and decision-maker on top of AI output, rather than the producer of everything from scratch. This is a meaningful shift in how the work feels day to day, and it favours CSMs who can exercise judgement over those whose value was in throughput.

What gets created

New work appears that did not exist before. Someone has to design the AI-augmented workflows, calibrate the signal detection, decide what gets automated and what stays human, and manage the customer experience of AI-facing interactions. CS operations grows in importance. New roles emerge at the intersection of CS, data, and AI that did not have names a few years ago. The function does not shrink; it changes shape, much as it does at each scaling transition from the previous chapter.

The New CS Leader Skillset

As AI reshapes the work, the skills a CS leader needs shift accordingly. The shift does not abandon the old skills; relationship judgement and organisational design remain central. It adds a layer on top of them.

Knowing what to automate and what to protect

The most important new judgement is about where automation helps and where it harms. A leader who automates the wrong things degrades the customer experience in pursuit of efficiency; a leader who automates nothing leaves leverage on the table and falls behind more efficient competitors. The judgement about which is which, account by account and interaction by interaction, is the central new skill. It cannot be delegated to the vendor, because the vendor's incentive is to automate as much as possible.

Data literacy without data specialism

A modern CS leader does not need to build models, but does need to understand enough about how AI signal detection works to know when to trust it and when to question it. The leader who treats AI output as oracular will be misled by its confident errors; the leader who dismisses it entirely will miss its genuine value. The skill is calibrated trust, which requires enough literacy to judge the system without enough specialism to build it.

Redirecting freed capacity

When automation frees CSM time, that time has to be redirected deliberately, or it dissipates. A leader who automates the administrative work and then does not consciously redirect the freed capacity toward higher-value relational and strategic work has captured the cost saving and missed the larger opportunity. The skill is in seeing the freed time as an investment to be allocated, not merely a saving to be banked.

A Framework for Evaluating AI Vendor Pitches

Every CS technology vendor now claims to be AI-powered, which makes the claim almost meaningless as a differentiator. A leader evaluating these pitches needs a way to cut through to the substance. A few questions do most of the work.

- ***What specifically does the AI do, in one sentence?***

  A vendor who can answer this concretely ("it predicts churn risk by analysing these specific signals") has a real capability. A vendor who answers in abstractions ("it leverages AI to drive customer outcomes") is selling marketing, not capability. The specificity of the answer is diagnostic.

- ***What would we do differently because of it?***

  If the AI capability does not change a decision or an action the team takes, it is not worth paying for, however impressive it sounds. The value is in the changed behaviour, not the capability itself. A vendor who cannot articulate what the team would do differently has not thought through whether their product produces value or merely output.

- ***How do we know when it is wrong?***

  Any AI system makes mistakes. A vendor who has thought seriously about their product has an answer for how errors surface and get caught. A vendor who implies the system is never wrong is either naive or dishonest, and either way the product cannot be trusted in the form they are describing.

These questions favour substance over polish, which is exactly the bias a leader needs when every pitch is polished and only some have substance underneath. The vendors with real capability welcome the questions; the vendors selling hype deflect them.

The Next Decade

Predicting the specific trajectory of AI in Customer Success is a fool's errand, but the broad direction is reasonably clear, and a leader can plan around the direction without betting on the specifics.

The tech-touch end of the customer base will be served increasingly by AI, and served better than it is today. The economics that currently make it impossible to give small customers much attention will shift, as AI makes it possible to provide a genuinely good experience at a scale and cost that human CS never could. This is good for customers and good for the businesses serving them, and it will become a competitive necessity rather than an advantage.

The high-value end of the customer base will remain human, augmented by AI rather than replaced by it. The enterprise CSM of the next decade will spend less time on preparation and administration and more on the relational and strategic work that justifies their cost. The role will become more senior on average, because the junior work that once filled a CSM's day will increasingly be automated, leaving the work that requires judgement and experience.

The function as a whole will not shrink. The pattern across the history of automation is that it changes the composition of work more than the quantity, and Customer Success has enough genuinely human work at its core to remain a substantial function. The leaders who thrive will be the ones who used AI to amplify the human work rather than to replace it, and who kept clear sight of the distinction between the two.

Closing the Chapter

AI is changing Customer Success, but more narrowly and more usefully than the hype suggests. It absorbs the routine and the data-intensive while leaving the relational and the judgement-heavy to people, which means it changes the function substantially without replacing it. The three layers (signal detection, automation, and customer-facing intelligence) mature at different rates and should be adopted in that order, from lowest risk to highest.

The work that disappears is mostly the administrative work no one valued. The work that changes becomes more about judgement on top of AI output. The work that gets created sits at the intersection of CS, data, and AI. The leader's skillset gains a layer concerned with what to automate, how much to trust the systems, and how to redirect the capacity that automation frees. And the vendor noise can be cut through with a few specific questions that favour substance over polish.

This chapter closes Part Two, and with it the operational core of the book. The conclusion that follows steps back from the execution to look forward: what a company that has built Customer Success into its foundations looks like over the long run, and what it costs the companies that do not.

THE BOTTOM LINE

AI **absorbs the routine and the data-intensive while leaving the relational and judgement-heavy to people**, which changes Customer Success substantially without replacing it. The three layers mature at different rates and should be adopted from lowest risk to highest. **The leaders who thrive will use AI to amplify the human work, not to replace it**, and will keep clear sight of the difference.

CONCLUSION

The CS-Driven Company

The companies that win the next decade of subscription business will not be the ones with the best Customer Success teams. They will be the ones where Customer Success stopped being a team and became a way the whole company operates.

This book has spent fifteen chapters on the strategy and execution of Customer Success as a function: how to make the case for it, how to model its economics, how to design it, staff it, run it, and scale it. That framing is correct, and it is where every company has to start. But it is not where the most successful companies end up. In the businesses that get this right, Customer Success eventually stops being a department and becomes something closer to an operating philosophy, embedded in how the entire company thinks about its customers.

This conclusion is about that destination. It is deliberately not a summary; the chapters have made their own arguments and the Bottom Line panels have captured them. Instead, it looks forward, at what a company with Customer Success in its foundations looks like over the long run, and at the slow, compounding cost paid by the companies that treated it as an afterthought.

What It Looks Like When It Is in the Foundations

In a company where Customer Success has become part of the foundations, the function is almost the least of it. The visible CS team still exists and still does its work, but the more important change is everywhere else: in how the product is built, how sales is compensated, how the board spends its time, and how the company decides what matters.

The product team builds with retention in mind, not just acquisition. Features are evaluated partly on whether they deepen the dependence of existing customers, not only on whether they win new ones. The roadmap reflects what keeps customers succeeding over years, because someone in the room always represents that interest and is listened to.

Sales is compensated in a way that aligns with customer success, not against it. The incentive to close a poor-fit customer who will churn in a year has been engineered out, because the company learned that a bad-fit logo costs more than it brings. The handoff from sales to Customer Success is clean, because both functions are measured on outcomes that depend on it being clean.

The board asks about net revenue retention before it asks about new logo growth, because it has internalised which number predicts the company's long-term value. The leadership team treats the existing customer base as the primary asset it is, rather than as a static thing to be defended while the real work happens at the top of the funnel.

None of this is dramatic. There is no single moment when a company becomes CS-driven. It happens gradually, through a hundred decisions that each weight the long-term health of the customer base a little more heavily, until one day the orientation is simply how the company works.

The Difference Compounds

The reason this orientation matters so much over time is the reason that gives this book its title. The advantages of getting Customer Success right do not add up; they compound. And compounding differences, small in any single year, become enormous over the span of years that a company actually operates.

Consider two companies that start in the same place, with the same product and the same revenue. One retains and expands its customer base a little better each year; the other loses a little more than it keeps and replaces the difference with new logos. In year one, the two look almost identical. The difference in their net revenue retention is a handful of percentage points, easy to dismiss as noise.

By year five, the two companies are not in the same business. The first has a base of customers who have grown with it, refer others to it, and cost almost nothing to retain, funding a growth rate the second cannot match at any acquisition budget. The second is on an acquisition treadmill, running harder every year to replace the customers leaking out of the bottom, its growth increasingly expensive and its margins increasingly thin. The gap between them is the compounding of a difference that looked trivial when it started.

> *A few points of net revenue retention, sustained over five years, is the difference between a company that grows effortlessly from its own base and one that runs an ever-faster acquisition treadmill just to stay level. The difference is invisible in any single quarter and decisive over any five years.*

This is the strategic heart of the book, restated at its end. Customer Success is not a cost to be managed or a service to be provided. It is the engine of a compounding advantage that, given enough time, separates the companies that endure from the ones that merely survive.

The Cost of Treating It as an Afterthought

The companies that do not build this orientation pay for it, but they pay slowly, which is what makes the cost so easy to ignore until it is large. There is rarely a crisis that announces the failure to invest in Customer Success. There is just a gradual erosion that becomes visible only in hindsight.

The pattern is consistent. A company grows quickly on the strength of its product and its sales engine, and retention is good enough not to alarm anyone. Customer Success is under-resourced, treated as support with a nicer title, because the growth numbers look fine and the churn is buried in the aggregate. The company keeps acquiring, keeps growing, and keeps not noticing that its net revenue retention is quietly below where it should be.

Then growth slows, as it always eventually does, and the acquisition engine can no longer hide the leak in the bottom of the bucket. Now the company discovers that its customer base was never as healthy as the top-line numbers suggested, that the relationships were thin, that the customers who looked retained were retained by inertia rather than value. The work that would have prevented this had to be done years earlier, when there was no apparent crisis to justify it, which is exactly why it was not done.

The cruelty of the pattern is that the bill comes due at the worst possible moment, when growth has already slowed and resources are already tight. The company that under-invested in Customer Success during the good years has to build it during the hard ones, which is slower, more expensive, and less certain to work. The compounding that could have worked in its favour has been working against it instead.

Where to Begin

If there is a single action this book argues for, it is to start treating the existing customer base as the strategic asset it is, before the numbers force the issue. That does not require building everything in this book at once. It requires only the decision to weight the long-term health of the customer base in the choices the company makes, and then the patience to let that decision compound.

For a founder or CEO, this means asking about retention and expansion with the same seriousness brought to acquisition, resourcing Customer Success as an investment rather than a cost, and making the structural decisions in this book deliberately rather than by default. For a Customer Success leader, it means building the function on the strategic logic of Part One, executing on the operational substance of Part Two, and making the case internally in the language of economics that the rest of the leadership team will hear.

The work is not glamorous. It is the slow, compounding accumulation of customers who succeed, who stay, who grow, and who bring others with them. It rarely produces a dramatic moment. It produces something better: a company whose foundations get stronger every year, built on customers who compound.

THE LAST WORD

Customer Success, done well, is not a department. It is the discipline of building a company on customers who compound. **The advantage it creates is invisible in any single quarter and decisive over any five years**, which is precisely why the companies that understand it start building long before the numbers demand it, and why the ones that do not are left building during the hard years what they should have built during the good ones.

APPENDIX

+

Reference Material

A glossary of every term used in the book, a set of working templates for the artefacts the chapters describe, and a short list of further reading. The glossary is the single source of truth for terminology; where a definition here differs from a casual usage elsewhere, this is the one that was intended.

Glossary

The terms that recur throughout the book, defined once, definitively. Acronyms are given with the full term and a plain-language explanation.

ACV *Annual Contract Value*

The annualised value of a customer's contract. The standard way to compare deal sizes across customers and segments, and the primary input to touch-model and segmentation decisions.

Activation Rate

The percentage of customers who reach the first value moment at all, within a defined window. Distinct from time-to-value, which measures how fast the successful ones got there.

Churn

The loss of customers or revenue. Splits into three structurally different types: involuntary (outside the customer's control), voluntary (an active choice to leave), and downgrade (staying but reducing spend).

CLTV / LTV *Customer Lifetime Value*

The total revenue or margin a customer is expected to generate over the entire relationship. The figure that justifies what a business can afford to spend acquiring and retaining customers.

Cohort Analysis

Grouping customers by a shared starting characteristic (usually the month they signed) and tracking how each group behaves over time, to separate the performance of newer customers from older ones.

CSM *Customer Success Manager*

The individual contributor responsible for a portfolio of customer relationships, accountable for retention, adoption, and (often) expansion within that portfolio.

CSP *Customer Success Platform*

Software built specifically for CS work: customer health views, playbook execution, renewal management, and alert workflows. Worth buying once the function is large enough to need shared workflows.

Expansion

Revenue growth from existing customers, through additional seats, tier upgrades, cross-sell of other products, or increased usage. The primary growth engine of mature subscription businesses.

GRR *Gross Revenue Retention*

The percentage of recurring revenue retained from existing customers over a period, excluding any expansion. Cannot exceed 100 per cent. Measures how much revenue the business holds onto before any growth from the base.

Health Score

A composite measure estimating how likely a customer is to renew, expand, and stay. Built on usage, sentiment, and relationship strength. Real only if it predicts churn and changes what someone does.

LTV:CAC *Lifetime Value to Customer Acquisition Cost*

The ratio of what a customer is worth over their lifetime to what it cost to acquire them. A core measure of the efficiency of a subscription business's growth.

Multi-Threading

Deliberately building relationships with several people across a customer organisation, so the account survives the departure of any single champion. The principal defence against sudden, surprising churn.

NRR *Net Revenue Retention*

The percentage of recurring revenue retained from existing customers year over year, including expansion and minus churn and contraction. Above 100 per cent means the customer base grows on its own. The metric investors watch most closely.

Onboarding
The structured process of taking a new customer from signature to first value. The highest-leverage stage of the customer journey, because the beliefs formed here shape the entire relationship.

PLG *Product-Led Growth*
A motion in which users adopt the product directly, often free, and convert to paid through usage. CS in a PLG context is data-driven and in-product rather than relationship-led.

QBR *Quarterly Business Review*
A periodic structured meeting between the CS team and a customer's stakeholders to review progress, align on goals, and plan ahead. Done well, it is about the customer's business; done badly, it is a status update the customer forgets.

Renewal
The point at which a customer decides whether to continue their subscription. The event where a year of accumulated experience is settled, though the outcome is largely determined by the eleven months preceding it.

Segmentation
Dividing the customer base into groups (usually by revenue, sometimes by strategic value or complexity) so that different levels of service can be matched to different levels of customer value.

Time-to-Value *TTV*
The elapsed time between a customer signing and the moment they first experience the value they bought the product for. The single most predictive metric in onboarding.

Touch Model

The level of human contact a CS function provides to a customer, ranging from high-touch (named CSM, regular contact) through low-touch and tech-touch (automated, no named CSM). Chosen by customer value.

First Value Moment

The point at which a customer first experiences the core benefit the product exists to deliver. Distinct from setup completion, which is technical readiness rather than value.

Tech-Touch

A service model in which customers are supported through automation and in-product guidance rather than a named human CSM. The economically necessary model for lower-value segments, and one that can be genuinely good if designed well.

Working Templates

Starting points for the artefacts the book describes. Each is meant to be adapted to the specific business, not adopted verbatim. They capture the structure; the content is yours to supply.

CUSTOMER HEALTH SCORE WORKSHEET

Usage signals (weight: ___%)

- ☐ Login frequency relative to expected
- ☐ Breadth of feature adoption
- ☐ Depth of use in core workflow
- ☐ Number of active users vs. licensed

Sentiment signals (weight: ___%)

- ☐ Most recent survey or NPS response
- ☐ Support ticket tone and volume trend
- ☐ Mood of recent interactions

Relationship signals (weight: ___%)

- ☐ Number of engaged contacts
- ☐ Seniority of champion
- ☐ Single-threaded or multi-threaded
- ☐ Recency of meaningful contact

QBR AGENDA

→ *Open on the customer, not the product*

- ☐ The customer's goals this quarter and progress against them
- ☐ What has changed in the customer's business

→ *Bring a point of view*

- ☐ Where the product is and is not helping them
- ☐ A recommendation: what we would do in your position
- ☐ Opportunities they have not yet considered

→ *Close on decisions and actions*

- ☐ Agreed next steps, owned by both sides
- ☐ Confirmation that the right people will attend next time

ONBOARDING CHECKLIST

- ☐ Sales-to-CS handoff received: goals, commitments, stakeholder map, risk profile
- ☐ First value moment defined and agreed with the customer
- ☐ Kickoff completed; success criteria documented
- ☐ Value sequence planned (not a feature checklist)
- ☐ Momentum monitored; stalls flagged within days, not weeks
- ☐ First value moment reached and confirmed
- ☐ Transition to ongoing relationship; owner assigned

THE CS SLIDE FOR THE BOARD DECK

Headline metrics

- ☐ Net revenue retention (the number to lead with)
- ☐ Gross revenue retention
- ☐ Logo retention

Composition

- ☐ Churn broken down: involuntary, voluntary, downgrade
- ☐ Expansion as a share of growth

Forward view

- ☐ Renewal forecast for the coming period, by confidence band
- ☐ At-risk revenue and the plans to address it

→ *One slide. The detail lives in the appendix to the deck, not on the slide.*

A note on the board slide: its job is to communicate the health of the customer base in the time a board gives it, which is short. Lead with net revenue retention, show the composition behind it, and give the forward view. Resist the temptation to show everything; a board that wants more will ask, and the detail belongs in an appendix to the deck rather than on the slide itself.

Further Reading

A short, deliberately selective list. The field generates a great deal of writing; these are the categories worth a leader's limited reading time, described by what to look for rather than prescribed title by title, because the best current work in each category changes faster than a printed book can track.

- **The foundational subscription-economy texts**
  The books that made the original case for recurring-revenue business models and the function built to serve them. Dated in their specifics now, but still the clearest articulation of why the post-sale relationship became central.
- **The SaaS metrics literature**
  The writing, much of it from investors and operators publishing online, that established how subscription businesses are measured. Worth reading for the economics in Chapter 4 in more depth than any single chapter can provide.
- **Primary research on retention economics**
  The academic and consulting work on the relationship between retention and profitability. The widely-cited statistics about retention's value originate here; reading the sources is worthwhile before repeating the numbers, as the original studies are narrower than the way they are often quoted.
- **Current practitioner communities**
  The most current thinking on execution lives in practitioner communities and their writing, not in books. For the operational detail of Part Two as the tools and tactics evolve, these are where the field's working knowledge actually accumulates.

A closing note on sources. This book has deliberately favoured durable principles over current specifics, because specifics age and principles compound. The reading above is offered in the same spirit: not a syllabus to complete, but a set of directions to explore as the questions in your own business demand. The best next thing to read is usually the one that answers a question you are actually facing.