

TECH RECRUITMENT 101

10 Metrics Every Recruitment Team Should Track



PART OF THE *TECH RECRUITMENT 101* SERIES

10 Metrics Every Recruitment Team Should Track

*A practical dashboard for seeing what is really happening in
your hiring process*

CONTENTS

1 Funnel metrics

2 Quality metrics

3 Outcome metrics

4 1. Stage conversion

5 2. Time in stage

6 3. Response rates

7 4. Hiring manager feedback speed

8 5. Source quality

9 6. Interview pass-through rate

10 7. Onsite-to-offer ratio

11 8. Calibration consistency

12 9. Offer acceptance rate

13 10. Time to fill

14 Keep the cadence boring

15	A working template	
16	The action question is the part most teams skip	
17	Quality of hire	
18	Pipeline aging	
19	Requisition load	
20	Cost per hire	
21	Candidate experience score	
22	When to add these	

Introduction

Your team may have a dashboard full of metrics in your ATS, and still no good answer to a simple question: when are we finally closing that ML Engineer role? The numbers are there. The reports exist. But nothing on the screen tells you what is stuck, why, or what to do about it.

This book gives you ten metrics that actually help you act. I picked these ten because they are the ones I have seen change how recruiting teams operate. Each one shows a specific problem, points to a cause, and tells you what to look at next.

In Chapter 1, I explain what makes a metric useful and why most dashboards fail that test. Chapter 2 gives you all ten metrics at a glance. In Chapters 3 through 5, we go deep: four funnel metrics that show where hiring slows down, four quality metrics that keep your pipeline honest, and two outcome metrics that show whether the whole system worked. In Chapter 6, I help you assemble these into a dashboard your team will actually use. Chapter 7 covers other metrics worth knowing once the foundation is solid.

This mini-book is part of the Tech Recruitment 101 series. I wrote it for recruiters getting into tech hiring who want a clear operating system for the work. Metrics will not make hiring easy.

They can make it much easier to see what is actually happening.

What metrics are for, and what they are not

Metrics should help you make a decision. If a number does not tell you what to fix, where to look, or which question to ask next, it is decoration.

Recruitment teams miss this all the time. They track emails sent, calls made, applicants received, resumes screened, and interviews booked. The dashboard looks busy, but hiring still stalls.

The problem is simple: activity is not the same as movement. A recruiter can send 500 outreach messages and create no real pipeline. A role can attract 300 applicants and still produce one weak shortlist. Volume matters only if it turns into progress.

I treat metrics like a flashlight, not a trophy. A good metric shows friction in the process — it tells you where candidates disappear, where hiring managers delay, or where a source looks strong at the top and fails later. A weak metric just tells you something happened, while a useful one tells you what to do about it.

This matters even more now because AI makes activity cheap. Tools can generate more outreach, more scheduling, and more admin output than any team could produce manually, but that

means old activity metrics stop meaning much. If outreach doubles and qualified screens stay flat, you did not improve the funnel — you automated noise. So pair volume with conversion, and pair speed with outcomes. Remember this: faster and worse is just a more efficient mistake.

You also need to separate leading indicators from lagging indicators. Leading indicators move earlier in the process: response rate, stage conversion, time in stage. They help you act now. Lagging indicators arrive later: offer acceptance, quality of hire, new-hire retention. They tell you what the whole system produced.

That distinction matters because recruiters often get judged on outcomes they do not fully control. A declined offer can mean pay was off, the process dragged, the hiring manager did not sell the role, or the candidate never trusted the loop. By the time that outcome lands, the damage happened earlier. Leading indicators let you spot the problem while you can still do something about it.

Be skeptical of any metric that flatters the team but explains nothing. Keep a metric if it helps you act. Drop it if it only helps you perform competence in a meeting.

The ten metrics at a glance

Before we go deep, here is the full list. These ten metrics split into three groups: funnel metrics that show where hiring slows down, quality metrics that keep your pipeline honest, and outcome metrics that show whether the whole system worked.

Funnel metrics

1. Stage conversion is the share of candidates who move from one stage to the next. It shows where the pipe narrows. If many candidates pass recruiter screen but few survive the technical interview, that stage needs a closer look.

2. Time in stage is the median number of days candidates spend in each step. It shows where work sits still. A seven-day wait between final interview and decision is specific enough to fix. Overall time to hire is not.

3. Response rates cover both sides of the conversation. Candidate response rate tells you whether your outreach earns a reply. Recruiter response rate tells you whether your team follows through once a candidate engages. Both matter, and they diagnose different problems.

4. Hiring manager feedback speed is the time from interview completion to written feedback or a decision. When this is slow, everything downstream stalls — recruiters cannot move candidates forward, and strong people accept other offers.

Quality metrics

5. Source quality judges each source by what it produces, not by how many candidates it sends. A job board that delivers 80 applicants and zero finalists is not a strong source. A referral channel that delivers 6 candidates and 3 finalists is.

6. Interview pass-through rate is stage conversion viewed as a quality check. If too many candidates fail at a specific stage, screening before that stage may be too loose. If almost nobody fails, the stage may not be testing much at all.

7. Onsite-to-offer ratio measures how many candidates who reach the final interview stage receive an offer. A low ratio usually means the recruiting team and the hiring team disagree on what a strong candidate looks like.

8. Calibration consistency tracks whether interviewers evaluate candidates the same way. One interviewer who approves everyone and another who rejects unless the candidate can rebuild the internet from memory are both problems. This metric finds the drift.

Outcome metrics

9. Offer acceptance rate is the share of extended offers that candidates accept. When it drops, the cause is almost always upstream: compensation, slow decisions, weak closing, or a process that surprised the candidate at the end.

10. Time to fill measures how long a role stays open from requisition approval to accepted offer. It is useful as a signal and dangerous as a target, because pressure to fill fast can push teams to lower the bar.

The ten metrics at a glance

#	METRIC	GROUP	WHAT IT ANSWERS
1	Stage conversion	Funnel	Where do candidates drop?
2	Time in stage	Funnel	Where does work sit still?
3	Response rates	Funnel	Is outreach working? Is the team following up?
4	HM feedback speed	Funnel	How fast are decisions made?
5	Source quality	Quality	Which sources produce finalists?
6	Pass-through rate	Quality	Is screening calibrated?
7	Onsite-to-offer ratio	Quality	Do we agree on what good looks like?
8	Calibration consistency	Quality	Are interviewers aligned?
9	Offer acceptance	Outcome	Are we closing well?
10	Time to fill	Outcome	How long do roles stay open?

In the next three chapters, we will look at each group in detail: how to measure these metrics, what the benchmarks say, and what to do when the numbers point to trouble.

Funnel metrics that show where hiring slows down

Once you know what a useful metric looks like, the next step is to find where the process slows down. That is what funnel metrics are for. These four metrics cover the core of every recruiting pipeline: where candidates drop, where they wait, whether communication works on both sides, and how fast decisions get made.

1. Stage conversion

Stage conversion is the share of candidates who move from one stage to the next. It shows where the pipe narrows and helps you find the stages that need attention.

In practice, the tricky part is defining stages consistently. Tech hiring pipelines vary by role. A backend engineer search might include a technical screen, a system design interview, and a coding round. A QA engineer search might skip system design but add a test automation exercise. Some roles include take-home assignments, others do not. A senior hire might have an extra bar-raiser conversation. These differences are normal and expected.

The problem is that when every role has a slightly different pipeline, comparing conversion rates across roles becomes difficult. If one pipeline has three interview steps and another has five, their stage conversion numbers are not directly comparable.

An easy way to deal with this in practice is to use high-level reporting buckets that stay the same across all roles, even when the detailed steps inside them differ. For example: New, Screening, Interview, Offer, Hired. Inside the Interview bucket, one role might have two steps and another might have five, but for reporting purposes they all roll up into the same stage. This gives you a consistent funnel you can compare across roles, teams, and time periods while still allowing each search to run the process it actually needs.

Keep those bucket definitions fixed. If you rename stages or merge steps halfway through the quarter, your trend line turns into fiction.

Once the stages are stable, the numbers tell a clear story. If 100 sourced candidates become 20 replies, your outreach response rate is 20%. If 20 recruiter screens become 14 hiring manager screens, that conversion is 70%. If 14 candidates enter the interview stage and only 2 move forward, something in that stage needs a closer look.

Here are some common patterns and what they usually point to:

What low conversion between stages usually means

STAGE TRANSITION	LOW CONVERSION MAY MEAN
New → Screening	Job description attracts the wrong audience, or sourcing targets too broadly
Screening → Interview	Recruiter and hiring manager are not aligned on requirements
Interview → Offer	Interview process is too long, panel is not calibrated, or the role is poorly defined
Offer → Hired	Compensation is off, competing offers, or closing happened too late

When recruiter screens convert well but interviews do not, the problem is often alignment. The role may be poorly defined, the screening questions may not match what the interview panel tests, or the hiring manager may say "senior" and mean something different from what the recruiter understood. A sharp drop at one stage is usually not random.

2. Time in stage

Time in stage gives you the second half of the picture. Conversion tells you where candidates fail to move, while time in stage tells you where they wait. Track the median days in each stage, not just overall time to hire. A seven-day median delay between technical interview and decision is specific enough to fix.

This is where many teams miss the real bottleneck. If candidates pile up after interviews, sourcing often gets blamed because it is visible and easy to criticize, but a late-stage queue is rarely a sourcing problem. In tech hiring, SmartRecruiters' 2025 technology benchmark (<https://www.smartrecruiters.com/resources/article/technology-benchmark-recruiting-metrics/>) found companies take 10 days longer than average between interview and offer. The clog often sits with decision-making, not candidate supply.

3. Response rates

Response rates cover both sides of the conversation, and they diagnose different problems.

Candidate response rate shows whether your targeting and messaging earn a reply. If candidate response is weak, check audience and message quality before you blame compensation. A backend engineer may ignore a vague note about an "exciting opportunity" but reply to a short message that names the stack (the main technologies the team uses, such as Python, PostgreSQL, and AWS), scope, and reporting line. If open rates look fine but replies stay low, your subject line may work while the body does not. Ashby's 2026 startup hiring report (<https://www.ashbyhq.com/talent-trends-report/reports/startup-hiring>) says startup email outreach gets about an 81% open rate, but opens can flatter you — replies tell the truth.

Recruiter response rate shows whether your team follows through once a candidate engages. If a candidate replies and waits three days for a screen invite, you are losing people you already earned. This metric is easy to overlook because it sits inside the team's own workflow, but slow follow-up is one of the most common reasons good candidates disappear early.

4. Hiring manager feedback speed

Hiring manager feedback speed deserves its own line on the dashboard. Measure the median time from interview completion to written feedback or a decision. If interviewers take three days to submit notes, recruiters cannot move candidates forward, and strong people accept other offers while your team deliberates.

This metric is also one of the easiest to act on. A clear expectation — feedback within 24 hours of an interview — costs nothing and removes one of the most common delays in the process.

What each funnel metric helps you diagnose

#	METRIC	WHAT IT SHOWS	COMMON REAL CAUSE
1	Stage conversion	Where candidates drop	Weak targeting or misaligned screening
2	Time in stage	Where work sits still	Slow decisions or overloaded interviewers
3	Response rates	Whether outreach and follow-up work	Poor messaging or slow recruiter follow-up
4	HM feedback speed	How fast decisions get made	Low urgency or unclear ownership

Quality metrics that keep your pipeline honest

A busy pipeline can still be a bad pipeline. Recruiters can hit outreach targets and keep interviews moving while sending the wrong people into the process. Everyone feels productive right up to the meeting where someone declares that the market is

impossible. Often it is not the market — it is poor targeting. These four metrics help you catch that before it becomes a pattern.

5. Source quality

Do not judge a source by applicant volume. Judge it by what happens later: how many candidates from that source pass recruiter screen, reach final stage, get offers, and accept?

A source that sends 80 applicants and produces no serious finalists is not a strong source — it is a loud one. A source that sends 8 candidates and produces 3 final-round interviews may be doing the real work. In tech hiring, that happens often with referrals and internal talent. SmartRecruiters' 2025 technology benchmark says referrals make up 16% of tech hires, well above the global average (SmartRecruiters' technology benchmark (<https://www.smartrecruiters.com/resources/article/technology-benchmark-recruiting-metrics/>)).

A simple source-quality view makes this easy to track:

Judge sources by outcomes

SOURCE	ENTERED PIPELINE	REACHED FINAL STAGE	OFFERS	QUALITY READ
Job board	42	1	0	Low signal
Referral	6	3	1	High signal
Outbound search	15	2	1	Promising

Review this monthly and let the data guide where you spend sourcing time and budget. A source that consistently delivers volume but not outcomes is costing you interview hours without producing results.

6. Interview pass-through rate

This is the same stage-conversion idea from the funnel chapter, but viewed as a quality check on how well your screening and interviews identify strong candidates. If many candidates fail at the first hiring-manager interview, your recruiter screen may be too loose. If almost nobody fails, the screen may be too strict, or the interview stage may not be testing much at all.

Pass-through rates only make sense with context. A backend engineer role and a product designer role will not convert the same way. Neither will entry-level and staff-level searches (very senior individual contributor roles above senior level). Compare like with like, or the metric will lie to you politely.

7. Onsite-to-offer ratio

Onsite-to-offer ratio, or final-stage success rate, measures how many candidates who reach the final interview stage receive an offer. Here, "onsite" means the final interview round or panel, whether it happens virtually or in person.

This is one of the clearest tests of alignment between recruiting and the hiring team. If five people reach the final stage for every one offer, one of two things is usually wrong: either the recruiter is bringing in weak finalists, or the hiring team does not agree on what a strong candidate looks like.

8. Calibration consistency

Quality problems often hide inside inconsistent interviewers. One interviewer says yes to everyone, while another says no unless the candidate can rebuild half the internet from memory. Neither is helpful, and both distort the funnel.

Track interviewer patterns over time. Which interviewers pass far more or far fewer candidates than the panel average for the same role family? Which hiring teams reject many candidates at final stage after strong earlier feedback? Which interviewers leave vague notes that cannot support a decision? This is not about policing people — it is about finding drift before it damages the funnel.

LinkedIn's 2025 Future of Recruiting report (<https://business.linkedin.com/talent-solutions/resources/future-of-recruiting/>) notes that quality of hire is usually measured through outcomes such as job performance, retention, and hiring-manager satisfaction, not through a single ATS field. That matters because pipeline quality is a leading indicator, not the whole story. These four metrics help you catch quality issues early, while you can still do something about them.

When quality metrics point to trouble

- ✓ High applicant volume, weak final-stage flow
- ✓ One source dominates volume but not offers
- ✓ Strong recruiter screens, weak manager screens
- ✓ Final-stage rejection patterns vary by interviewer
- ✓ New role opens and panel cannot define must-haves

Use these metrics to start better conversations with hiring teams. Which sources create finalists? Where do strong candidates start to fail? Do we all mean the same thing when we say "strong"?

Outcome metrics that show whether the process worked

Funnel and quality metrics show where the process bends. Outcome metrics show what happened at the end. They summarize the combined effect of every earlier decision, which means when one moves in the wrong direction, you should not rush to fix the metric itself — go looking upstream.

9. Offer acceptance rate

Offer acceptance rate is the share of extended offers that candidates accept. The formula is simple:

$$\text{offer acceptance rate} = \frac{\text{accepted offers}}{\text{total offers extended}}$$

A declined offer can mean pay was off, the manager did not sell the role well, the process dragged, the team changed the brief halfway through, or the candidate never trusted the process. The number itself does not tell you which one — it tells you to investigate.

Track offer acceptance by team, role family, level, and location, because a single company-wide number hides too much. If backend engineers accept at 82% and Site Reliability Engineers (SREs, the people who keep production systems running) accept at 48%, you do not have one problem — you have a market-specific problem.

When offer acceptance drops, ask narrow questions. Are we losing people on compensation? Are offers taking too long after final interview? Are candidates meeting the manager early enough? Are we surprising them at the end with office rules, leveling, or scope? "Candidates are getting picky" is not analysis — it is a shrug in business casual.

A useful supporting signal here is **process completion**: how many candidates who start the interview process actually finish it. If candidates drop out late, your loop may be too long, too repetitive, or too slow between steps. Ashby's 2026 startup hiring report (<https://www.ashbyhq.com/talent-trends-report/reports/startup-hiring>) found that hired candidates at startups spend roughly 2.5 to 3 hours interviewing, while larger startups can spend about 21 to 29 interviewer hours per technical hire. When that gets heavy, candidates start looking for the exit before they ever reach your offer.

10. Time to fill

Time to fill is useful and dangerous in equal measure. It is useful because open roles cost time and team capacity. It is dangerous because people use it as a blunt instrument and then act shocked when quality drops.

Define it before you report it. Greenhouse's glossary (<https://www.greenhouse.com/resources/glossary/what-is-time-to-fill>) distinguishes **time to fill** (from requisition approval to accepted offer) from **time to hire** (from candidate entry into the pipeline to accepted offer). These are different numbers, and treating them as the same will quietly break your trend lines.

Use benchmarks as a reference, not a weapon. SmartRecruiters' 2025 benchmark report (<https://ta.smartrecruiters.com/rs/664-NIC-529/images/Recruitment-Benchmarks-2025-Report.pdf>) reports a global median time to hire of 38 days. That is a starting point for conversation, not a target your team must hit regardless of role complexity.

In tech hiring specifically, watch the gap between final interview and offer. SmartRecruiters' 2025 technology benchmark (<https://www.smartrecruiters.com/resources/article/technology-benchmark-recruiting-metrics/>) found that technology companies take 10 days longer than average at this step. That delay is rarely about candidate supply — it usually points to slow debriefs, unclear interview signals, or a compensation process that starts too late.

How to read outcome metrics

#	METRIC	WHAT IT TELLS YOU	WHAT IT DOES NOT TELL YOU
9	Offer acceptance	Whether the team closed well	Which exact issue caused declines
10	Time to fill	How long the req stayed open	Whether speed improved quality

When you discuss these numbers with hiring managers, stay out of defensive mode. Say what the metric shows, what it cannot show, and what you want to inspect next. For example: "Our offer acceptance rate for this role is lower than the team average. Most declines mention compensation and timeline. I want to review our close, our offer package, and how long we take after final interview."

Building a dashboard your team will actually use

A recruiting dashboard fails when it tries to prove that recruiting is busy. It works when it helps you make a better decision this week.

Build your dashboard around movement and quality. I use a simple rule: each metric needs five things before it earns a place on the page. It needs a definition, one owner, one reporting cadence, one benchmark, and one question that forces action. If a metric cannot answer "what do we do next?" it belongs in a spreadsheet graveyard.

Keep the cadence boring

Review your dashboard weekly if hiring volume is high, or every two weeks if it is not. Use the same format every time: look at the trend, identify one exception, assign one action. That is the entire ritual. Anything more elaborate becomes a meeting people dread and a dashboard people stop trusting.

A working template

Here is a simple template that turns the ten metrics into something you can actually run:

Working dashboard template

#	METRIC	OWNER	CADENCE	ACTION QUESTION
1	Stage conversion	Recruiting lead	Weekly	Are we screening too loose or too tight?
2	Time in stage	Ops or lead recruiter	Weekly	Which stage needs a rule change?
3	Response rates	Recruiter	Biweekly	Is the message weak, or is the team slow to follow up?
4	HM feedback speed	Hiring manager	Weekly	Who is blocking decisions?
5	Source quality	Ops or lead recruiter	Monthly	Which sources earn more budget?
6	Pass-through rate	Recruiting lead	Biweekly	Where is screening misaligned?

#	METRIC	OWNER	CADENCE	ACTION QUESTION
7	Onsite-to-offer ratio	Recruiting lead	Biweekly	Do we agree on what good looks like?
8	Calibration consistency	Recruiting lead	Monthly	Which interviewers have drifted?
9	Offer acceptance	Recruiting lead	Weekly	Are we losing on comp, speed, or trust?
10	Time to fill	Recruiting lead	Weekly	Where is delay growing?

That is enough. Do not add seven more tabs because the ATS can export them.

The action question is the part most teams skip

Without it, the dashboard becomes historical decoration. "Offer acceptance is down" is not a decision. "Offer acceptance is down for senior platform roles in one region, and we need tighter compensation ranges before interview loop starts" is a decision.

Keep the dashboard cheap to maintain. If a metric takes hours to reconcile and nobody changes behavior because of it, cut it. Real dashboards have to survive contact with real life.

Dashboard test

- ✓ Each metric has one owner
- ✓ Each metric has one fixed definition
- ✓ Benchmarks are marked as local starting points
- ✓ Every metric has an action question
- ✓ Anything nobody uses gets removed

If your dashboard helps you see the bottleneck, decide what to change, and check whether it worked, keep it. If it mainly creates debates about whose number looks bad, start over.

What other metrics can be useful?

The ten metrics in this book are the foundation. They give you enough signal to run a healthy hiring process without drowning in reports. Once the foundation is solid, a handful of other metrics become worth adding — not because every team needs them, but because each one answers a question that the core ten do not.

Quality of hire

Quality of hire is the metric every leadership team wants and few can define. It tries to answer the most important question in hiring: are the people we hire actually good?

The challenge is that quality only shows up over time. You cannot know whether a hire was strong until they have been in the role for six to twelve months. LinkedIn's 2025 Future of Recruiting report (<https://business.linkedin.com/talent-solutions/resources/future-of-recruiting/>) notes that talent teams most often measure it with a combination of job performance ratings, new-hire retention, and hiring manager satisfaction. I treat it as a composite, not a single number pulled from an ATS field.

Add this once your funnel is healthy and you have enough hires to look at patterns. Before then, you do not have enough data to tell signal from noise.

Pipeline aging

Pipeline aging tracks candidates or roles that have sat in the pipeline too long without moving. Unlike time in stage, which measures delay inside one specific step, pipeline aging shows where work is getting old across the overall process.

This becomes useful when your team is running many roles in parallel and candidates start falling through the cracks. A simple report that flags any candidate who has been in the same stage for more than, say, ten days can surface problems faster than waiting for the weekly review.

Requisition load

Requisition load is the number of open roles assigned to each recruiter. It is a capacity metric, not a pipeline metric. It answers a question the core ten cannot: is anyone drowning?

When requisition load is uneven — one recruiter juggling twelve roles while another has three — the overloaded recruiter's metrics will drop across the board, not because they are underperforming but because they cannot give each role enough attention. Watch this when time in stage and response rates slip for one person but not the team.

Cost per hire

Cost per hire is a finance-friendly metric: total recruiting costs divided by the number of hires in a given period. It includes sourcing tools, agency fees, job board spend, referral bonuses, and recruiter time.

This metric matters most to the people holding the budget. It can surface wasteful sources and help justify or trim recruiting investment, but it tells you nothing about whether the hires were any good. Pair it with source quality and quality of hire before making decisions from it.

Candidate experience score

Candidate experience score, usually collected through a short survey after the process ends, measures how candidates felt about the way you ran the loop. A low score is a warning sign that shows up before offer acceptance drops.

The catch is that most teams do not measure it well. Response rates are low, surveys ask the wrong questions, and the results end up in a slide deck no one acts on. If you are going to track this, commit to reading and responding to the feedback, not just collecting it.

When to add these

I would start with the ten metrics in this book. They will tell you where your process is healthy and where it is not. Add these other metrics only when a specific question comes up that the core ten cannot answer — not because a vendor demo made them look interesting. Every metric you add costs attention, and attention is the scarce resource in recruiting, not data.

Conclusion

The point of these ten metrics is not to build a prettier report. It is to see the process clearly enough to improve it.

If I had to reduce the whole book to a few habits, they would be these: track conversion instead of volume, track time in stage instead of vague delay, judge sources by outcomes instead of noise, and always ask what the number is telling you to do next. When offer acceptance drops, go upstream. When final-stage interviews pile up, look at feedback speed and decision quality. When outreach opens look strong but replies stay weak, fix the message, not the spreadsheet.

Here is a concrete next step. Open your current dashboard this week. Cut every metric that does not help you act, then take the ten from this book and set them up properly: one definition, one owner, one cadence, and one question beside each. It will take an

afternoon. The result is a dashboard that actually answers the question you started with — when are we finally closing that ML Engineer role? — and tells you exactly what to do about it.

If this was useful, the other books in the Tech Recruitment 101 series will help you build the rest of the system around it. Metrics matter, but only when they serve better hiring decisions. That is the standard I would keep.

About This Book

About the Author



Natalia Romanova

Natalia is the founder of Talantrix, an applicant tracking system for tech recruiters. She began her career as a software engineer, later moved into engineering management, and eventually led several software companies, including ones she founded. Along the way she repeatedly had to build technical recruiting capability, and over time ended up learning a lot about hiring engineers—more than she ever expected to.

About Talantrix

Talantrix is an intelligent applicant tracking system built by recruiters, for recruiters. Designed specifically for technical hiring, it helps recruitment teams automate the busywork — from CV parsing and skill-based search to pipeline management — so they can focus on what matters most: connecting with top engineers and making great hires.

Learn more at <https://talantrix.com> (<https://talantrix.com>)

The Tech Recruitment 101 Series

A practical book series for recruiters who work with engineering teams. Each volume tackles a specific challenge in technical hiring — from screening candidates to understanding code — with clear explanations, real-world examples, and actionable frameworks.

Revision 2.3 — 6 April 2026

Tech Recruitment 101

A practical book series for recruiters who work with engineering teams. Each volume tackles a specific challenge in technical hiring — from screening candidates to understanding code — with clear explanations, real-world examples, and actionable frameworks.