

A CONCORD CONSULTING GUIDE

Beyond the Resume

Why executive hires fail, and how to build ones that last

Most executive mis-hires aren't a talent shortage. They're a modeling failure: nobody defined what success in the role actually required before the search started. This guide walks through the pattern behind repeat mis-hires, why credentials alone can't predict performance, who actually owns a search timeline, and why retention, not time-to-fill, is the metric worth tracking.

Read time: About 12 minutes

Who this is for: Boards, CEOs, CHROs, and CPOs evaluating a leadership hire, especially a role that's been filled more than once

In This Guide

PART 1

The Mis-Hire Pattern

Why the same leadership role keeps getting filled, and vacated, and what's actually driving it.

PART 2

Credentials vs. Capability

What a resume can't tell you, and why pedigree is a weak predictor of performance.

PART 3

The Job Model Methodology

How defining success before sourcing changes the entire search, and how we build one.

PART 4

Search Timeline Ownership

Why the slowest part of most searches isn't the market, it's the people approving each stage.

PART 5

Retention as the Real KPI

Why time-to-fill is the wrong number to optimize, and what to track instead.

Part 1: The Mis-Hire Pattern

When a leadership role gets filled, and vacated, two or three times in a few years, it's tempting to chalk it up to a tough market, a shallow talent pool, or bad luck. In practice, it's almost always a specific, diagnosable pattern. The market rarely deserves the blame it gets.

We've seen this pattern often enough to name its four most common root causes. Rarely does a repeat mis-hire trace back to just one. Usually it's two or three compounding each other, which is exactly why it keeps recurring instead of resolving itself on the next try.

An unclear mandate

Most searches start with a job description: a list of responsibilities, years of experience, and a handful of preferred qualifications. What they rarely include is an agreed definition of what success in the first year actually looks like. When that definition is missing, or only informally understood by whoever happened to be in the room when the search started, different stakeholders end up evaluating the eventual hire against different, unstated standards. The board wanted a change agent. The CEO wanted stability. The hire, reasonably, tried to satisfy both and satisfied neither. This is rarely a failure of the person hired. It's a failure to agree on the target before aiming.

Assessment that doesn't test the right thing

Resume review and conversational interviews are efficient, and they measure the wrong things for a senior role. A resume shows credentials and pedigree. A conversational interview measures how well someone performs in an interview, which is itself a specific, coachable skill only loosely related to how they'll perform in the job. Structured assessment, evaluating a candidate's actual track record of decisions made under real pressure, with specific situations rather than general impressions, catches what conversational interviews miss. Without it, a hiring team is often selecting for confidence and articulateness, not the judgment the role actually requires.

Stakeholders who were never aligned

A search can survive a difficult market. It rarely survives a board, CEO, and hiring committee who each walked in wanting something different and never resolved it. Misalignment usually doesn't surface during the search itself, everyone can agree in the abstract that a strong candidate is strong. It surfaces six months in, when the new hire's decisions start reflecting one stakeholder's priorities and disappointing another's, and the search that looked successful at signing starts to unravel.

Missing transition support

Even a well-modeled, well-assessed hire can fail without support in the first 90 days. Executives dropped into a new organization with no structured onboarding, no defined early wins, and no regular check-in are far more likely to misjudge the culture, alienate a key stakeholder, or simply leave, than the assessment process would have predicted. The mis-hire narrative often blames the person hired, when the actual failure point was what happened, or didn't happen, after they started.

Self-check: has this specific role been filled more than once in the last three years? If so, which of the four causes above sounds most familiar? That's usually the place to start fixing it, not the next search itself.

Part 2: Credentials vs. Capability

A resume tells you where someone has been. It tells you almost nothing about how they'll perform under the specific pressures of your business. This distinction sounds obvious stated plainly, and yet most hiring processes, including many run by search firms, are built almost entirely around credentials: titles held, company names, years of tenure, degrees earned.

Two candidates can have nearly identical resumes, similar titles at similar-sized companies, similar tenure, similar academic pedigree, and produce wildly different outcomes in the same seat. The reason is that the thing which actually predicts performance, judgment under ambiguity, how someone makes decisions when the data is incomplete and the stakes are real, rarely shows up on paper. It shows up in how someone has actually handled specific hard moments, and that requires asking about specific hard moments, not scanning a career history for the right shaped credentials.

Why pedigree bias is so persistent

Pedigree is easy to evaluate and feels objective: a candidate either worked at a recognizable company or didn't, either has the target years of experience or doesn't. Capability is harder to evaluate because it requires structured effort, specific behavioral questions, reference checks that go beyond a friendly call to someone hand-picked by the candidate, and a clear model of what you're actually testing for. Under time pressure, hiring teams default to the easier signal, even when everyone involved would agree, if asked directly, that pedigree isn't what they actually care about.

What actually predicts performance

The traits that predict success in a specific leadership role are almost never generic. "Strong communicator" and "strategic thinker" appear on nearly every executive job description and explain almost nothing about who will succeed in this particular seat, at this particular company, under this particular set of pressures. The traits that matter are specific: comfort operating alongside a founder who's still deeply involved in the function, judgment about when to escalate versus when to decide alone, the ability to rebuild trust with a team that's been through two prior leadership changes. None of that is visible on a resume. All of it is discoverable through structured assessment, if you know what you're testing for before the interviews start.

Self-check: for your last few leadership hires, could you name the specific traits, not generic ones, that predicted whether they'd succeed? If not, the search likely leaned on credentials by default, not by decision.

Part 3: The Job Model Methodology

Everything in the first two sections points to the same conclusion: most mis-hires trace back to something that was never defined clearly enough before the search began. The fix isn't a better interview question or a sharper recruiter. It's building a job model before anyone is contacted at all.

A job model is a structured definition of what success in a specific role actually requires: the traits, skills, and experience that matter, built with the hiring team before sourcing begins, not inferred afterward from whoever happens to apply or however a search firm's existing network happens to be composed.

How we build one

Building a job model means sitting down with the board, CEO, or hiring committee and getting specific about questions a job description never answers: What kind of decisions will this person make weekly, and under what pressure? What has caused past hires in this seat, or similar ones elsewhere, to succeed or fail? Which of those factors are teachable on the job, and which need to already be present on day one? Where does this role sit in the org's specific politics, and what judgment does navigating that require?

The output isn't a longer job description. It's a working standard: a short, specific list of the traits and experience markers that separate someone who will succeed in this exact seat from someone who merely looks qualified for it.

What changes once the model exists

Once a job model exists, it becomes the standard every candidate is measured against, not their resume, not their pedigree, but whether the specific evidence in their background suggests they have the traits the model calls for. Structured assessment stops being a generic competency framework applied to every search the same way, and becomes a direct test of the specific things that matter for this role. Reference checks stop being a formality and start probing the exact situations the model flagged as predictive.

This is also, worth being direct about, not a claim that we reach candidates who aren't actively looking. We search the active market, people who are open to a move or already engaged in one. The rigor is in how deeply each candidate is assessed against the job model, not in how wide or discreet the sourcing net is cast. A precise model searched narrowly beats a vague one searched widely.

Self-check: before your next leadership search, could your team write down the job model in a page or less, specific traits, not generic ones? If that page doesn't exist yet, that's the actual starting point, before any candidate is contacted.

Part 4: Search Timeline Ownership

When an executive search drags on, it's easy to point at the market: not enough qualified people, everyone's happy where they are, compensation expectations have shifted. Sometimes that's true. Just as often, the real bottleneck is internal, and it's rarely named as such because it's less comfortable to say out loud.

Where searches actually stall

Indecisive interview panels are the most common cause. A candidate who's genuinely strong can lose momentum, and eventually interest, waiting two weeks for five stakeholders to independently schedule a 45-minute call and another week for their feedback to be collected and synthesized. Slow feedback loops compound this: a candidate left waiting without a clear next step reasonably starts entertaining other conversations. And a board or hiring committee that can't align on the mandate, discussed at length in Part 1, will relitigate that disagreement at every single finalist stage, turning what should be a decision meeting into another round of the same debate.

None of these causes are about candidate availability. They're about the speed and discipline of the people running the process on the client side, and a search can only move as fast as those people allow it to.

What ownership actually looks like

The boards that run the tightest searches aren't the ones with the easiest roles to fill. They're the ones who treat their own responsiveness as part of the timeline, not a variable outside their control. That means committing to specific feedback windows before the search starts, not negotiating them mid-process. It means designating one person, not a committee, to own scheduling and momentum. And it means resolving mandate disagreements once, in Part 1's job model conversation, rather than relitigating them with every finalist.

A well-run search firm can shorten a timeline. It cannot fix a client-side bottleneck the client hasn't acknowledged. The single highest-leverage thing a board can do to speed up its own search is to be honest, in advance, about where the last one actually got stuck.

Self-check: in your last executive search, what caused the longest delay between finalist interviews and an offer? If the honest answer involves your own team's scheduling or alignment, that's fixable before the next search starts.

Part 5: Retention as the Real KPI

Time-to-fill is easy to measure and easy to report to a board. It's also the wrong number to optimize if the hire doesn't last, and most organizations optimize for it anyway, because it's the number they're used to tracking, not because it's the number that predicts a good outcome.

A search that closes in six weeks and produces a hire who leaves in ten months has cost far more, in re-search fees, in lost momentum, in the team's confidence in leadership, than a search that took fourteen weeks and produced a hire who's still thriving three years later. Almost no board actually tracks this comparison, because time-to-fill is recorded the day an offer is signed, and retention outcomes only become visible well after the search team has moved on to the next mandate.

What to track instead

Eighteen-month retention is the single most useful number most boards aren't tracking: is this person still in the role, and still effective, a year and a half after starting? It's far enough out to filter out early honeymoon effects and close enough to still be actionable, so a pattern can be caught and addressed before a third or fourth search repeats it. Alongside it, a simple "would rehire" score from the hiring manager or board chair at the 12-month mark captures something numbers alone miss: whether, with hindsight, this was the right call.

Why this closes the loop

Retention tracking isn't just a better scoreboard. It's the feedback signal that makes the rest of this guide self-correcting. If a job model produced a hire who's thriving at 18 months, that model worked, and it's worth reusing the next time a similar role opens. If it didn't, that's specific, useful information about what the model missed, not just a vague sense that the market was tough this time. Without retention as a tracked KPI, a board has no way to tell the difference between a search process that's actually improving and one that's simply gotten faster at repeating the same mistake.

Self-check: for your last three executive hires, do you know where each of them stands right now? If you had to guess rather than check a record, that's a sign retention isn't being tracked as a real metric yet.

Where to Go From Here

None of the five patterns in this guide are unusual, and none of them require a bigger candidate pool to fix. They require defining what success actually looks like before the search starts, testing directly for it, keeping stakeholders aligned through the process, owning the pace of the search internally, and measuring the outcome by whether the hire lasted, not by how fast the role closed.

If you want to see where your own situation stands, three short tools can help without requiring a call:

- Executive search readiness checklist — is a structured search the right model for this role?
- Leadership vacancy risk assessment — how exposed is the organization right now, with this seat open?
- Mis-hire pattern diagnostic — for a role that's been filled more than once, what's actually driving the pattern?

Ready to talk through a specific mandate?

[Book a confidential conversation with Concord Consulting.](#)