

BPO HIRING GUIDE

Funnel Math

WHERE CANDIDATES DROP OUT AND HOW TO FILL EVERY CLASS ON TIME



Table of Contents

To Fill a Class Is to Forecast	03
Funnel Math 101: The Forecast Is Division	04
The Candidate-to-Hire Ratio	05
The Two Kinds of Drop-Off	07
Plugging the Leaks	08
Metrics That Matter	09
The Funnel Doesn't End at Day 1	10
Quality Improves With the Right Machine	10
Common Funnel-math Mistakes	11
References & Methodology	12



To Fill a Class is to Forecast

Class-based hiring changes the job from “find good people” to “find enough good people by a specific date.” That one change turns hiring into arithmetic.

Every training class has a fixed number of seats and a fixed start date. That gives you exactly two ways to miss: show up short, or show up on time with the wrong people who wash out through training and the first 90 days. Either way, you're backfilling seats you already celebrated filling. Underneath every class are two numbers you actually control:



Volume — how many. How many candidates you put into the top of the funnel. Too few, and the math can't reach your seat count no matter how clean the process is.

Velocity — how fast. How quickly you move candidates through it. Move slowly and good people go cold and leak out before Day 1, for reasons that have nothing to do with fit.

Almost every missed class is a volume problem, a velocity problem, or both — and they look identical from the top of the funnel while costing very different amounts to fix. Knowing your stage-by-stage rates is what lets you tell them apart weeks before a start date, while there's still time to act.

What's ahead, in order. The equation that turns seats into an applicant target:

- The candidate-to-hire ratio that sets your volume
- The two kinds of drop-off that set your velocity
- The highest-leverage ways to plug the leaks
- And why the funnel doesn't end on Day 1 — because a full classroom isn't the same as a full floor at day 90.

The Forecast is Division

The math on this page is simpler than it looks. Before the formula, here are a few terms to know.

Picture hiring as a funnel: a large group applies at the top, each step lets some people through and screens others out, and a much smaller group actually starts on Day 1.

- **Stage** — one step in the process. Applying, finishing a pre-screen, completing an assessment, interviewing, accepting an offer, and starting on Day 1.
- **Conversion rate** — the share of people who make it from one step to the next. If 100 reach the pre-screen and 50 finish it, that step converts at 50%. Your stage-by-stage conversion rates are just that percentage for every step.
- **Applied-to-start yield** — the bottom line: of everyone who applies, the percentage who make it all the way to a Day-1 start. You get it by multiplying the stage rates together.

Start with the only two inputs that matter: how many seats you need, and by when. Everything else is a division problem. Multiply your stage conversion rates together to get your applied-to-start yield, then divide the seats you need by that yield.

Applicants needed = Seats ÷ Overall Yield

Your overall yield is each stage's pass-through rate, multiplied together

Example 25-seat class

Using conversion rates from a real high-volume program, here's the full funnel from applied to seated:


Pipeline Stage	Conversion	Candidates Needed	Dropped at this Stage
Sourced/ Applied	--	~250	--
Pre-screen complete	50%	~125	125 dropped
Assessment complete	60%	~75	50 dropped
Interview/ offer	50%	~36	39 dropped
Offer accepted	70%	~25	11 declined
Day-one Start (after fallout)	85%	~21	4 no-shows

Multiply the rates: $0.50 \times 0.60 \times 0.50 \times 0.70 \times 0.85 \approx 8.9\%$ applied-to-Day-1 yield. Notice the example lands at 21 starts, not 25. That's normal offer-to-start fallout. To actually seat 25, run the division: $25 \div 0.089 \approx 280$ applicants, call it 300 with a buffer. Your other option is a higher yield. Both are valid levers; only one of them costs more every single month.

Here is the same funnel as a bar graph. Each bar is what survives to the next step — the shape of the division above, and the shape of every class you run.

The funnel narrows fast — that's the division at work

Candidates remaining at each stage, 25-seat class



STEP ONE IS A CONVERSATION, NOT A SOURCING SPRINT

The first thing to do with your yield isn't to source harder — it's to tell the client what's realistic. "To seat 25 on March 4, we need about 300 applicants in the top of the funnel by January 21." That's a commitment the math can actually keep. Agreeing to a date your funnel can't reach just moves the miss to the most expensive possible moment: the week the class was supposed to start.

The Timeline is Addition

The division tells you how many. The date takes addition.

Pull the median time-in-stage for each step from last quarter's data — the same per-program cut you used for your conversion rates. Add them and you have the median applied-to-hired journey: how long a successful candidate actually spends in your funnel.

A median means half your candidates take longer — and you're not seating one candidate, you're seating twenty-five. Plan around the 75th percentile, or add 20% if you don't track percentiles yet.

Finally, add sourcing runway. Three hundred applicants don't appear the day the requisition opens; they accumulate at whatever rate your channels produce. Divide your applicant target by your weekly applicant flow to get the weeks the top of the funnel needs to be open before the first candidate can even start the journey.

Requisition-open date = Class start – (Sourcing Runway + Candidate Journey)

Sourcing Runway= How long (in days) it takes for the candidate to enter the funnel

Candidate Journey= The time (in days) from applied to hired

Example:

Stage medians summing to 24 days, padded to 28 at the 75th percentile. An applicant target of 300 at 150 applicants per week is another 14 days of runway. That's 42 days — count back from a March 4 start and you land on January 21. The date in the client conversation isn't a guess; it's addition. And it gives you a hard tripwire: for every week you open the req past that date, the math forces you to either lower the bar or miss the class.

The Candidate-to-Hire Ratio

The equation only works if you know your rates — so start by measuring the one ratio that drives your volume: **how many applicants it takes to make one hire**. Here is a real, high-volume BPO program — a global customer-service requisition, trailing twelve months — normalized to 1,000 applicants and shown against its actual counts:

Stage	% of all applicants reaching this stage	Per 1,000 applicants	Actual Count
Applied	--	1,000	68,701
Reached pre-screen	96%	964	66,253
Advanced to assessment	51%	509	34,997
Reached recruiter interview	12%	119	8,171
Hired	--	58	3,980

Real program data, anonymized. Endpoint is hired — not Day-1 start, and not still-in-seat at day 90 (page 10 shows what happens next). Figures rounded.

That's roughly **17 applicants for every hire** in this program — and it is not a benchmark you can borrow. A second BPO program in the same industry ran closer to **27 applicants per hire** and leaked in a completely different place: most of its loss happened in the early application steps, before the funnel proper ever saw the candidate. Two programs, same industry, roughly 2× apart, leaking in different spots. Only your own stage-by-stage rates can tell you where your money is going.

Not all of that loss is bad. Some of it is the funnel doing its job — screening out people who shouldn't be hired is the point. The expensive loss is the part that isn't selective: strong candidates who leak out during slow scheduling, long assessments, and silent gaps. That loss is a velocity problem, and it's what the next page is about.

The Two Kinds of Drop-Off

Candidates who aren't screened out leave your funnel for two reasons — and the fixes are different for each


- **Step drop-off — friction.** The parts of the process people don't want to do: a long application, a heavy assessment, a scheduling back-and-forth.
- **Time drop-off — waiting.** Even candidates who would finish go quiet the longer they wait between steps. This one is invisible in your stage and counts until the day nobody shows up.

The clock starts at first contact. The published research on lead-response time (Oldroyd, McElheran & Elkington, 2011), extended into recruiting by practitioners, is unambiguous: how fast you respond is one of the strongest predictors of whether a candidate converts.

Time to First Contact	What It Costs You
Within 1 Hour	Highest Conversion: the candidate is still engaged
1-4 Hours	Moderate decline: they may have started another application
4-24 Hours	A 30-50% Reduction: they've likely applied to 2-3 other employers
48+ Hours	Most have moved on: outreach feels cold

The way to cut drop-off is to be fast

Share of candidates still actively responsive, by days since first contact





THE TWO-SPEED FUNNEL

In our example program, candidates who got hired moved from application to hire in under six days on average. Meanwhile, the median candidate parked in the assessment step had been sitting there for 47 days, deep in the dead zone of the curve above, where fewer than one in five candidates is still engaged. Same requisition, two different funnels: the fast one produces hires; the slow one accumulates inventory that has already evaporated. If your pipeline counts look healthy but your classes aren't filling, this is usually why.

Plugging The Leaks

Volume fills the top of the funnel; velocity determines how much survives to the bottom. The highest-leverage fixes all share one property: they **remove manual steps without removing rigor**. Speed and quality are only a tradeoff in a manual process.

Where the Leak is	What's Usually Causing It	The Fix
Application → first contact	Recruiters batching outreach; first touch measured in days	Automated acknowledgment with an immediate next step, within the hour
Application completion	Long early steps candidates abandon before the funnel proper begins	Shorten and mobile-optimize the early steps; move heavier asks later
Scheduling	Back-and-forth coordination while candidates go cold	Self-scheduling for screens and interviews — candidates book instantly
Assessment stage	Long, late, or manually scored evaluations — where Program A's median parked candidate sat for 47 days	Shorter, validated assessments with instant scoring, placed early
Offer → Day 1	One to three weeks of silence; 10–20% offer fallout	A pre-boarding cadence: welcome within 48 hours, weekly touchpoints
After Day 1	Hiring decisions invisible to training and ops	Track cohorts to 30/60/90 days and feed results back into screening

One caution as you tune: don't plug leaks by lowering the bar. Advancing marginal candidates to fill Monday's seats trades a visible problem — an empty seat — for the far more expensive invisible one on the next page. Filter before Day 1, not after.

Anecdote: What Removing Manual Steps Actually Buys You

A healthcare-focused BPO with 20,000+ agents needed to hire thousands of licensed insurance agents on tight contract deadlines — high volume with regulatory complexity on top — without growing its recruiting team. It rebuilt the workflow around automation: a short pre-screen, auto-triggered job simulations and job-fit assessments, self-scheduled interviews, and a feedback loop connecting hiring data to post-hire outcomes.

The result was a 65% reduction in hiring costs and a 50% improvement in time-to-hire — velocity gains that came from cutting manual work, not from lowering the bar. More broadly, Journeyfront clients typically see a 50%+ increase in effective recruiter capacity, because automating screening, scoring, scheduling, and status updates frees recruiters to spend their time only where human judgment adds value.

Metrics That Matter


You can't manage a funnel you only look at monthly. Track these five operational metrics weekly to keep classes on schedule:

Metric	What to track	Why it matters
Pipeline fill rate	Candidates at each stage vs. target, by class, against deadline	The core early warning — flags a class in trouble weeks before start date
Time-in-stage	How long candidates sit at each stage	Finds the bottlenecks where candidates go cold — the 47-day parking lot
Offer-to-start fallout	% of accepted offers that don't show on Day 1	Directly shrinks class size; signals pre-boarding gaps
Recruiter utilization	Active pipeline volume per recruiter, across programs	Catches overload and drift before they become missed targets
Class fill forecast	Projected fill from pipeline and historical conversion	Lets you intervene weeks before a miss, not days

Then extend measurement past Day 1, by cohort: training completion, 30/60/90-day attrition, time-to-proficiency, early quality scores, and source effectiveness tracked all the way to day-90 outcomes — because two channels can each deliver 500 applicants and produce wildly different agents at day 90. The turnaround on page 10 wasn't found by intuition; it was found by exactly this kind of cohort measurement. Put the reviews on the calendar — five minutes per active class weekly, a cohort review monthly — because metrics nobody looks at are decoration.

The Funnel Doesn't End at Day 1

A full classroom on Day 1 is not the finish line — it's the starting line. Here is when attrition actually happens, from a global BPO program across twenty-four months of hire cohorts:




Three out of four people who leave are gone within 90 days of starting. The funnel through Day 1 tells you whether your machine generates enough volume; the funnel through day 90 tells you whether it generates enough value — and most reporting stops exactly where the expensive losses begin. At a typical sunk cost of \$4,000–\$17,000 per early quit, the seats you celebrated filling can quietly become the most expensive hires of the quarter.

Quality Improves With The Right Machine

Here's what happens when one program starts measuring past Day 1 and tunes its upstream funnel against those outcomes — same job family, cohorts compared through equal 90-day exposure windows:

90-day retention rose ~25 points — while volume hit an all-time high

90-day retention, earliest cohorts vs. most recent fully-observed cohorts



Same program, same job family. Only cohorts old enough to have fully lived their first 90 days are included.

Retention improved by roughly 25 points **while hiring volume scaled to its largest cohorts ever** — the exact condition under which quality is supposed to collapse. The difference wasn't slower hiring; it was measurement and tuning: connecting post-hire outcomes back to the upstream funnel, cohort by cohort, and adjusting the screens against what actually predicted staying. At \$4,000–\$17,000 per early quit, a swing that size is worth a mid-seven-figure annual gain for a large program — value invisible to any scorecard that stops counting on Day 1.

What Good Looks Like

Use these as starting points for diagnosing where your machine is breaking down — the improved cohorts above cleared them. Treat them as directional, not definitive; every BPO differs by size, industry, role complexity, and geography.

Benchmark	Target
Class fill rate	95%+
Offer-to-start conversion	85%+
Training completion rate	90%+
90-day-retention	75%+



[Access The Worksheet Here](#)

Open the worksheet and create a copy to input your programs data.

Common Funnel-Math Mistakes

- **Optimizing cost-per-hire** instead of cost-per-productive-agent. The cheap hire who quits in week 3 is the most expensive hire you'll make this quarter.
- **Sourcing to an exact target.** A funnel with zero buffer fills classes only when nothing goes wrong — and something always goes wrong.
- **Averaging conversion rates across programs.** A blended rate hides the program that's bleeding. Run the math per program, per market.
- **Measuring volume instead of conversion.** Application counts can rise while yield falls; rates are the early warning, counts are the lagging headline.
- **Stopping the funnel at Day 1.** Three-quarters of attrition hits in the first 90 days — downstream of the only milestone most dashboards celebrate.



Get the full BPO Hiring Playbook

Written by Daniel Ash, founder of Journeyfront, drawing on a decade of work with the largest BPO hiring operations in the world.

To request your copy or learn more, email sales@journeyfront.com

References & Methodology

Oldroyd, J. B., McElheran, K., & Elkington, D. (2011, March). The short life of online sales leads. Harvard Business Review, 89(3), 28. <https://hbr.org/2011/03/the-short-life-of-online-sales-leads>

About the data in this guide. The candidate-to-hire funnel (page 6) and the attrition and retention exhibits (pages 9–10) are drawn from real, anonymized high-volume BPO programs running on the Journeyfront platform. The primary program is a global BPO's evergreen customer-service requisition observed over a trailing twelve-month window; the comparison figure references a second program running offshore English-language accounts observed over nine months. The 20,000-agent example on page 9 is a licensed-healthcare BPO program. Funnel endpoints are hired unless otherwise noted. Retention cohorts are compared through equal exposure windows — only cohorts old enough to have fully lived the measured period are included. The lead-response and engagement-decay figures (page 7) and the example (pages 4–5) are illustrative estimates, not measurements. All figures are rounded. Your programs will differ — substitute your own rates at every stage before relying on any number here.