Propel

TELEHEALTH · BFCM 2026

The Telehealth BFCM Pre-Flight Checklist

Six decisions that are far easier to make in October than to discover in December, and the send-ramp math behind the first one.

Built for
Telehealth & DTC health brands

Read time
About 20 minutes

From
Propel · trypropel.ai

BEFORE YOU START

Who this is for, and how to use it

This is a planning document, not a theory document. It assumes your BFCM offer is close to locked and your calendar is close to drafted, and that what is still open is everything underneath them.

Most BFCM preparation stops at two questions: what is the offer, and when do the emails go out. For a telehealth brand, those are the easy questions. The hard ones sit in the layer beneath, and they tend to surface as problems only after the sale, when the revenue has already been decided.

A GLP-1 or HRT purchase is a considered medical decision compressed into a short, high-urgency window. The discount can buy attention. It cannot answer whether treatment is right for someone, whether it is safe, or whether now is the moment. That gap between urgency and trust is where a telehealth BFCM is won or lost, and none of it shows up as a line item in a campaign calendar.

How to work through it

Each of the six checks follows the same shape: the question worth answering honestly, what typically goes wrong, and what a stronger plan tends to do instead. Read it with your actual calendar open. The point is not to agree with the framing, it is to find the two or three places where your plan does not yet have an answer.

Check one carries the most math, because it is the one with a hard physical deadline. If your sending cannot reach peak volume in the days you have left, no amount of creative fixes it. That check has its own worked example and a live planner on the page you downloaded this from.

> A NOTE ON HONESTY
>
> Where this document gives numbers, they are illustrative planning defaults, not promises. Deliverability outcomes depend on your sender reputation, list composition, engagement history and infrastructure. Anywhere we have used a rule of thumb, we have said so.

CONTENTS

What is in here

01

Why telehealth BFCM breaks differently

Retail BFCM is an impulse business. Telehealth BFCM is a consideration business wearing an impulse costume.

In a typical DTC sale, the discount is the argument. Someone wanted the thing, the price fell, they buy. The lifecycle's job is mostly logistics: get the message in front of them, remove friction, close.

A telehealth purchase does not work that way. Before someone starts a GLP-1 or HRT protocol they are weighing whether treatment is appropriate for them, whether it is safe, what the first eight weeks actually feel like, what happens if they stop, and whether this provider is credible. A price cut does not touch any of those. It only shortens the time they have to resolve them.

That produces three consequences that shape everything else in this document.

One — the value has to be carried by the lifecycle, not the offer

If the discount is doing all the persuading, you attract the people for whom price was the only obstacle. Those are rarely your best patients. They convert in November and churn in January, and they drag your retention curve into the next year.

Two — volume spikes harder than in retail

Telehealth lists are usually built for steady nurture, not for peak retail sending. The gap between a normal Tuesday and BFCM Friday is therefore much wider, which is exactly the pattern mailbox providers treat as suspicious.

Three — the damage outlives the sale

Your December and January programme is refills, renewals, adherence and win-back. All of it is email. A reputation dented during peak week quietly taxes every one of those sends for months.

> **THE THROUGH-LINE**
>
> Every check that follows is a version of the same question: is this plan spending something in November that you still need in January?

02

Check one
Can your sending reach peak without wrecking your domain?

Most telehealth brands do hit their volume peak during BFCM. The problem is how they get there. They push from 30k or 40k a day up to 150k or 200k in a matter of days, and mailbox providers read that spike as a change in behaviour. Filtering kicks in, complaint rates climb, and a slice of the best sale of the year lands in spam.

It is worth being precise about the mechanism, because it is not punitive. Gmail and the rest are not judging your sale. They are pattern-matching: a domain that has sent 30,000 a day for months and suddenly sends 200,000 looks identical, from the outside, to a domain that has just been compromised. The safe default for them is to slow you down and watch.

The damage does not end when the sale does. A domain that gets burned during peak week carries a weaker reputation into December and beyond, so the refill, renewal and win-back emails you send after BFCM open at a lower rate too. You spend the rest of the quarter paying for one week of over-sending.

> **WHAT A STRONGER PLAN DOES**
>
> Works backward from peak volume and the date it begins, ramping in steps (a common approach is roughly **1.5x every three days**) so you reach peak a few days **before** the sale, and protect the reputation you still need after it.

The reason this check comes first is that it is the only one with a hard physical deadline. Creative can be rewritten the night before. Segments can be rebuilt in an afternoon. A ramp cannot be compressed past a certain point without spending the thing it is protecting, and the days you need are days you either have or do not have. Which is what the next two pages are for.

03

The RAMP method, and the math behind it

Four inputs, one output: whether the calendar can carry the volume.

The planner on the landing page runs the model below. It is deliberately simple, because it answers one question before anything else gets decided: **do you have enough days?**

The four inputs

Current daily sends	Your recent daily average. Not list size, not your best day.
November peak target	The highest single-day volume the calendar will require.
Start date	The first day you can realistically begin increasing volume.
Sale start date	The day the offer goes live, not the day the teaser goes out.

The rule it applies

Volume steps up by about **1.5x**, and each step is **held for three days**, roughly how long it takes for engagement signals from a new tier to settle. A **three-day buffer** at the end puts peak shortly before the sale, not on the morning of it.

Days needed is the number of 1.5x steps from current volume to target, times three. Days available is the gap to your sale date, minus the buffer. The margin between them is the answer.

Reading the verdict

ROOM	Margin is zero or positive. A steady three-day-hold ramp lands on time. Do not accelerate because it looks slow.
TIGHT	Up to six days short. Reachable on two-day holds, which leans harder on reputation. Starting sooner or trimming the peak buys the margin back more cheaply.
RISKY	More than six days short. Compressing to hit target is where a slice of the sale lands in spam. Start sooner, lower the peak, or split across a second warmed domain.

WHAT THE MODEL DOES NOT CLAIM

It illustrates the ramp; it does not guarantee inbox placement. It knows nothing about your sender reputation, complaint history, list hygiene or authentication. A green verdict means the calendar is possible, not that the send is safe.

04

A worked ramp, start to peak

20,000 a day today. 200,000 at peak. What the schedule actually looks like.

Crossing from 20,000 to 200,000 at 1.5x per step takes six steps. At three days a hold, that is eighteen days of ramp, plus a three-day buffer before the sale: **twenty-one days of runway** from the day you start.

STEP	DAYS	DAILY SENDS	INCREASE
1	**1–3**	**20,000**	Start
2	**4–6**	**30,000**	+50%
3	**7–9**	**45,000**	+50%
4	**10–12**	**67,500**	+50%
5	**13–15**	**101,250**	+50%
6	**16–18**	**151,875**	+50%
Peak	**19**	**200,000**	**Peak reached**

Two things are worth noticing. First, the early steps look almost pointlessly small, and that is where most teams lose their nerve and jump ahead. The compounding does the work; skipping step two to "save time" costs you the stable base that makes steps five and six survivable.

Second, the runway is longer than it feels. Twenty-one days before a late-November sale means starting in the first week of November at the latest. If you are reading this in mid-November with a 10x jump planned, the useful move is not a faster ramp, it is a lower peak or a second warmed domain.

> **WHILE YOU RAMP**
>
> Add volume by sending to your **most engaged segments first** and widening outward. A ramp filled with dormant addresses raises volume and lowers engagement at the same time, which is the exact signal you are trying to avoid sending.

05

Check two
Will your audience hear about the sale only a few days before it starts?

BFCM does not begin the day the offer goes live. If the first message about the sale is the sale itself, launch day has to create demand and capture it at the same time, on the noisiest inbox day of the year. You are launching to a cold list on your most important morning.

For a considered medical purchase this is particularly expensive. Someone deciding whether to start treatment is not going to resolve that question between opening your email and the timer running out. They need to have been thinking about it already. Demand-building is not a warm-up act here, it is the part that does the convincing.

Demand-building should start roughly **three weeks out**. A workable shape: teasers and "something is coming" from about the **second week of November**, a named early-access or waitlist push in the **week before**, then launch to a list that is already leaning in.

> **WHAT A STRONGER PLAN DOES**
>
> Opens the sale to a **warmed waitlist** built over the prior three weeks, so day one converts existing intent instead of teaching people the sale exists. A calendar add or early-access signup is a simple way to capture that intent early.

The quiet second benefit

A waitlist is also a deliverability instrument. It gives you a segment that has raised its hand in the last fortnight, which is the healthiest possible audience to be leaning on while your volume is climbing. The demand work and the ramp work reinforce each other, and brands that treat them as one project tend to arrive at peak week with both in better shape.

06

Check three
Are you waiting until BFCM to learn which subject lines and offers work?

A test started inside the sale never reaches significance before the sale is over, so you spend your biggest budget guessing. Peak week is the worst possible time to be discovering what resonates.

The arithmetic is unforgiving. A four-day sale gives you perhaps six sends. Splitting any of them to test a variable means you are running your highest-stakes campaign at half confidence, and the result arrives after the last opportunity to use it. You will have bought a learning at the price of a conversion.

Decide the winners in **October**, on lower-stakes sends: subject-line angles, the hero message, discount framing, and which segments respond to what. By the time peak week ships, it should carry only proven creative and offers.

> **WHAT A STRONGER PLAN DOES**
>
> Treats October as the test window and BFCM as the rollout, so the biggest sends run on **evidence you already have** rather than in-the-moment guesses.

Worth testing in October

- [ ] **Urgency framing versus value framing**
  Does "last chance" or "here is what changes in eight weeks" move your list further?
- [ ] **Discount presentation**
  Percentage, dollar amount, or months-of-treatment framing.
- [ ] **Sender identity**
  Brand, founder, or clinician. This one moves more than people expect in health.
- [ ] **Segment responsiveness**
  Which of lapsed, never-started and active-but-underdosed actually respond to price.

07

Check four
Has the discount become the entire reason to buy?

A GLP-1 or HRT decision is a considered medical purchase, not an impulse buy. A discount can get attention, but it does not answer whether treatment is right for someone, whether it is safe, or whether it is worth acting on now. Lead on price alone and you attract discount-hunters who churn in January while your serious patients feel sold to.

There is a second cost that is easy to miss. Price-led messaging quietly repositions the treatment. A protocol presented at 40% off for four days reads as a commodity, and the patient who was cautiously weighing a medical decision now has to reconcile that framing with the seriousness of the choice. Some of them resolve it by not buying.

> WHAT A STRONGER PLAN DOES
>
> Keeps the **clinical value, the results timeline, and the cost of pausing treatment** carrying the message, with the discount as the nudge that accelerates a decision the patient was already weighing.

A test you can run on your own copy

Take your strongest planned BFCM email and delete every reference to the discount. If what remains still gives a reader a reason to act, the offer is doing its proper job of accelerating a decision. If what remains is an empty template, the discount is carrying the entire argument, and the people it converts will be the people who leave when the next brand undercuts you.

The fix is rarely more copy. It is usually restoring one specific thing the sale edited out: what the first eight weeks look like, what happens to progress if treatment pauses, or who is actually behind the prescription.

08

Check five

Will everyone see the same offer and message for the whole sale?

A flat discount running the full window trains your best buyers to wait, because nothing improves by acting now. Opens and clicks fade after the first couple of days, and the people who did not act on day one are given no new reason to act on day six.

The offer itself should move. For example, hold **30% off** through the main window, then push to **40% off in the final 48 hours** with a clear "last chance" message. That gives the high-intent people who were on the fence a fresh, concrete reason to convert today, and it reliably lifts the back half of the sale, which is usually where the best buyers sit.

> **WHAT A STRONGER PLAN DOES**
>
> Lets the offer and message **escalate toward the close** rather than repeat, so each phase (early access, launch, reinforcement, final push) gives a new reason to act instead of asking the same question for ten days.

The four phases, and what each one is for

PHASE	JOB	WHAT IS NEW
Early access	Reward the waitlist	Getting in before everyone else
Launch	Convert existing intent	The offer is live
Reinforcement	Answer the hesitation	Proof, clinical detail, objection handling
Final push	Close the fence-sitters	A better offer and a real deadline

Note that only two of the four phases lean on price. The middle of a sale is where trust gets built or lost, and it is the phase most calendars fill with a resend of day one.

09

Check six
Do you have backup domains warmed and standing by for the sale?

During peak week, send volume and complaint rates both spike. Even a well-run sale sees more people mark messages as spam simply because everyone is sending more, and that pushes your primary sending domain's reputation down at the worst possible moment. Once it starts landing in spam mid-sale, you are stuck: you cannot warm a fresh domain in the middle of BFCM, because warming takes weeks you no longer have.

That is why the time to prepare is now. Brands that protect the sale keep **two or more backup domains already warmed before BFCM begins**, dedicated to the sale, so if the primary degrades you can shift volume to a domain that is ready rather than watch the revenue leak.

> WHAT A STRONGER PLAN DOES
>
> Warms **2+ dedicated sale domains in the weeks before BFCM**, so a reputation dip during peak week has somewhere to go instead of taking the whole sale down with it.

What "warmed and standing by" actually requires

- [ ] **Authentication complete and aligned**
  SPF, DKIM and DMARC configured on each domain, and verified as passing, not merely present.

- [ ] **Its own ramp, run in parallel**
  A backup domain that has never sent is not a backup. It needs the same stepped warm-up as the primary.

- [ ] **Sent to engaged segments during warm-up**
  Reputation is built by people opening and clicking, not by volume alone.

- [ ] **A decided trigger for switching**
  Agree in advance what deterioration justifies a move, so the call is made on evidence at 6am rather than instinct.

A backup domain is insurance with a lead time. It is cheap to arrange in October and impossible to arrange on the Friday.

10 The ten-week countdown

The same six checks, laid against a calendar. Dates assume a late-November sale.

WHEN	WHAT HAS TO BE TRUE BY THE END OF IT
10–9 weeks out Late Sept	Peak volume target agreed. Current daily average measured honestly. Ramp runway calculated, so you know whether the calendar is possible before anything else is committed.
8–7 weeks out Early Oct	Backup domains bought, authenticated and beginning their own warm-up. October test plan live: subject lines, framing, sender identity, segment response.
6–5 weeks out Mid Oct	Offer architecture decided across all four phases, not just the headline discount. List hygiene done: suppress the genuinely dormant before you start adding volume.
4 weeks out Late Oct	Test results read and winners locked. Primary domain ramp begins. Waitlist mechanism built and live on site.
3 weeks out Early Nov	Teaser sequence running. Ramp at roughly a third of peak. Creative for all four phases in production.
2 weeks out	Named early-access push to the waitlist. Ramp past half of peak. Switching triggers for backup domains agreed and written down.
1 week out	Peak volume reached and held. Everything scheduled. Final-push creative already approved so nobody is writing it mid-sale.
Sale window	**Monitor deliverability daily against the agreed triggers. Escalate the offer on schedule. Change nothing that was not planned, unless the triggers say to.**

IF YOU ARE ALREADY INSIDE FOUR WEEKS

Do these in order: recalculate the ramp against the days you actually have, lower the peak to what the runway supports, and get the waitlist live this week. A lower peak reached cleanly beats a higher peak reached through the spam folder.

11

The pre-flight checklist

Print this page. If you cannot tick a box, that is the work.

- [ ] **We know our current daily send average and our November peak target as numbers**
  Not list size, not best day. The real recent average.
- [ ] **We have calculated the ramp runway and it fits the calendar**
  Days needed versus days available, with a buffer before the sale.
- [ ] **The ramp adds engaged segments first and widens outward**
  Volume and engagement climbing together, not in opposite directions.
- [ ] **Two or more backup domains are authenticated and already warming**
  SPF, DKIM, DMARC verified as passing. Their own ramp underway.
- [ ] **We have written down what triggers a switch to a backup domain**
  Decided in advance, in daylight, not at 6am mid-sale.
- [ ] **Demand-building starts about three weeks before the offer goes live**
  Teasers from mid-November, early access the week before.
- [ ] **There is a live mechanism capturing early intent**
  Waitlist, early-access signup or calendar add, already on site.
- [ ] **Subject lines, framing and segment response were tested in October**
  Winners decided before the stakes got high.
- [ ] **Our strongest email still gives a reason to act with the discount deleted**
  Clinical value carrying the message, price as the nudge.
- [ ] **The offer escalates across the window instead of repeating**
  Four phases, each with something genuinely new.
- [ ] **The final-push creative is approved before the sale opens**
  Nobody is writing the most important email while it is happening.
- [ ] **Someone owns deliverability monitoring daily during the sale**
  A named person, with the triggers in front of them.

12 What to measure, and how to read it honestly

A plan you cannot read during the sale is a plan you cannot correct. But peak week also produces the least trustworthy numbers of the year, so it is worth agreeing in advance which metrics you will actually act on.

Open rate deserves a warning label

Privacy pre-fetching, most visibly Apple Mail Privacy Protection, registers opens that no human performed. That inflates reported open rates, and the inflation is not evenly distributed: it varies with the mix of mail clients in whichever segment you happen to be sending to. During a ramp, that mix is changing every few days by design.

This does not make open rate useless. It makes it a **relative** measure. Read it as a delta on a consistent measurement basis, comparing like segments week to week, and never quote it as an absolute. Any uplift figure worth acting on, including any figure we would show you, should be stated that way.

What to watch instead, daily

SIGNAL	WHY IT MATTERS DURING A RAMP
Complaint rate	The metric mailbox providers weigh most heavily. Rising complaints during a ramp is the clearest early signal to hold the current step.
Bounce rate, hard and soft	A soft-bounce climb as volume rises often means throttling has begun, before placement visibly degrades.
Click rate by segment	Harder to inflate than opens. The honest read on whether the newly added segments are actually engaged.
Placement, where you can see it	Seed testing or provider postmaster data. Imperfect, but directional, and better than inferring from opens.
Revenue per thousand sent	The number that survives all of the above. If it falls while volume rises, the ramp is outrunning the list.

DECIDE THE THRESHOLDS NOW

Write down the complaint rate and bounce rate at which you hold the ramp, and at which you move to a backup domain. Thresholds chosen calmly in October get followed. Thresholds invented mid-sale get argued about.

13

Five ways this goes wrong

Every one of these is a decision that looked reasonable at the time.

One — the ramp starts on time and then gets skipped

The early steps feel slow, a campaign gets added, and two steps are merged to catch up. The compounding that made the plan work is exactly what breaks when a step is missed. If you are behind, lower the peak rather than widen a step.

Two — the ramp is filled with the wrong people

Volume is found by reactivating the dormant, because that is where the addresses are. Volume rises, engagement falls, and the signal you send is the worst of both worlds. Widen outward from your most engaged, and suppress the genuinely dead before you start.

Three — the waitlist is built but never used

Early access gets collected and then receives the same launch email as everyone else. The promise was the reward, and breaking it costs you the segment you most need engaged while volume peaks.

Three and a half — the backup domain exists on paper

It is bought and authenticated but never warmed, so when the primary degrades there is nowhere to go. A cold backup domain under peak-week load fails faster than the domain it was meant to rescue.

Four — the escalation is announced in advance

Telling everyone on day one that it becomes 40% off on Friday guarantees nobody buys at 30%. The escalation works because it is a new reason to act, not a published schedule.

Five — nobody is watching on the days that matter

The sale is scheduled, the team goes quiet for the holiday, and a placement problem that started Wednesday is discovered Monday. Someone owns the dashboard daily, including the days off.

14

If you only do three things

Ranked by how expensive they are to fix later.

01

Calculate the ramp this week

It is the only check with a hard physical deadline, and it constrains everything else. If the runway does not fit, you want to know while lowering the peak is still a calm decision rather than a concession.

02

Start warming backup domains now

The lead time is the whole product. A domain started in October is insurance; the same domain started in November is a line item you cannot use. It is the cheapest protection available for the revenue at stake.

03

Get a waitlist live before the teasers

It converts your launch day from demand creation into demand capture, and it hands your ramp a recently engaged segment at exactly the moment volume is climbing. One mechanism, two problems solved.

AND ONE THING TO STOP

Stop planning to test during the sale. Move every open question into October, and let peak week carry only what you already have evidence for.

Everything in this document is easier to fix in October than to discover in December. The six checks are not exotic. What makes them hard is that none of them appear as a line item on a campaign calendar, so they get planned last or not at all, in the weeks when there is least room to change anything.

WORK WITH PROPEL

Want this pressure-tested against your real numbers?

We will run the same six areas against your actual programme and calendar, and show you exactly what we would pressure-test before the biggest sends go out.

- [ ] **Your real send-ramp math**
  Against your BFCM dates, your sender reputation and the days you actually have.
- [ ] **Whether your backup domains are ready**
  Authenticated, warmed, and able to take load if the primary degrades.
- [ ] **Where demand, testing and offer progression have gaps**
  The three weeks before launch, and the four phases inside the window.
- [ ] **How patient value is carrying through the discount**
  Or getting lost behind it.

Get my detailed BFCM audit

A working session, not a pitch. You leave with the areas worth fixing either way.

Book the audit →

trypropel.ai/contact-us · Lifecycle & retention for telehealth & DTC health brands