




# The recruiting leader's 8-step guide to capacity planning



# Table of Contents



Introduction .....	3
1. Get Involved in the Headcount-Planning Process .....	6
<b>Thoughts from Viet Nguyen, Customer Talent Advisory at Gem .....</b>	<b>8</b>
2. Separate Out Roles .....	9
3. Determine Historical Productivity.....	11
<b>Thoughts from Mike Moriarty, Talent Advisor @ Human Capital .....</b>	<b>13</b>
4. Use Historical Productivity to Forecast Future Performance .....	14
5. Factor in Attrition and a Buffer .....	15
<b>Thoughts from Bernardo Sosa, Director, Talent Sourcing, Ops &amp; Programs @ Box.....</b>	<b>16</b>
6. Check in with Your Leaders.....	18
7. Calibrate the Gaps.....	19
<b>Thoughts from Joel Torres, Talent Acquisition Manager @ Twilio.....</b>	<b>20</b>
8. Validate and Socialize Your Model .....	22
How Gem Can Help.....	23

# Introduction

---

Below, we'll walk you through building a model that uses historical hiring data to determine if you're properly-resourced to meet imminent headcount needs, recalibrate with functional leaders, determine how to prioritize new resources or accelerate your productivity-per-resource (PPR), and more.

**If there's an ongoing, mission-critical question being asked on every talent acquisition (TA) team, it's this: Are we currently on track to hit our hiring goals this quarter? Every member of your team is asking themselves this daily. Yet as a talent leader, your version of this question likely looks a bit different: How do our headcount goals compare to the current capacity of our recruiting team? In other words, it's one thing to ask if you're meeting your goals; it's another to ask if you're even equipped to meet them.**

Capacity planning, or capacity modeling, is the means through which an organization determines how much work it can do given employee headcount and time constraints: does current supply meet current demand? In a recruiting context, it's ensuring you have enough resources on your TA team to meet the hiring needs of the business. It's the passage from a transactional, reactive model to a *proactive* one. It ensures you don't say 'yes' to every new headcount request that comes your way without *knowing* you have the resources to take it on—a reflex that risks overburdening an already-stretched-thin team and sets you up to miss those ever-changing goals, because you didn't have a forecast in place to flag the point at which you'd need new resources.

Capacity planning allows you to say to your CTO, when she tells you the org now needs to hire 24 more engineers next quarter: "Let's look at the data I shared with you in Q1 about how many reqs each of our recruiters can handle per month, based on historical metrics. *Here's* how many recruiters it would take to hire 8 engineers per month. We have x many recruiters now, all working at 90% capacity. Those 24 new eng hires would put our current team at well over 100%. I want to help you hire that headcount; but we'll need x much time and x many more resources to do so."






That response is *still* a 'yes'! But it's data-driven rather than reactive; it doesn't over-promise and under-deliver; and it hinges on the business understanding that TA requires resources to do what it does.


The best capacity model looks in two directions: backwards at historical data (how many engineers can a recruiter hire in a quarter?) to accurately forecast what your org can do in the future. Keep in mind that your capacity model will be more-or-less complex depending on the size of your org, or whether this is your first time modeling or you're iterating on a model that's already in place. Know, too, that there's no *one* way to build a capacity plan. Below, we'll walk you through building a model that uses historical hiring data to determine if you're properly-resourced to meet imminent headcount needs, recalibrate with functional leaders, determine how to prioritize new resources or accelerate your productivity-per-resource (PPR), and get the business on board with your model.

You can do this on a spreadsheet or with the data your recruiting solutions automatically pull for you. You can begin it during your next headcount planning cycle, or *off-cycle* if you're trying to reconcile your gap-to-goal and understand where you need more resources. Capacity planning is an iterative process. If you're working toward a data-driven model that's even *starting* to meet your needs, you're already headed in the right direction.




### The benefits of creating a capacity model

- **Appropriate goal-setting:** better manage hiring managers' and executives' expectations about what's possible
- **Building credibility:** not overpromising on unrealistic hiring expectations and then missing goals
- **Eliminate the consequences of poor hiring quality due to req overload:** skipped intake meetings, no passive talent sourcing, poor candidate experience, etc.
- **Eliminate burnout, improve morale, and decrease turnover** since recruiters aren't working above capacity, burning candles at both ends
- **Discover what's working and what's not in your hiring process,** better understand your sourcers' and recruiters' skill sets and strengths, create better budgets, benchmark your own performance year-over-year, and ultimately hire more intelligently



# 1. Get involved in the headcount-planning process

You'll know that you're viewed as a strategic member of the "TA—Finance—functional leaders" trifecta when you're meeting with these stakeholders every other day during the headcount planning cycle. Your goal in these conversations will be to help Finance and leadership understand what your team needs in order to support their respective agendas.



It's a great practice to start with "the knowns"—business goals and accompanying headcount—and work backwards from there. (For you leaders who've been tracking productivity-per-resource, this historical "known" is *also* a great place to start. But if it's your first time building out a capacity model, headcount is as fine a place to start as any.) Here's why: if you know in December what your company's revenue goal is for next year, you'll understand why your Sales org says they need a team of x to hit that goal. And if the revenue goal changes mid-year, you can predict headcount shifts in the Sales org before anyone even stops by your desk to tell you. The same goes for product roadmap plans and engineering headcount: you stay ahead of the curve. So pay attention to broader business goals and initiatives as they pertain to product, to revenue, and to accompanying headcount.

Ideally you'll be a part of headcount conversations as they're happening. If you're not yet, start building relationships with Finance and with functional leaders (i.e. Head of Sales or Head of Eng if you're a smaller org, Directors and VPs if you're larger). You'll know that you're viewed as a strategic member of this trifecta (TA—Finance—functional leaders) when you're meeting with them every other day during the headcount planning cycle.

Your goal in these conversations will be to help other stakeholders understand what your team needs in order to support their respective agendas. If headcount negotiations happen *without* you, you run the risk of Finance and leadership assuming that whatever headcount targets they land on, you're capable of hitting. So think of your presence in these conversations as a form of checks-and-balances, *as well as* a desire to understand their goals.

If you're just starting to build relationships within this trifecta, we recommend starting with Finance. Talking to each functional leader individually can get messy, because you may get wishlists of headcounts that haven't been approved. Partnering with Finance *first* may mean asking them to go have conversations with functional leaders early in the planning cycle. Through those conversations, Finance will come up with not only a headcount *number* (50 CSMs); but also pacing (5 CSMs in Q1, 20 in Q2, 25 in Q3) and location (U.S.? International? Remote? etc). This is the granularity you'll need when it comes to inputs for your capacity model.

Conversations you might have with functional leaders in the meantime might include these questions—though remember there's a difference between a *wished-for* headcount and an *approved* headcount:

- What are your department-specific goals for the coming year, and what will your headcount requirements be in order to meet those goals? Which reqs are currently approved/open, and which aren't yet?
- What specific skill sets will you need in order to meet those goals? What are the skills gaps on your current team?
- What do the timelines look like on current and expected reqs? In which quarters do you need those hires?
- Over the next 6-12 months, who's getting promoted? Who do you think might leave or be let go? Overall, what do you expect attrition and retention to look like?
- How do you expect the business to impact you over the next 12 months? How do you expect to impact the business?



## Thoughts from Viet Nguyen

Customer Talent Advisory

One of the reasons capacity planning can be tricky is that recruiting teams often don't get their headcount targets until after the fiscal year begins. So if your team doesn't get their targets until February or March, you've got 10 or 11 months to do 12 months of work. *Then* it takes 2-3 months to make the new hires you'll need to hit the targets you've been given, and another 2 months to onboard. When all is said and done, you'll only get about 6 months of production from your new hires in a given year—a ratio you'll have to account for in your capacity model.

This is why, if you're a TA leader, strong relationships with Finance are essential: to encourage them to have those conversations with functional leaders early, understand headcount needs, and send over targets on a timeline that gives you as much of the year to work with as possible. If there *isn't* a solid foundation between TA, functional leaders, and Finance, headcount negotiations will happen between those two groups without you. You'll get pulled in once they've landed on numbers that feel right to them, and they'll be like, Hey, you can hit these numbers, right? This is the rat race every recruiting team is in until they put their foot down and say, Hey, I need to be a part of this conversation from the beginning. Because there are resources my team needs in order to hit our goals.

What I'm trying to answer with a capacity model is how many recruiters do I need to be able to meet the needs of the business? Use two years of data where possible. You may not always have it, and a good workaround is to identify a few recruiters you can use to model after. If Jane is an ideal recruiter for sales and she produced 25 hires last year, I'm making the assumption that she'll work in the same fashion next year. Of course it's important that I know Jane's been working at a reasonable capacity, otherwise I'm not modeling for sustainability. But I'll *also* use her numbers to assume that a standard sales recruiter should be able to hire 25 people, year-over-year. Once you get there, it can be incredibly valuable to leverage your network to sanity-check your numbers.

If you've built a balanced team, that 25-hires-per-recruiter will average out. One person might hire 27, another 23. That's not to say don't model in recruiter experience. In fact, the larger your org gets, the more complex your calculations might be, the more dimensions they might contain. Maybe you ultimately say: We can expect x number of hires from our senior recruiters, x number of hires from our junior recruiters. But you don't want to unnecessarily complicate things from the beginning if this is your first time building a model.

*Definitely* factor in attrition—at least an annual number. If you want your model to be even more accurate, factor it in by quarter. Factor in a buffer for unforeseeable circumstances. The reality is that you're losing resources both on the recruiting team and in the broader org *even as* you're hiring for the year. 10%-15% are good averages for both; but HR should be able to give you historical data on attrition to work with. It will mean that your recruiters are operating under capacity sometimes. But you actually want that. You want to have that room to be agile, to scale and expand where you need to.

10%-15% are good averages for both attrition and a buffer. It will mean that your recruiters are operating under capacity sometimes. But you want that room to be agile, to scale where you need to.





## 2. Separate out roles, and consider assigning weights to them

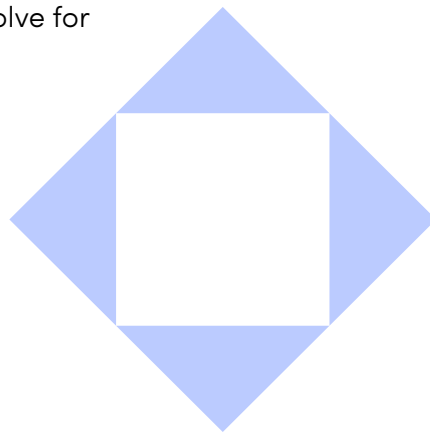
Most organizations we work with at Gem keep this step pretty high-level. Maybe you separate tech roles and non-tech roles, since tech roles are generally twice as hard to fill as non-tech roles are. Maybe you separate out executive roles and university hires, since they're *more* and *less* difficult to hire for, respectively, than your typical ICs or entry-level managers. Solve for what's left *first*.

Not all roles are equally hard to fill; and for the sake of your capacity model, you shouldn't treat them as though they are. You know as well as we do that you can't simply take the final headcount Finance gives you and assign each of your recruiters the same number of reqs: ICs working more high-volume roles will look like they're killing it, while those who get harder-to-fill roles will experience drops in morale. Some reqs (niche roles, roles your org has never hired for before) are simply more demanding than others. So treat them as such.

As we've mentioned, a best practice if this is your first time building out a capacity model is to start simple. In fact, most organizations we work with at Gem keep this step of the process pretty high-level. Maybe you separate tech roles and non-tech roles, since tech roles are generally twice as hard to fill. Maybe you separate out executive roles and university hires, since they're *more* and *less* difficult to hire for, respectively, than your typical ICs or entry-level manager roles. Even starting with *these* distinctions will ultimately get you more clarity around productivity-per-resource.

A more complex way to ensure an even workload is to assign each headcount a weight by difficulty. Beyond the difficulty of the particular role, you might consider these criteria:

- **Source-of-hire.** For example, there's a difference in passthrough rates—and ultimately in offer-accept rates—between active and passive talent. *You'll* determine how much you want the sourcing function to contribute to the total number of offers your org extends; but take conversion rates by source into account.
- **How new the role is to the org.** If your team is hiring for a role it's never hired for before, the recruiter and hiring manager may need more meetings than usual to align, and a lengthier kickoff to ensure clarity about the role and its impact.
- **Sourcer/recruiter experience.** This may be as simple as: if a recruiter is ramping up, they're only expected to hire at 50% capacity. Further down the line, you may have different targets, say, for L2s and L4s—but again, this is *down the line!*



“At Dropbox we messed around with a weighting system, though we only ended up using it for performance reviews. But ultimately it could be a resourcing model. I’d be like, I need a thousand points across this org, so every recruiter should get 30 points on average. What I like about that model is it incentivizes the recruiter to own their own destiny. They might say, Okay, I’m going to hire a bunch of three-pointers, but I need 10 of them. Or I can challenge myself with some hard-to-fill reqs, but fewer of them.”

**Mike Moriarty**

Talent Advisor @ Human Capital

- **Location.** Historical data will tell you if your offer-accept rates are lower in New York than in Denver. You might take location-based conversion rates into account if you’re hiring in a handful (or more) of geos.
- **Which hiring managers are on which roles,** and their historical responsiveness and engagement.
- **Bottlenecks in specific pipelines.** Hiring manager responsiveness is one example of this. Solutions that track historical data will show you where the bottlenecks are in certain funnel stages for certain roles. Are candidates dropping out more often at a certain stage for one of your roles than for others? If you haven’t worked out *why* yet, take those time-to-fill differences into consideration.
- **Seasonality.** Are some roles more difficult to hire for in some quarters than in others?

Other considerations might include diversity requirements for certain roles, which can sometimes slow down searches; what the talent market looks like for a given role right now; or company reputation (is big news forthcoming which may make it easier to fill roles from an employer branding perspective?). Taking these criteria into consideration, figure out the point system that makes the most sense for your team. Before you bake it into your model, align on it with your ICs and with leadership.

We can’t stress enough the importance of building a capacity model that derives the necessary insights *without* being overly-complex. There’s a considerable difference between asking “How many recruiters do I need to meet our hiring targets this year?” and “How many hires can I expect from each recruiter while *also* considering the level of difficulty and location of each role?” If that first question gets you all the insights you need, start *there*.

### 3.

## Determine historical productivity

---

Take the # of hires made per department over the last two years, and divide them by the # of dedicated, full-time recruiters who supported each of those departments. *This* is your PPR. Did your technical recruiting team also have a sourcer? A recruiting coordinator? It doesn't have to get much more complicated than this. You know that a quartet of 2 recruiters, 1 sourcer, and 1 RC = 4 technical hires per quarter. *That's* a ratio you can work with.

You can't predict future performance without knowing historical productivity and performance metrics. If you don't have historical data to work with (or if your data isn't as clean as it could be), use industry benchmarks, or ask your peers at similar-stage companies what numbers *they've* used to model. You can also make educated estimates. For example, if you're trying to figure out how many Sales hires you can expect from a single recruiter and a lot of your recruiters are new, pick the recruiter who's been at your org the longest and use *their* data as your proxy. Just know that the more your model is based on best-guesses rather than on your own data, the less accurate your model—and your forecast—will be.

We recommend collecting two years' worth of data to account for, and smooth out, seasonal trends. Of course, remove any anomalies—if there was a hiring freeze one quarter, extract that quarter's data from your calculations.

If you use a recruiting solution (like Gem!) that collects that data *for* you in a centralized source-of-truth, this step will be simple: pull up the number of hires made per department over the last two years, and divide those numbers by the number of full-time recruiters who supported each department. If you see, for example, that your average number of hires per quarter last year was 8.33 in Engineering and 9 in Sales, and you know that there were only two dedicated recruiters for each of those departments, you know you can expect around 4 hires per recruiter, per quarter in Eng and Sales. *This* is your PPR.

Did your technical recruiting team also have a sourcer? Did either team have a recruiting coordinator? It doesn't have to get much more complicated than this. You know that a quartet of 2 recruiters, 1 sourcer, and 1 RC equaled 4 technical hires per quarter. *That's* a ratio you can work with.

If you *don't* have a solution like Gem in place, this step of the process may be a bit more labor-intensive, since it will require you to reach out to your ICs (or have individual talent leaders reach out to their ICs) and find out how much time they spent supporting their various business units. Ask how many offers—extended your sourcers saw and how many offers—

---

Another way of arriving at PPR is to look at funnel efficiency. Maybe you find (or your recruiting analytics solution *tells you*) that 3,000 applications created → 300 pre-interviews → 24 phone screens → 6 onsite → 2 offers-extended → 1 offer-accept. How much *total* activity, then, occurred over the last two years to make the # of hires you made per department, per quarter? Divide that by the number of full-time recruiters who supported each of those departments to calculate PPR.

accepted your recruiters saw. If your recruiters work across business departments, ask what percentage of their time they spent on searches and hires for each unit. In your first go-round, these may just be estimates; but from here on out, ICs should be tracking this data for more accurate forecasts in the future.

Another way of arriving at PPR is to look at funnel efficiency: start at your Offer-Extend → Offer-Accept conversion rates and move backwards from there. Let's say your Sales roles see a 50% offer-accept rate, meaning you need 2 offer-extends for 1 offer-accept. And let's say your Onsite → Offer Extend conversion rate is 33%. That means you need 6 onsite to see 2 offers-extended and 1 offer-accept. *Now*, how many phone screens did it take to get those onsite? And so on—all the way back to outreach if your recruiting strategy involves outbound (and it should!).

All these data points will get you to the same place: how much activity occurred over the last two years to make the number of hires you made per department, per quarter? Divide that by the number of full-time recruiters who supported each of those departments to calculate PPR.

Note that if you're experimenting with the points system we mentioned above, you'll assign each headcount you hired over the last two years a number (based on difficulty) within the schema you created. From there, you can determine how many points each recruiter carried quarter-over-quarter.

Generally speaking, previous performance predicts *future* performance: if a recruiter hired, on average, 4 engineers per quarter last year, they'll likely be able to hit similar numbers *this* year. So once you've determined historical PPR, multiply that number by the number of recruiters on your team to forecast how many reqs your *existing* team can expect to fill per quarter in the coming year.



## Thoughts from Mike Moriarty

Talent Advisor

When it comes to capacity planning, the metric I live and breathe by is productivity-per-resource. With PPR, you're essentially looking at funnel efficiency: How many onsites do we need to get the number of offer-extends we need to see one offer-accept? And so on, up the funnel. Maybe you find that you need 300 pre-interviews and 24 phone screens to get an offer-accept. And you also know that you need 50 offer-accepts in department x this year. So there's your math.

Of course, there are variations within the funnel. I split headcount into two verticals: tech and non-tech, because you get different PPR outputs in each. At Dropbox we saw around double the production in non-tech that we saw in tech. *Then* you take other things into consideration. Active versus passive candidates, for example. Offer-accept rates are higher among active candidates, but passive talent is a critical element of any recruiting strategy. So how many of your hires are going to come from passive engagement? At Dropbox we aimed for 60%.

Consider channel. Our accept rate for referrals was typically around 80%, while our accept rate for sourced talent was below 60%. Consider level. At Dropbox, L2s were our lowest accept rates. Consider location. We found that San Francisco had a much higher decline rate than Austin did. And so on. These are patterns

you'll discover in your data over time. Of course, the model works best when everyone is working on the same roles and they've got a pipeline.

Ultimately you're looking to answer this question: How many offers-accepted did the average resource get last year? If the answer is five, and you need to hire 100 people, you need 20 recruiters. From *there*, you can determine the number sourcers based on your own sourcing goals, the number of interviewers, the number of recruiting coordinators. Work with an RC or look at their metrics: how many interviews can a person schedule in a day? Then do the math: how many RCs do we need—knowing that we have a 20% decline rate and so we have to schedule at 120%—to get the offer-accepts we need?

Monitor quarter-over-quarter. Some quarters our PPR went from 5.5 to 4.9. We'd miss our targets that quarter, dig into what happened. Like I said, we found that L2 software engineers in San Francisco were

Consider level. At Dropbox, sourced L2s were our lowest accept rates. Consider location. We found that San Francisco had a much higher decline rate than Austin did. And so on.

really hard for us to hire. But what's great is you have the data to say, Look, it's not that we're having trouble hiring software engineers *generally*. We're having trouble hiring *this* subset of *this* subset. We put in new processes to optimize for that subset. We figure out how to decrease time-to-fill there, which is a metric with a direct relationship to PPR.

If you're a recruiting leader, you know that you're always tinkering with *something*. I might say to a functional leader, Listen, we've historically had an 80% close rate. It's currently 50%. *That's* our problem. And here's why; and here's what we're doing to combat that. But as soon as your accept rates are on pace, you look at the data and see that no one's referring anyone despite the fact that you banked on a third of your hires being referred. Which is why it's so important to have a daily relationship with your data. There's also the fact that, in the end, Finance is what you're solving for. If Finance sees the recruiting org as a cost, I've failed to give them a data-backed reason for why we're an investment instead.

## 4. Use historical productivity to forecast future performance

---

Once you've determined historical PPR, multiply that # by the # of recruiters on your team to forecast how many reqs your *existing* team can expect to fill, per quarter, in the coming year. Factor in recruiting coordinators with the recruiter-to-RC ratio that's historically worked for your org. Factor in sourcers based on things like volume and quality of your inbound, and current pipeline for imminently-opening roles.

Of course, you'll have recruiters who are more senior and others who are more junior. But if you're building a balanced team, you're working on the assumption that everything will average out. If your PPR is 10, one recruiter might make 8 hires while another makes 12. But you're *still* averaging 10 hires per recruiter, as your model predicted.

When it comes to other roles, you'll factor these in as you need. Perhaps you already have a recruiter-to-recruiting-coordinator ratio in place; keep that ratio for as long as nobody feels overextended. When it comes to sourcing support, *you'll* decide how much you want sourcing to contribute to your total number of offers extended. In part this may depend on the volume (and quality) of inbound you receive; in part it may depend on your best source-of-hire: how well have sourced candidates performed in your pipeline historically? Do you need to build pipeline in Q1 to prepare for a hiring spree in Q3? This might be a reason to invest in more sourcers—but perhaps temporarily, through an RPO, rather than as headcount. And so on.

You know your hiring funnel better than we do; so *you* know whether your top, middle, or bottom-of-funnel needs more support. This will determine the sourcer-to-recruiter ratio that's supporting a given department in a given quarter or year.

## 5. Factor in attrition and a buffer

---

Factoring attrition into your capacity model means recognizing that talent will leave your org throughout the year *even as* you're hiring to hit your gross targets. Factoring in a buffer prepares you for unforeseen circumstances (sick leave, parental leave, termination), and ensures you're not in a situation in which everything has to go perfectly all year for the team to meet its hiring goals.

There's at least one adjustment—and perhaps there are two—that you'll want to make to your model once you've determined your PPR. The first is attrition, which is a number you'll get from HR or from your FP&A team. If your goal is to end the year with 80 engineers and your current headcount is 50 (so 60% growth), that means your gross target is 30 engineers. But if your company's average attrition rate is 15% and you don't model that in, you'll be short the 7-8 engineers who left your org in the meantime. So bake 38 hires into your plan in order to net out on your end-of-year target of 80.

The other thing to consider adding to your capacity model is a buffer for unforeseen circumstances. Sick leave, parental leave, termination—all of these things cause you to lose resources for periods of time. Building in a buffer ensures you're not in a situation in which everything has to go perfectly all year for the team to hit its hiring goals.

Granted, if you're modeling using your organization's historical data, there's already a buffer of sorts in there, since your team probably *did* experience unforeseen circumstances at some point over the last two years. But if that's not the case, consider building in a 10% or 15% buffer—whatever the business is comfortable with. (If you have difficulty getting Finance to agree to a buffer, ask them this question: If we *don't* hit our targets, what is the risk to the business?) A buffer means that, for some of the year, your recruiters will be operating under capacity. And that's okay. Because you want to have room to expand and be agile if you need to.

## Thoughts from Bernardo Sosa

Director, Talent Sourcing, Ops & Programs



At Box, our fiscal year starts in February; we start capacity planning in September. That's plenty of time to think deeply about what changes we want to make to the model, or if we want increased complexity from last year's inputs. It also gives us time to work with Finance to have conversations with functional leaders. At our size it's messy to talk to functional leaders individually about headcount, because we'd get wishlists that haven't been approved. So we partner directly with our corporate Finance team and ask them to get granular: not only what we're hiring, but also when and where. Q1, Q3, remote, international? All of these are important inputs for my model. Of course, headcount is a moving target even after you start the fiscal year. Your model needs to allow for that flexibility. But Finance-approved headcount is the strongest thing you'll get at any stage.

I tend to see hires in three buckets. The biggest bucket is ICs or entry-level managers. Then you have executive hires—leadership and above, senior directors. I take these out of my model because they're white-glove; they take a lot of time to hire. We have a separate exec recruiting team that takes care of those roles. The third bucket is university hiring, and I remove those high-volume reqs from my model

as well. It wouldn't be fair to keep them in because our recruiting program manager has such high productivity numbers. What I solve for is that bucket in the middle. I care about those other two buckets. But we have to think about resourcing differently with them.

Then it's a question of historical productivity. Granted, some roles are trickier to fill than others because they're new or niche. But I don't distinguish there, because the less complexity there is in the model, the easier it is to tell the story. In engineering, a full-cycle recruiter can typically close five to six hires per quarter. For G&A, that number is twice as high. Go-to-market falls somewhere in the middle—between six and ten. But those numbers may change depending on any number of circumstances, which is why I always talk with my leadership team to see if they're comfortable using the productivity assumptions I've made based on historical data. Maybe eng recruiting isn't resourced to deliver five hires per recruiter next quarter. They're building pipeline in a new place; they're hiring international in a new market; the team is ramping. Whatever the case may be.

I tend to see hires in three buckets. The biggest bucket is ICs or entry-level managers. Then there's executive hires and university hiring, which are white-glove and high-volume, respectively. I take these out of my model, and solve primarily for that IC bucket.

I should say that what's hard-coded into our model is the number of full-cycle recruiters. Sourcers are more flexible because we decide how much we want sourcing to contribute to the total number of offers we extend. Of course we base that on historical; but it's my decision whether I want sourcing to represent 20% or 50% of offer-accepts. That determines how many sourcers I need. Always ask if it's a top, middle, or bottom-of-funnel problem you're trying to solve in a given pipeline. Sometimes it's just, Hey, we need to build pipelines ahead of a hiring spree in six months. So let's invest in sourcers now. And at that point, we'll engage a sourcer RPO.



*Thoughts from Bernardo Sosa continued*

Which brings us to resource-expansion. At some point in the planning cycle we start to get a clear idea of how much we need to expand the team by. Have RPO vendors or contingent workers that you can plug in where you need. If I know we'll need extra support in the next six months, we'll engage a contractor. The model has to include that, because it's the easiest thing to pull back or scale.

My biggest piece of advice is to start with a very basic model. Don't try to solve everything the first time around, because it'll get too complex too quickly. Nobody will understand it. And you won't be able to craft an easy story that will bring Finance along with you. Sophisticate it over time. Because there will come a point when you'll need to know whether you need a sixth recruiter or a third sourcer. And you'll need to explain the rationale for your decision.

Have RPO vendors or contingent workers that you can plug in where you need. If I know we'll need extra support in the next six months, we'll engage a contractor. The model has to include that, because it's the easiest thing to pull back or scale.

## 6. Check in with your leaders: are they comfortable with your numbers?

---

If you have a talent leadership team, you'll want confirmation from your leaders that their respective teams are capable of doing what they've historically done. New recruiters might be ramping, pipeline might need to get built in new markets. Get feedback from individual leaders about where their teams are *now*, and make adjustments to your model where necessary.

Of course, if your org is small enough that you *are* the leader, consider this step done. But if you have a talent leadership team, you'll want confirmation from your leaders that *their* respective teams are capable of doing what they've historically done. Do your leaders feel comfortable using these productivity assumptions? Is their team currently equipped to deliver x hires per recruiter?

Your eng recruiting leader might tell you that her new team in Poland means recruiters are still ramping and faced with a new market where your company has less brand recognition. She can't commit to the number you brought to her, but if you knock it down by four, she can commit her team to *that*. That's critical feedback. Maybe your sales recruiting leader reminds you that his team has been burning the candle at both ends over the last four quarters, and they can't continue at this capacity—otherwise you're burning them out for the foreseeable future. Maybe your biz recruiting leader reminds you that she's hiring in Boise for the first time, and the org doesn't have any pipeline in Boise. So she can only promise 90 hires next quarter rather than 100... but she can commit to 100 the *following* quarter, because the team will have built pipeline by then. And so on.

On the other hand, perhaps a leader tells you that their team has increased efficiency by 10% over the last two quarters. They onboarded a recruiting automation solution, or recruiters are now fully ramped. Get that qualitative feedback from your leaders. Make adjustments to your capacity model where they propose them.

## 7. Calibrate the gaps

Two great strategies for when your team is stretched are to increase your hiring capacity (through short-term contingent workers, for example), and to optimize your funnel for speed. Lowering hiring goals, drawing out hiring timeframes, and tech-stack additions that increase your PPR are other options. Collaborate with leadership to decide what's best for the business.

By the end of step 6, you have your capacity plan. With the data now in one place, you can see what the gaps are between your team's current capacity and the business' hiring goals. There are a number of ways to close those gaps:

- **Increase your hiring capacity**, whether that means growing the team (if you'll need that additional headcount for the foreseeable future), augmenting with agencies or RPOs, or plugging in contingent workers for the short-term. The latter are the easiest to scale and pare, though third-party firms are expensive. Still, it's always good to have a shortlist of RPO vendors to replace headcount for the short-term—whether that's covering maternity leaves or sudden spikes in hiring.
- **Lower hiring goals**—for instance, cutting any nice-to-haves from the headcount plan.
- **Draw out hiring timelines**. Does that volume of hires *really* need to be made within the timeframe you've been given? If not, modify the timeframes for certain hires.
- **Consider tooling**. Decreasing time-to-hire (and thereby increasing your PPR) is another way of calibrating the gaps between your team's capacity and its hiring goals. Outreach automation, scheduling software, and assessment software are all tech-stack additions that will increase efficiency, shorten time-to-hire, and give your team time back to make *more* hires.
- **Optimize your funnel for speed**, which is another way of increasing PPR. Maybe you skip a stage—that second phone screen or that additional interview—to move candidates through the funnel more quickly. Maybe you incorporate batch days—all-day events in which multiple candidates interview for the same role/s—since a few focused days of interviewing is more efficient than months of drawn-out process. All of these are ways of moving faster.

Other gap-closing strategies include internal hiring, leveraging employees as recruiters (i.e. through an employer branding or referral program), and more. Get creative. Collaborate with leadership to determine the best strategy for your business. Of course, at any time you find you're operating *under* capacity, consider using those extra resources for process improvements, L&D, pipeline building, attending to branding or diversity initiatives, and more.

## Thoughts from Joel Torres

Talent Acquisition Manager



Twilio has seen such high growth that capacity planning has been especially tricky, since headcount keeps increasing mid-year. Most recently we wanted to take advantage of the talent market during COVID. Now growth is complicated by expected attrition from “The Great Resignation.” We need to consider worst-case scenario, in which net growth might be a lot higher than gross growth. I wouldn’t say we have the equation down. But we’re certainly considering all of its elements.

Some things have shifted as the company has grown. One is that we started dividing things out to plan more accurately. We’ve got three different divisions: G&A, R&D, and Go-to-Market. We further segmented those divisions: for example, splitting out very niche roles within R&D. Those go to a specific team, as do exec and leadership roles. By taking out executive, taking out niche, recruiters can focus on, say, L2 to L5 engineering talent. They can deliver at a volume and know how to pace. Otherwise they’ll get slowed down by those harder-to-hire roles.

Another thing is that historically, Twilio had a lot of niche roles across eng and product. As you scale,

The questions I’m asking ahead of every planning cycle include: *When in the year do we want those hires? Some things you know historically. R&D likes to front-load, for example. If you resource with contractors that’s useful knowledge, because you may not have enough work for those contractors later in the year.*

specialization wanes a bit, and talent acquisition adjusts to support a different model. This may impact your sourcing function, for example. Some orgs hire sourcers who are specialists, rather than pipeline people who work on evergreen roles. Others think in terms of channels and are constantly turning on and off the proverbial faucet for those different channels. They don’t need the specialists so much; they need volume folks. Twilio has had a hybrid approach; our sourcers could do both volume and niche. But the less niche our roles get, the more strategy and predictability we can have. It’s no longer: everyone pull out every skill that you have to tackle these reqs.

The questions I’m asking ahead of every planning cycle are: Where will we be growing? *How* will we grow? What products are we focusing on? What kind of talent are we looking for? And *when* in the year do we want those hires? Some things you know historically. R&D likes to front-load, for example: Q1, Q2. If you resource with contractors that’s useful knowledge, because you may not have enough work for those contractors later in the year. Once we have the answers to these questions—which we get from HR, Finance, Directors and VPs who are planning their orgs for the next year—we look back two years to see what our capacity to hire has been. How many hires can one person make per quarter? Within R&D I can pretty safely say that, *without* talent sourcer support, a recruiter can make five hires per quarter. If we have 1,000 hires to make at 250 hires per quarter, that’s 250 divided by five.

I personally like to staff my team to 120% capacity, because 80% of the team will hit their goals and 20% won’t. Staffing to 120% means you’re still hitting 100% of your goals team-wide. Of course, the surprises will still come, especially when staffing a global organization. A functional leader will say, Great, I got

*Thoughts from Joel Torres continued*

this headcount for Austin. Except I need more than one engineer. So I'm gonna move that headcount from Austin to Columbia, because I can hire *three* people in Columbia rather than one. So how are *you*, as a TA leader, going to structure your org both to hit those goals—because sometimes Finance will say, we just don't have the money for the resources you need—and to be agile?

I often imagine Roman Centurions in classical battle. Ideally you'd have one-to-one "combat": one recruiter per role. But you don't have the budget, so you have to think about the tactics available to make up for the fact that it *isn't* a one-to-one situation. Maybe you need an archery team. Maybe you need people who can shoot cannons and "take out" 20 roles at once. Maybe that's a hiring event—in which case, who loads the ball, who lights the fuse, and who aims it; and should they be three different people? It's a helpful metaphor for me to ask: What's the most effective and efficient way for us to win this recruiting battle?

A functional leader will say, Great, I got this headcount for Austin. Except I'm gonna move that headcount from Austin to Columbia, because I can hire *three* people in Columbia rather than one. So how are you, as a TA leader, going to structure your org both to hit your goals and to be agile?



## 8. Validate and socialize your model

---

If Finance sees the numbers in your model as a cost versus an investment, you haven't fully explained what the business will ultimately *get* out of supporting additional recruiting resources. Speak to Finance in *their* language. It's one thing to say that a program manager will standardize your interview process; it's another to say that standardization will decrease your time-to-hire by four days on average—which means x additional dollars in revenue because your Sales hires will start four days sooner.


Now that you've built your capacity model and determined the best solution/s for calibrating your gaps, your final step is getting buy-in from the business. It's not uncommon for a leader outside of TA to experience sticker shock when you share how many resources you'll need to meet hiring needs. By now, you've had HR validate your inputs. Your department leaders understand that, because they have x% of the headcount plan, they'll have x% of your team's resources month-over-month.

Your final hurdle is Finance (whom ideally you've been in conversation with from the beginning). Luckily, your capacity model explains your forecasting and resourcing logic objectively, which should help them get on board. But expect that you may have to iterate on the model with Finance until you all align.


If Finance sees the numbers in your model as a cost versus an investment, then you haven't fully explained what the business will ultimately *get* out of supporting additional recruiting resources. However you can tell the story of TA as a profit-generating resource rather than a cost center, do *that*. The logic of scaling recruiters is generally straightforward: more recruiters = more hires. But be especially careful to speak to Finance in *their* language when it's time to go to bat for the program managers and ops roles you've put into your model. It's one thing to say that a program manager will standardize your interview process; it's another to say that standardization will decrease your time-to-hire by four days on average—which means x additional dollars in revenue because your Sales hires will start four days sooner.

We could add a step 9 here, which would be the iterative step of adding accuracy and complexity to your model over time. But fiscal years and capacity models are cyclical, so that iteration is ultimately in steps 2 and 3, above. As you gather data year-over-year, you'll begin to see how you're trending over time. You'll begin to better understand how, and when, to extend your model to include sourcers, coordinators, contractors, program managers, ops roles, and so on.

Rinse and repeat regularly: while you may only iterate *wholly* on your model once a year, you should perform a monthly reforecast to ensure your initial (or most recent) inputs were accurate. Adjust where you need to. And be patient with yourself. Capacity planning isn't an exact science. But it will certainly put you in a position to be a *much* stronger partner to your business.



# How gem can help



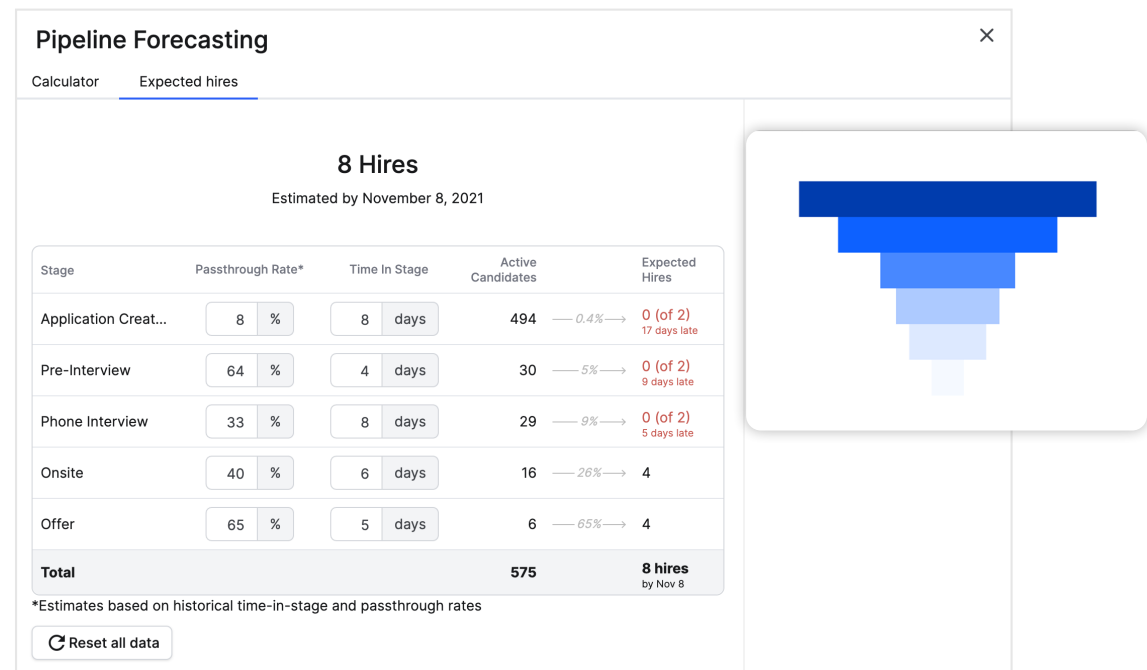
# Introducing Gem's Talent Compass

## Forecasting & planning

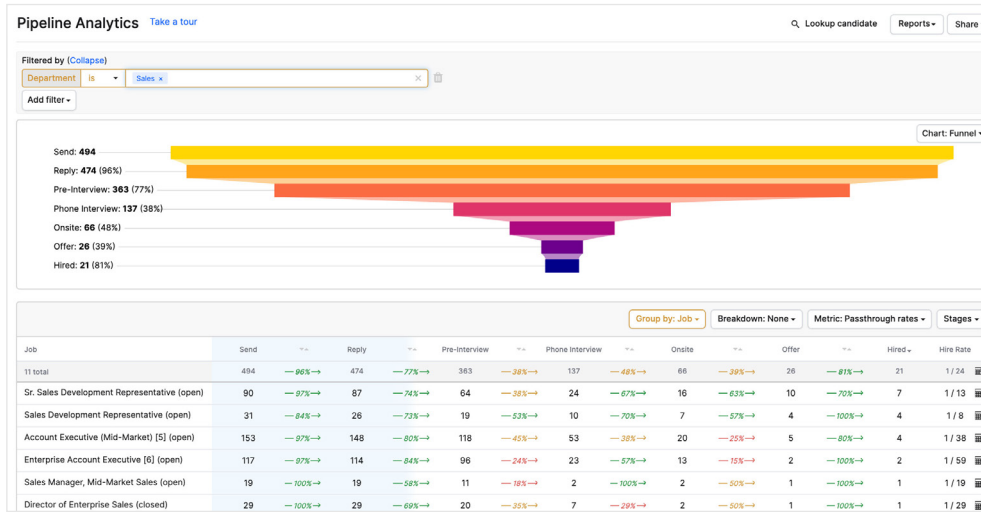
Talent Compass includes a forecasting calculator that projects not only how many hires you can expect to make, but also by when. This can be forecast at the individual job req or at an aggregate level, like department.

Passthrough rates and time-in-stage are based on historical data, but can be adjusted. This is particularly useful when teams want to understand the impact of potential process improvements. The calculator also works backwards: users can input the final number of hires they would like to make and the calculator will output the number of candidates and applicants required at each stage in order to meet final goals. This feature is particularly useful for teams looking to capacity plan in the face of ever-growing hiring goals.

For talent acquisition teams looking for insights that offer strategic guidance—including when it comes to capacity planning—Gem offers Talent Compass. Talent Compass is equipped with full-funnel visibility, hiring forecasts, performance metrics, and executive reporting that TA teams use to plan ahead and guide their recruiting strategy.







## End-to-end visibility

Tap into every dimension of the recruiting funnel. With access to all recruiting data from Gem and the ATS from one source of truth (without raising a support ticket), TA teams can track everything outreach to hire and identify key trends across the process. They can then slice and dice metrics the way their business thinks: by role, department, geo, demographics, and more. For example, recruiting teams can uncover outreach strategies that lead to more replies for specific roles and industries. They can also track how these candidates passthrough the funnel and spot any troubling dropoffs or bottlenecks.

## Performance optimization

Within the Pipeline Analytics module, recruiting managers and leaders can monitor individuals and roles to better understand how their teams are faring. Reviewing specific job breakdowns allows managers to observe pipeline risks across roles or even spot any troubling candidate bottlenecks and/or drop-offs. Within dashboards, managers can even create widgets with specific metrics for their individual recruiters (e.g., phone screen to interview ratio, offer acceptance rate, etc.), serving as a performance scorecard.

Job	Send	Reply	Pre-Interview	Phone Interview
Total	495	474	363	137
▼ Sr. Sales Development Representative (open)	90	87	64	24
Gisela Faretta	66	64	47	19
Bella Gutkin	7	6	4	2
Georgena Frazier	1	1	1	1
Aileen Tseng	1	1	1	-
Bruce Lang	14	14	10	2
(no recruiter set)	1	1	1	-
Market [5] (open)	31	26	19	10
Bella Gutkin	30	25	18	9
Deb Feldman	1	1	1	1
Aileen Tseng	1	1	1	-
John Evans	1	1	-	-

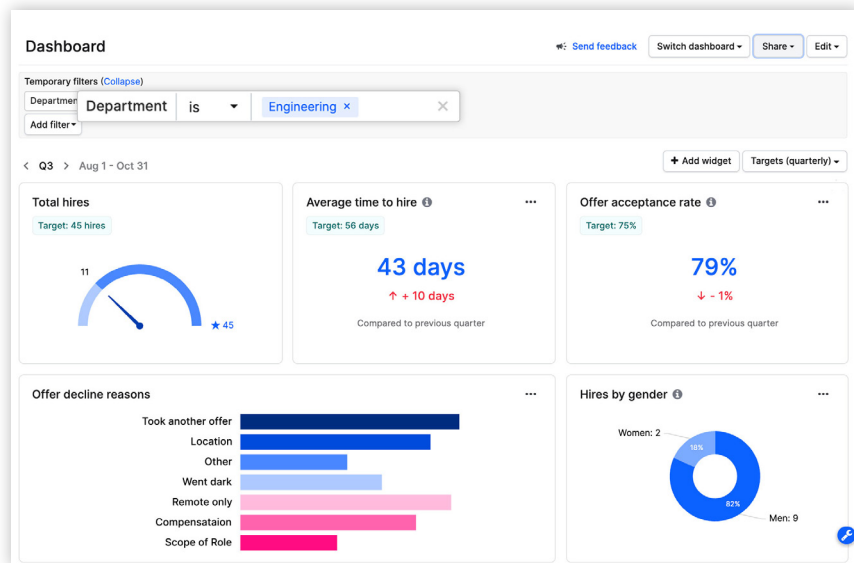
“[Our] talent acquisition [team] has even more credibility now because they are no longer speaking anecdotally; they’re speaking with data.”

Candice Tang  
Director of Talent Acquisition



## Executive reporting

TA teams often need to report and collaborate with their functional counterparts. Within Talent Compass, users can launch configurable, presentation-ready reports to visualize KPIs. With easily digestible metrics on hand, talent acquisition teams can come prepared to discuss how they are progressing and where there are opportunities to improve. Reports can be customized with built-in widgets and filters, thereby surface tailored, actionable insights.



## Diversity recruiting insights

Talent Compass also offers visibility and insight into how an organization recruits diverse talent, from initial outreach to eventual hire. At the top of the funnel, teams can review the progress they are making in reaching out to diverse talent. By pinpointing specific content and strategies that drive greater engagement, teams adjust their sourcing strategy towards building diverse applicant pools. With visibility into how candidates from underrepresented groups pass-through the recruiting funnel, hiring teams can uncover any disproportionate drop-offs and adjust their process accordingly.

Group by: Project | Breakdown: Race/Ethnicity x Gender | Metric: Passthrough rates

Project	Send	Reply	Pre-Interview
202 total	5,654 (-20%)	1,137 (-18%)	208 (-56%)
Black Women	358 (6%) (-18%)	66 (6%) (-15%)	10 (5%) (-30%)
Black Men	288 (5%) (-28%)	80 (7%) (-25%)	20 (10%) (-50%)
Black, Unknown gender	6 (0%) (-33%)	2 (0%) (-0%)	-
Hispanic/Latino Women	285 (5%) (-18%)	50 (4%) (-14%)	7 (3%) (-57%)
Hispanic/Latino Men	266 (5%) (-19%)	51 (4%) (-12%)	6 (3%) (-50%)
Hispanic/Latino, Unknown gender	2 (0%) (-0%)	-	-
White Women	1,025 (18%) (-16%)	162 (14%) (-27%)	43 (21%) (-70%)
White Men	937 (17%) (-22%)	207 (18%) (-19%)	40 (19%) (-73%)
White, Unknown gender	37 (1%) (-14%)	5 (0%) (-20%)	1 (0%) (-100%)
Other	1,282 (23%) (-17%)	218 (19%) (-14%)	30 (14%) (-40%)
Other, Unknown gender	1,145 (20%) (-25%)	289 (25%) (-18%)	51 (25%) (-49%)
Other, Unknown gender	4 (0%) (-25%)	1 (0%) (-0%)	-

Gem's end-to-end modern recruiting solution empowers talent acquisition teams to engage their entire talent network, optimize sourcing efforts, and uncover actionable insights that guide smarter, forward-looking decisions. Gem works alongside LinkedIn and other places that you source, while integrating with Gmail, Outlook, and your ATS. Find the talent you need to meet hiring targets and scale your teams with Gem.

To learn more and see a demo, visit [gem.com](https://gem.com)



---

**Lauren Shufan, Author**

Lauren is a content strategist with a penchant for 16th-century literature. When they're not trying to tap into talent teams' pain points, they're on their yoga mat or hiking the hills of Marin County. Come at them with your favorite Shakespeare quote.