

TECH & AI CASE STUDY - RECRUITING & STAFFING - 2026

An AI Copilot That Cut Time-to-Shortlist 65%

Based on a delivery engagement with Larkspur Staffing

-65%

TIME-TO-SHORTLIST

+41%

RECRUITER SUBMISSIONS

+18%

QUALIFIED SUBMITTAL RATE

-27%

COST PER HIRE

A fillable practice edition: the full story, 8 actions,
6 reflections and a 90-day plan - work it with your team.

Free - no email required
emersontechnology.com

MAKE THIS PRACTICE BOOK YOURS

Work it, do not just read it

This practice book turns the Larkspur Staffing case study into a working session for your team. Read the full story first - what happened, what was built and what it changed. Then tick the 8 actions you can commit to, write short answers to the 6 reflections, and finish with a 90-day plan. Every field in this PDF is fillable: open it in Adobe Acrobat Reader, Chrome, Edge or macOS Preview, type directly into the boxes and save.

YOUR NAME	ROLE
COMPANY	TEAM / DEPARTMENT
DATE	EMAIL (OPTIONAL)

MY GOAL FOR THIS PRACTICE SESSION

- 01 The story** - What happened at Larkspur Staffing - before, during and after the project.
- 02 The results** - The numbers that changed, plus industry benchmarks from public research.
- 03 The 8-action checklist** - Tick boxes as your organization completes each step. Owner and date fields included.
- 04 The 6 reflections** - Fillable questions that pressure-test your own roadmap.
- 05 My 90-day plan** - Three fillable phase cards that turn this case study into your roadmap.

RUN IT AS A 90-MINUTE TEAM SESSION

25 min READ Read the story together. Ask everyone to note the one number that surprised them most.	40 min PRACTICE Work the 8-action checklist and the 6 reflections. Assign an owner and a date to every action you commit to.	25 min COMMIT Fill in the 90-day plan. One goal per phase, one number that tells you it worked. Book a review date before you leave the room.
--	--	---

THE STORY

What happened, and why it worked

01 The Company

Larkspur Staffing is a mid-market recruiting firm that places roughly 1,300 candidates a year into healthcare and engineering roles for hospital systems, outpatient clinic groups, medical device manufacturers and industrial employers. The firm operates with 40 recruiters organized into specialty pods, each pod owning its own client accounts, requisition queue and submission targets. Its commercial model is unforgiving in a specific way: the firm that produces a credible, client-ready shortlist first usually wins the placement fee, so speed and selectivity are not competing priorities — they are the same objective measured from two directions. By 2026, Larkspur's leadership had concluded the firm was losing that race not because its recruiters lacked judgment, but because nearly all of their time was going to work beneath their judgment.

The company's applicant tracking system had been built in-house on Next.js and Prisma, hosted on AWS, and had grown into a genuine system of record: candidate histories, client notes, interview feedback and placement outcomes all lived in one place. That engineering investment turned out to matter enormously, because it gave the AI copilot a home to live in instead of forcing recruiters into yet another tab. Leadership had watched competitors pilot screening tools and had sat through plenty of pitch decks; what they wanted was throughput — more client-ready shortlists per recruiter per week — without diluting the placement quality that kept clients renewing year after year.

02 Where Things Stood

Before Emerson Technology engaged, recruiters spent 71% of their working week on two activities: reading resumes and writing first-touch emails. Eye-tracking research published by Ladders found that recruiters spend an average of 7.4 seconds scanning a resume before deciding whether to keep reading, and Larkspur's recruiters were no exception — but Larkspur's requisition volume made those seconds compound into days. Average time-to-shortlist ran 6.2 days from requisition to client-ready list, while clients increasingly demanded shortlists inside 48 hours. Every day of that gap was a day a competing firm could submit first and take the fee.

Outreach was the second leak. First-touch emails came from a small library of templates written years earlier, and candidate response rates sat at 14%. Personalization was reserved for the handful of star candidates a recruiter already knew; everyone else received the same paragraph with a name swapped in. The arithmetic was brutal: at a 14% response rate, roughly six of every seven outreach emails produced nothing, pipelines thinned silently between application and interview, and recruiters compensated by sending more volume rather than better messages.

The financial consequence showed up in the revenue Larkspur never booked. Leadership estimated the firm had walked away from an estimated \$3.4M in fees over the trailing year — requisitions declined or quietly lost to faster competitors — not because Larkspur lacked candidates, but because it could not staff the requisitions it already had fast enough to chase more. Growth had become a throughput problem, and throughput was capped by how many resumes forty people could physically read.

03 The Challenge

The deeper problem was variance. Screening criteria lived in individual heads, and they varied by recruiter: the same resume could pass on Monday and fail on Friday depending on who opened it and how their week was going. Client accounts inherited the inconsistency — one account manager's strong submittal was another's not-competitive — and the firm had no way to prove its standards to a client beyond assertion. The industry was moving at the same time: SHRM research reportedly finds about one in four US organizations now use AI to support HR activities, with recruiting and hiring leading the list, and Larkspur was watching the firms in that quarter compound an advantage in responsiveness that marketing spend could not buy back.

Any solution had to satisfy two constraints leadership stated on day one. First, judgment stays human: a machine could rank, summarize and draft, but people would interview, decide and send — partly for compliance reasons, and partly because screening judgment was the actual product Larkspur sold. Second, the tool had to live inside the ATS where recruiters already worked. Leadership had seen peer firms buy standalone AI tools that died within a quarter of context-switching fatigue, and it refused to fund another one.

04 Discovery and Diagnosis

Emerson Technology opened with a two-week discovery sprint across Larkspur's ATS data model, resume corpus, outreach templates and placement history. The most valuable find was three years of historical placements: every submitted candidate, the client feedback attached to each submission, and the final outcome. That history was, in effect, a labeled dataset of what good looked like, which meant any screening model could be validated against commercial reality before it ever touched a live requisition. The sprint also confirmed the arithmetic: with 71% of recruiter time going to reading and first-touch writing, even a partial recovery of those hours translated directly into more interviews and client calls per week.

The diagnosis quantified the variance problem. Emerson took thirty anonymized resumes from recent requisitions and had recruiters across the pods rate them pass or fail against the job description; agreement was far weaker than pod leads expected, with several resumes splitting the room exactly along pod lines. The audit further showed that none of the funnel numbers — time-to-shortlist, submittal rate, response rate — were tracked by stage, role or recruiter, which meant nobody could distinguish a model problem from a simple follow-through problem once automation arrived. Harvard Business Review has reported Unilever's experience with AI-assisted graduate screening: roughly 100,000 hours of recruiter time saved and a wider diversity of candidates reaching interviews. The lesson Larkspur took from that example was specific — the gains came when screening ran against structured criteria, not when a model was turned loose on resumes unattended.

05 The Solution

Emerson Technology designed and built an AI copilot embedded directly inside Larkspur's existing Next.js and Prisma ATS on AWS. The governing design principle was that the machine drafts and the recruiter decides; the copilot produces ranked shortlists with written reasoning, never verdicts. The architecture has three layers.

Structured scorecards

Every high-volume role family was converted into a weighted scorecard before any model ran: licenses and certifications, years of relevant experience, setting exposure, shift flexibility and compensation fit, each weighted to match what clients actually rejected candidates for. The LLM screening service parses every incoming resume, maps the candidate's experience against the relevant scorecard, and returns a ranked list with a written rationale and a confidence level for each candidate. Recruiters approve, adjust or reject each recommendation, and nothing moves to a client without their decision. The scorecard layer is what turned screening from taste into criteria — the same criteria for every recruiter, on every resume, every day of the week.

Validation and guardrails

Before launch, a Python evaluation harness replayed three years of historical placements through the ranking model and compared its ordering against what clients had actually accepted. Where the model ranked known winners poorly, the team corrected the scorecard weights and re-ran the harness — before candidates could ever feel the error. Screening output and generated text now run through a monthly audit for age, gender and school-pedigree signals, and scorecard weights are version-controlled so every screening change stays auditable and reversible. These guardrails were not ceremony; they were the condition Larkspur's enterprise clients would eventually ask to see evidence of.

The recruiter experience

The copilot lives inside the ATS as a panel on the requisition view, not as a separate product. Outreach drafts pull specifics from the candidate's actual background — the specific license, the unit or plant environment they have worked in, the certification the client requires — instead of a template, and every send requires recruiter approval. Funnel analytics track time-to-shortlist, submittal rates and response rates by role and by recruiter, giving pod leads their first stage-level view of where pipelines thin. The workflow inversion is the point: recruiters now start from a ranked, reasoned shortlist and spend their recovered hours on interviews and client conversations.

06 The First 90 Days

Weeks one and two converted the ten highest-volume role families into scorecards, working with pod leads so the weights reflected client rejections rather than recruiter folklore. Weeks three through six ran the pilot: two pods, twelve recruiters, with the copilot screening in shadow alongside the normal process. Every recommendation was logged against the human decision, and the weekly disagreement review became the pilot's most valuable ritual — it caught scorecard weights that over-penalized career changers, and it gave recruiters visible evidence that the system earned its confidence rather than demanding it.

Weeks seven through ten rolled the copilot out to all 40 recruiters: a 45-minute onboarding for each, a one-page playbook describing the exception workflow, and a named contact for questions during the first month. The team agreed on a simple rule for the transition: when a recruiter rejected a recommendation, the reason was logged, and sound rejections fed directly into scorecard tuning. Weeks eleven and twelve switched on personalized outreach, turned up the response-rate dashboards, and established the monthly bias-audit cadence and internal equity reporting.

By the end of the first quarter, screening variance had effectively disappeared — the same resume produced the same verdict regardless of who opened it — and new recruiters were ramping faster because the scorecards encoded institutional knowledge that previously took a year to absorb. The weekly shortlist meeting stopped relitigating criteria and started debating which candidates to prioritize for client calls.

07 The Results

Within two quarters the improvements were unambiguous. Time-to-shortlist fell from 6.2 days to 2.2 — a 65% reduction — putting Larkspur inside the 48-hour window its clients were demanding. Recruiters' monthly submissions rose 41% as reading hours converted into submittals, and qualified submittal rates improved 18% as the scorecards enforced consistent criteria across all 40 recruiters. Cost per hire fell 27%, response rates on personalized outreach doubled from the 14% baseline within the first quarter, and clients noticed the faster shortlists before Larkspur got around to telling them.

METRIC	BEFORE	AFTER
Time-to-shortlist	6.2 days	2.2 days (-65%)

METRIC	BEFORE	AFTER
Recruiter submissions	-	+41%
Qualified submittal rate	-	+18%
Cost per hire	-	-27%

The qualitative changes compounded the quantitative ones. Recruiters spend the recovered hours on interviews and client calls, which is where placement decisions are actually won. Screening conversations with clients changed character: account managers can now show the scorecard behind a shortlist, which turned trust-us claims into evidence. The firm also stopped leaving fees on the table — with time-to-shortlist inside client windows, requisitions leadership once declined for capacity reasons are back in play, which is where the recovered share of that \$3.4M is starting to show up.

08 What Your Business Should Do

The pattern behind Larkspur's result is repeatable without a research lab. Start by writing scorecards before screening anything: convert each role into weighted requirements before a single resume gets read, because scorecards turn screening from taste into criteria a machine and an auditor can both apply. Validate the model against your own historical placements before it goes live — replay past hires and known good placements, and if it ranks known winners poorly, fix that before candidates ever feel it. Keep humans on every decision: let AI rank, summarize and draft, while people interview, decide and send. That division is both the compliance baseline and, in hiring specifically, the actual quality advantage.

Then instrument the funnel end to end — time-to-shortlist, submittal rate and response rate by stage, role and recruiter — because without stage-level metrics you cannot tell model quality apart from simple follow-through. Audit generated text and rankings for biased signals monthly, and publish equity metrics where lawful, because measurement is how AI screening stays an improvement instead of becoming a liability. Personalize outreach from the candidate's actual background rather than a template; generic templates produced your current response rates, and specific outreach is the measurable fix. And integrate inside your ATS rather than shipping a separate tool, because context-switching kills adoption faster than any model shortcoming ever will.

Close by asking the questions Larkspur's leadership still reviews quarterly: how many hours a week your recruiters spend on work a model could draft for approval, what every day of screening delay costs in placements lost to faster competitors, whether all your recruiters apply the same criteria and how you would prove it, which part of screening is genuine judgment versus pattern matching you have always done by hand, whether you could show a client evidence of fairness today, and what your response rates would look like if every first-touch email referenced the candidate's real background. Work through the eight actions in this practice book with your team, and revisit those reflection questions each quarter to keep the capability honest.

REAL-WORLD PROOF & BENCHMARKS

Unilever has publicly reported that AI-assisted screening for graduate hiring saved roughly 100,000 hours of recruiter time and widened the diversity of its interview pool. - Harvard Business Review

Eye-tracking research published by Ladders found recruiters spend an average of just 7.4 seconds scanning a resume before deciding whether to keep reading. - Ladders, 2018

SHRM research reportedly finds about one in four US organizations now use AI to support HR activities, with recruiting and hiring leading the list. - SHRM, 2024

THE PRACTICE PAGES

The 8-action checklist

Work through the list with your team. Tick each box as your organization completes the step, and fill in an owner and a target date so every action has a name and a deadline.

☐

1. Write scorecards before screening

Convert each role into weighted requirements before any resume gets read. Scorecards turn screening from taste into criteria a machine and an auditor can both apply.

OWNER

TARGET DATE

STATUS

--	--	--

☐

2. Keep humans on decisions

Let AI rank, summarize and draft — people interview, decide and send. Keeping judgment human is both the compliance baseline and the actual quality advantage in hiring.

OWNER

TARGET DATE

STATUS

--	--	--

☐

3. Validate against historical placements

Test the model on past hires and known good placements before going live. If it ranks known winners poorly, fix that before candidates ever feel it.

OWNER

TARGET DATE

STATUS

--	--	--

☐

4. Audit for biased language

Screen generated text and rankings for age, gender and school-pedigree signals every month. Bias reviews are becoming table stakes for enterprise clients and regulators alike.

OWNER

TARGET DATE

STATUS

--	--	--

☐

5. Personalize outreach from resumes

Reference the candidate's actual experience in every first-touch draft, not just the role title. Generic templates produced your current response rates; specific outreach is the measurable fix.

OWNER

TARGET DATE

STATUS

--	--	--

☐

6. Instrument the funnel end to end

Track time-to-shortlist, submittal rate and response rate by stage, role and recruiter. Without stage-level metrics you cannot tell model quality apart from simple recruiter follow-through.

OWNER

TARGET DATE

STATUS

☐

7. Integrate inside your ATS

Embed the copilot inside the ATS where recruiters already work rather than shipping a separate tool. Context-switching kills adoption faster than any model shortcoming ever will.

OWNER

TARGET DATE

STATUS

☐

8. Publish equity metrics monthly

Report pass-through rates by demographic band where lawful, and act on gaps. Measurement is how AI screening stays an improvement instead of becoming a liability.

OWNER

TARGET DATE

STATUS

AFTER THE CHECKLIST: OUR TOP 3 PRIORITIES

From the eight actions above, write down the three you will start first - and the single number you expect each one to move. Three priorities with owners beat eight without any.

PRIORITY 1	PRIORITY 2	PRIORITY 3
NUMBER IT SHOULD MOVE	NUMBER IT SHOULD MOVE	NUMBER IT SHOULD MOVE

THE PRACTICE PAGES

Reflections for your team

Use these questions in your next planning session and type short answers into the boxes. The gaps you find are usually the fastest wins.

Q1. How many hours a week do your recruiters spend on work a model could draft for approval?

Q2. What does every day of screening delay cost you in placements lost to faster competitors?

Q3. Do all your recruiters apply the same criteria to the same resume — how would you prove it?

Q4. Which part of screening is genuine judgment, and which is pattern matching you have always done by hand?

Q5. If a client asked how your AI screens candidates, could you show evidence of fairness today?

Q6. What would your response rates look like if every first-touch email referenced the candidate's real background?

THE PRACTICE PAGES

My 90-day plan

Finish the session by committing to a plan. Three phases, one goal each, one number that tells you it worked. Bring this page to your next leadership review - or to a free consultation with Emerson Technology, and we will pressure-test it with you.

Days 1-30 Focus: audit and alignment

PRIMARY GOAL

KEY ACTIONS

HOW WE MEASURE SUCCESS

OWNER

Days 31-60 Focus: build and prove

PRIMARY GOAL

KEY ACTIONS

HOW WE MEASURE SUCCESS

OWNER

Days 61-90 Focus: scale and measure

PRIMARY GOAL

KEY ACTIONS

HOW WE MEASURE SUCCESS

OWNER

Tip: save a copy per team so each group keeps its own answers. Revisit the reflections every quarter - the questions stay the same even as the numbers move.

BOOK THE REVIEW BEFORE YOU CLOSE THE BOOK

Plans without a review date drift. Before you close this practice book, agree on when you will come back to these pages, who will run that session, and when you will check the numbers again. If you want an outside pair of eyes, bring this page to a free consultation with Emerson Technology and we will pressure-test the plan with you.

REVIEW DATE	REVIEWED BY	NEXT CHECK-IN
------------------------------------	------------------------------------	--------------------------------------

Ready for a transformation like Larkspur Staffing's?

Emerson Technology designs and builds the data, AI, cloud and automation systems behind results like these. On a free consultation call we will map this exact playbook to your business - your data, your constraints, your numbers.

Book a free consultation

support@emersontechnology.com

or start online: emersontechnology.com/contact