WHITE PAPER · TALENT LEADERSHIP

Capacity planning for talent acquisition teams

How to work out how many recruiters you need: a two-method capacity model, a worked example, what 2026 benchmarks measure and why they disagree.

R **Ryzon**

Published 27 November 2026 Updated 27 November 2026 7 min read

IN SHORT

Work out demand from your own hiring plan, segmented by role, and compare it with the hours your recruiters really have. Check the answer a second way, using how many roles will be open at once. Use published benchmarks to question your assumptions, not to replace them, because providers define the same measures differently.

- Model demand by role family, including backfill and roles that restart or close unfilled.
- Count the hours recruiters really have, after leave, other work and ramp time.
- Check the result with a second method before you commit to headcount.
- Benchmarks with the same name often measure different things. Match definitions before comparing.
- Agree in advance what triggers extra help when the plan changes.

A hiring plan is a promise made on behalf of a talent team. Capacity planning is how you find out whether the team can keep it. This paper sets out a practical model for heads of talent and people leaders at technology scale-ups, a worked example and a guide to reading the benchmarks.

How this paper was put together

We reviewed benchmark reports and capacity-planning methods published by recruiting software providers, professional bodies and official statistics, using the latest versions we could access on 27 November 2026. Some full reports are gated and were not read. This is a synthesis of other people's research, not new data from Ryzon. Most vendor benchmarks describe that vendor's own customers.

Why a single ratio will not answer the question

It is tempting to ask how many requisitions a recruiter "should" carry. Consider two figures that are both often reported as 14:

- Gem reports that the average recruiter in its 2024 data handled 14 requisitions at the same time, with averages for company-size bands ranging from 10 to 14.[1]
- Ashby reports that, at the 90th percentile of recruiter weeks, recruiters and sourcers worked on 14 jobs in a week.[2]

The first is an average of concurrent requisitions. The second describes a **busy** week and counts jobs with activity. Neither is a recommended load.

Role mix changes the picture too. In Gem's 2024 data, engineering roles took an average of 53 days to hire and 39 interviews per hire. Sales took 35 days and 21 interviews. Customer service took 26 days and 9.[1] A recruiter with ten engineering roles does not have the same job as one with ten customer service roles.

Start with demand

Build the hiring list

For the planning period, list expected hires by role family, level and location. Add:

- **growth roles** from the business plan, marked as certain or conditional;
- **backfills**, based on your own turnover, not a general average;
- **restarts and unfilled openings**, from your own history.

That last line matters. Greenhouse's UK data for January to August 2026 reports that 33% of openings closed without a hire.[3] It does not define the measure on the page and covers its

own customers, so use your own figure. The point is that work happens on roles that never produce a hire.

Estimate effort per hire

For each role family, estimate the recruiter hours needed from brief to signed offer: intake, sourcing, screening, coordination that falls to the recruiter, interviews and debriefs, offer and reporting. A few weeks of light time tracking is better than a guess.

Use stage volumes where you have them

If you know how many candidates move between stages, you can build the hours from the bottom up: forecast the volume at each stage, give each stage a time cost and add them up.[4] This also shows where better process would free the most time.

Then count supply

For each recruiter, start with the weeks in the period and remove holiday, public holidays, sickness and training. Multiply by weekly hours. Then take the share that goes on hiring work. Projects, employer brand, reporting, interviewer training and systems work all take time. Subtract that time **once**: if your hours-per-hire figure already includes some of it, do not remove it again.

Ramp time

A new recruiter is not at full capacity on day one. Gem's guide gives an example in which a ramping recruiter might be expected to hire at half capacity,[5] and another Gem account of one company's model assumed three months to ramp.[6] These are assumptions, not benchmarks. Set your own.

A worked example

Every number in this example is an assumption chosen to show the method. Replace each one with your own.

A company plans 24 hires over two quarters.

Role family	Planned hires	Assumed recruiter hours per hire	Hours
Engineering	10	50	500
Go-to-market	10	35	350
Operations	4	28	112
Total	**24**		**962**

Add 15% for restarts and openings that will not result in a hire: about **1,106 hours** of demand.

Supply per fully ramped recruiter over 26 weeks: remove 5 weeks of leave, holidays and sickness, leaving 21 weeks. At 37.5 hours a week, with 65% on hiring work, that is about **512 hours**.

1,106 ÷ 512 = about **2.2 fully ramped recruiters**.

Now suppose the team is one experienced recruiter and one joining on the first day, who works at half capacity for three months and fully after that. Over the six months the new joiner supplies three quarters of a ramped recruiter. The team has 1.75 × 512, about **896 hours**: roughly **210 hours short**.

Test the assumption that matters most

If real effort per hire is 20% higher, demand rises to about 1,330 hours, or about **2.6 ramped recruiters**. Run this test on each important assumption and see which one changes the decision.

Check it a second way

A useful cross-check is the number of roles open at the same time:

average open roles = openings in the period × average days open ÷ days in the period

This is a standard queueing relationship, and a recruiting software company, Dover, uses the same form to estimate concurrent searches.[7] Keep the units consistent: use calendar days throughout.

In the example, 24 hires with an assumed 60 days open, over 182 days, gives about **7.9** roles open on average, or about **9.1** with the 15% allowance. Split across 2.2 recruiters, that is

about four each, on average. Peaks will be higher, because hiring rarely arrives evenly. If the two methods tell very different stories, find out why before you decide.

A third lens: your own throughput

You can also ask how many hires each recruiter has actually made per period and compare that with the plan. Gem's capacity guide describes this approach, using past hires per recruiter and adjusting for ramp.[5] It is quick, but past output reflects past demand and past conditions. It can understate what a well-supported team can do, or overstate it if last year's hiring was easier.

How to use benchmarks

Published numbers are useful for asking whether your assumptions are reasonable. They are poor substitutes for them.

Hires per recruiter

Ashby reports 7.3 hires per recruiter per quarter in early 2026.[8] Its separate analysis of Europe, the Middle East and Africa shows a range of roughly 3.8 to 4.8 over more than three years.[9] Gem's pages report that hires per recruiter have fallen since 2021.[10] The datasets, periods and customer bases differ, and the sources do not reconcile them. Do not average them.

Interviews per hire

Source	Figure	What it counts
Ashby, early 2026	11.7 for business roles and 17.6 for technical roles[8]	The chart labels describe applicants interviewed per hire
Gem, 2024	20 overall[1]	All interviews used to fill a role, including with unsuccessful candidates

These are not two estimates of the same thing. A model that treats applicants interviewed as interview sessions will badly understate the work.

Applications per hire

Greenhouse reports 182.6 applications per hire in the UK for January to August 2026.[3] Ashby reports 291 in its wider dataset.[8] In the UK, CIPD's 2026 survey found that 58% of employers saw more unsuitable applicants and 33% saw more suitable ones.[11] More applications do not necessarily mean more qualified candidates, so do not assume screening effort scales with value.

Offer acceptance

Ashby reports about 80% acceptance at start-ups.[12] Gem's 2024 data shows 76% for engineering and 90% for customer service and support.[1] Use your own rate, split by role family, because it directly changes how many processes you need to run.

Plan for spikes

Hiring plans change with funding, launches and lost customers. Decide in advance what you will do.

Scenario	What changes	Decide now
Base	The plan as agreed	Team and priorities
Higher	More roles or earlier start dates	Trigger for contract, fractional or embedded help, and who approves it
Lower	Roles paused or removed	What the team works on instead

Because a new recruiter takes time to hire and to ramp, the trigger needs to fire well before the roles are open. A practical trigger is the number of approved roles per ramped recruiter, taken from your model, held for more than a few weeks.

Protect quality

A capacity model that only counts speed will reward the wrong behaviour. CIPD's 2026 report found that three in ten employers measure new-hire performance or retention in the first year.[11] Review quality with hiring managers alongside volume, and leave time in the model for it. See recruitment metrics that matter.

Run it as a routine

1. **Monthly:** compare planned and actual hires, roles open and hours per hire.
2. **Quarterly:** rebuild assumptions and rerun the scenarios.
3. **At any plan change:** show leadership what moves if a role is added, in dates rather than adjectives.

That last habit changes the conversation. “We are stretched” is easy to dismiss. “Adding these six roles moves the last four start dates into the next quarter unless we add help by this date” is a decision leaders can make.

Where Ryzon fits

Ryzon places senior talent partners inside scale-up teams, fractionally or full-time, which can cover a gap while a permanent team is built. For the wider plan behind the numbers, read strategic workforce planning for scale-ups.

Sources

1. Gem, 2025 recruiting benchmarks, 2025 (data 2021 to 2024).
2. Ashby, What is a recruiter’s capacity?, November 2025.
3. Greenhouse, UK recruiting benchmarks, 2026.
4. ERE, How a talent acquisition team can look at its own hiring plan, January 2017.
5. Gem, The recruiting leader’s guide to capacity planning, accessed 27 November 2026.
6. Gem, Gathering data for a recruiting capacity model, accessed 27 November 2026.
7. Dover, Recruiting capacity planning for fundraising rounds, updated September 2026.
8. Ashby, Recruiter productivity, 2026 talent trends, April 2026.
9. Ashby, EMEA recruiter productivity, September 2026.
10. Gem, Tech hiring guide, accessed 27 November 2026.
11. CIPD, Resourcing and talent planning report 2026, September 2026.
12. Ashby, The state of startup hiring, February 2026.

Common questions

How many open roles should one recruiter handle?

There is no reliable universal number. Published figures measure different things: an average number of requisitions worked at once, the number of jobs touched in a busy week, or requisitions over a whole year. Work it out from your own hours per hire and role mix.

How often should we update the model?

Review it monthly against actual hires and open roles, and rebuild the assumptions at least each quarter or when the hiring plan changes.

Should coordinators and sourcers be in the same model?

Give them their own lines. Their work and their peaks differ from a recruiter’s, and adding them together can hide a shortage.
