

BRIEFING 02 · FOR RECRUITERS

Less sourcing. More closing.

10 free credits. One JD. 2 minutes. No card. Written for recruiters who are tired of being pitched on AI tools that don't work.

10

FREE CREDITS TO START

1 credit = 1 qualified Japan candidate
matched to your role.

inside THIS BRIEFING

What follows, in ten sections.

A briefing written for one person: the recruiter who has been pitched too many AI tools that didn't work. The structure assumes skepticism, not enthusiasm — and gives you what you need to verify everything in two minutes on your own desk.

01 You should probably be skeptical.

Five reasons most AI recruiting tools fail — including the ones we built around.

02 The honest calendar.

Where your week actually goes when you measure it instead of estimate it.

03 What 10 credits actually show you.

The binary question your first list answers in two minutes.

04 The production data.

Q1 2026 numbers from our own desk: +38% meetings, +13.8% reply, +13.5% interview pass, +14% offer rate.

05 The math at your desk.

What the lift is worth to your personal take-home at 25%, 35%, and 45% comp shares.

06 Capability.

Five things Headhunt.AI does, named clearly.

07 The limits.

Four things it does not do — because pretending otherwise is what makes recruiters distrust AI.

08 Pushback.

Seven objections worth taking seriously, answered honestly.

09 The mechanic.

How the 10 free credits work, end to end.

10 The bottom line.

The honest take. The offer.

01 THE STARTING POSITION

You should probably be skeptical.

If you have recruited in Japan for more than two years, you have been pitched on at least four AI-powered recruiting tools. You have sat through demos that promised to transform sourcing. You have spent free trial credits on tools that produced lists of candidates your team had already worked through six months earlier. You have watched leadership get excited about the next thing while you went back to BizReach and LinkedIn Recruiter and quietly hit your number anyway.

We have sat in the same chair. ESAI is a software company born out of an agency — we built Headhunt.AI because the existing AI tools didn't work for our own desk and we needed something that did. From experience, most AI recruiting tools have at least one of these problems:

- **They improve only the longlist stage.** The names look better. Your reply rate stays the same.
- **They surface candidates you have already worked** — especially on roles you have been on for months.
- **They write outreach that reads like AI.** Candidates who deal with recruiters every week can spot it instantly, and so can your clients.
- **They are built around English-language profiles.** They underperform on Japan-specific signal — university tier, keiretsu lineage, bilingual capability, JP/EN career trajectory patterns.
- **They lock you into their database.** You build the asset; they keep it.

“Skepticism is the right default. The point of 10 free credits is that you do not have to take our word for any of this.

02 THE HONEST CALENDAR

What actually eats your week.

Ask ten recruiters where their hours go and most will tell you forty to fifty percent of the week is sourcing. The number sounds about right. Then audit a calendar in 30-minute blocks across five working days, and the real number lands consistently between 60 and 70 percent.

It is not just the searches. It is the surrounding work that doesn't show up on a calendar invite:

- Boolean searches that take 90 minutes and return 1,200 names you then have to triage.
- Profile review on candidates you will never reach out to, just to filter the noise.
- Drafting scout mails that take 15 minutes each and convert at 6–9%.
- Follow-up sequences, no-response chases, and the dozen small qualification steps before a meeting books.
- Calendar Tetris, time-zone work, candidate prep notes, the constant context-switching.
- Re-reading profiles to check whether someone you remember from six months ago is still at the same company.

All of that happens before a candidate ever joins a meeting. Then the meeting itself takes about three hours of total recruiter time once you count prep, the meeting, and the writeup. A recruiter running 30 meetings a month spends 90 hours on meetings alone — before any sourcing, client work, or closing.

02 CONT. · WHAT PAYS YOU

The number nobody quotes until they measure it.

Most recruiters spend 60–70% of the week on sourcing and qualification — not on meetings, client work, or closing. That percentage is exactly the work AI sourcing absorbs. Once it goes, the calendar capacity that was always there becomes visible — and the rest of the work is what you actually get paid for.

This is the framing that matters. Your fee does not come from sourcing. It comes from running good meetings, qualifying candidates well, and closing — work that requires judgment a tool cannot replace.

Sourcing is what is between you and the work that pays.

“The argument for AI sourcing isn’t “do your job faster.” It’s “spend less of your week on the parts that don’t pay you.”

03 THE TEST

What 10 free credits actually show you.

The free trial is structured to answer one binary question on your first search: are there candidates on this list that I haven't already worked?

Here is the mechanic. One credit equals one ranked candidate who passes your search criteria and scores 50 or higher on the ESAI Score. Sign up at headhunt.ai, get ten free credits, paste a JD for a real open role, and the platform returns a ranked list within 1–2 minutes. Each candidate carries an explicit score, written reasoning behind the score (tenure pattern, company tier, role fit, language signal, career trajectory), and an export path into LinkedIn Recruiter, Bullhorn, Salesforce-based ATSs, Zoho, or your equivalent.

Show the list to yourself. Or to a colleague who works the same segment. Look at the names. Read the scoring reasons. The question is whether the list contains candidates you haven't already approached, sourced, or written off.

1 OUT OF 10

One JD. One ranked list. One question:

“Are there names here I haven't already worked?”

If even one out of ten is new to you, the platform is finding people your current process is missing. If none are, you walk away. Cost: **zero**. Time: **2 minutes**.

When recruiters at peer agencies have run this test, 30–60% of the candidates on the list are typically new to them — even on roles their team has been working for months.

04 THE PRODUCTION DATA

The numbers from our own desk.

Headhunt.AI is the platform we use ourselves. ExecutiveSearch.AI K.K. has been running on it since 2018, and ESAI Agency K.K. uses it as the daily sourcing engine for every recruiter on the team.

In Q1 2026, we measured the lift our recruiters showed at every stage of the funnel — comparing the same recruiters working with Headhunt.AI to their previous quarters on manual sourcing.

Q1 2026 · PRODUCTION RESULTS · ESAI AGENCY · PER RECRUITER

The lift compounds at every stage of the funnel.

+38% **+13.8%** **+13.5%** **+14%**

MORE CANDIDATE
MEETINGS PER
RECRUITER

SCOUT REPLY RATE

INTERVIEW PASS RATE

OFFER RATE

Same recruiters. Same market. Same clients. Same fees. Better tools.

These are not lab numbers from a controlled trial. They are actual production results from our recruiters in Q1 2026. The comparison is the same people working in the same market on the same fee structures — the only variable is the sourcing layer.

Why a lift on every stage matters more than a single big number.

Most AI tools claim better top-of-funnel candidates. The hard part is what happens after. A bigger candidate list does not help if your reply rate stays at 6%, your interview pass rate stays flat, and your offer rate stays the same. You just send more messages and waste more time on candidates who don't convert.

Headhunt.AI lifts every stage. More candidates met means more pipeline. Higher reply rates mean less wasted outreach — and lower spam complaint rates against your domain, which matters more than most recruiters realize. Higher interview pass rates mean better fit reaching your clients. Higher offer rates mean clients trust your shortlists.

“Each lift on its own is small. Stacked across the funnel, they roughly double placement output for the same recruiter time.”

05 THE MATH AT YOUR DESK

What this means for your take-home.

The firm-level economics of AI sourcing are well covered elsewhere. What matters at the desk is what it does to your personal compensation.

Run the math on a typical Japan billing recruiter: ¥4M average placement fee (30–35% of a typical ¥10M–¥13M base), baseline output of ~0.75 placements per month, comp share of 25–45% depending on firm type and seniority. Apply the Q1 2026 lift compounded across the funnel and the math comes out to roughly +0.5 additional placements per month, conservatively.

COMP SHARE BAND	SHARE	ADDITIONAL ANNUAL TAKE-HOME
Junior to mid-level recruiter	25%	+ ¥6,000,000 / yr
Mid to senior biller	35%	+ ¥8,400,000 / yr
Senior biller / partner	45%	+ ¥10,800,000 / yr

Math: 0.5 additional placements/month × ¥4M average fee × 12 months × your comp share. Round figures, real magnitudes. Even halved, the lift is material.

“Nothing else in your control moves these numbers materially. The only lever fully inside your control is the volume of qualified meetings on your calendar.”

06 CAPABILITY

What Headhunt.AI actually does.

Five things, named clearly.

01 Scores 4M+ Japan-focused profiles against your specific role.

Not keyword matching. Real fit on role, company tier, tenure pattern, language signal, career trajectory. Each candidate gets an ESAI Score from 0–100 with explicit written reasoning.

02 Returns ranked shortlists in 1–2 minutes.

You see the top scorers first, with the full reasoning behind each ranking. You decide who to scout. The platform handles the ranking; you keep the judgment.

03 Drafts personalized scout mails in business Japanese or English.

Each draft references the candidate's actual profile, current role, and visible career signals — not template merge fields. Native-quality keigo for Japanese. Clean business English.

04 Exports cleanly into your existing workflow.

LinkedIn Recruiter CSV import for direct push into Projects. Bullhorn, Salesforce-based ATS, Zoho, or equivalent. Your data stays your data — workspace isolation at the tenant level, never shared, never used to train models for anyone else.

05 Builds your candidate database as a side effect.

Every search you run scores and structures candidates into your own database. Stale records get refreshed automatically. Old ATS dead weight turns back into a usable talent pool.

07 THE LIMITS

What this does not do.

The honest list, because pretending otherwise is what makes recruiters distrust AI tools.

01 It does not replace your judgment.

The ESAI Score ranks candidates. You decide which ones to scout, which to push to clients, and which to walk past. Strong scores are inputs, not verdicts.

02 It does not send mails on your behalf.

You paste, customize where needed, and send through your own LinkedIn account or CRM. Auto-send is exactly the kind of feature that gets agencies in trouble with candidates and platforms.

03 It does not work equally well across every niche.

AI scoring genuinely struggles on very narrow technical specialties where expertise is invisible from a public profile — certain hardware engineering subfields, deep regulatory niches. For bilingual finance, mid-tier IT, sales, commercial, supply chain, product, marketing, HR, legal, GTM, ops, and most engineering, the math works.

04 It does not automate the recruiter craft.

Meetings, qualification, candidate prep, client communication, offer negotiation, and closing all stay with you. The work that's actually paid stays human. The work that's between you and that work goes to the platform.

We make this list explicit because skeptics — correctly — assume any vendor list of features is incomplete on the things that don't work. We would rather tell you upfront where the limits are than have you discover them after you have signed.

Objections worth taking seriously.

Seven objections we hear most often, with direct answers — not deflections.

“I’ve tried AI sourcing tools. They were garbage.”

We have tried them too. Most of them are. The cleanest way to handle this objection is a side-by-side test rather than a debate. Run the same JD through Headhunt.AI and through whatever else you have — LinkedIn AI, your ATS scoring, anything. We benchmark against the leading global AI recruiting tools regularly. On Japan-specific candidate data, our quality lead is significant. The 10 free credits exist for exactly this reason. If our list is not visibly better on your roles, walk away.

“My ATS already has AI scoring.”

Most ATS scoring is keyword matching dressed up as AI. It is also limited to candidates already in your ATS — typically a small fraction of the active Japan market. Headhunt.AI scores against 4M+ Japan-focused profiles, including candidates whose records are not in your system yet. Different problem, different solution.

“I source from BizReach and LinkedIn already. I don’t need another database.”

Headhunt.AI does not replace those tools. It adds a Japan-specific scoring and ranking layer over a 4M+ profile database, then exports to LinkedIn Recruiter or your CRM for outreach. You stay in your existing workflow. We just take the longlist-and-triage work off your plate.

“I don’t trust AI scoring. I want to understand why a candidate ranked where they did.”

Every ESAI Score comes with explicit written reasoning — tenure pattern, company tier, role fit on specific dimensions, language signal, career trajectory observations. You see the evidence behind the ranking, not just the number. If you disagree with how a candidate was ranked, that is the kind of feedback the system is built to learn from over time.

“Will the AI scout mails embarrass me?”

Honest answer: drafts are good but not perfect. Native-quality business Japanese with proper keigo, tested constantly against bilingual recruiters and bilingual candidates. Clean business English. But you should still read every mail before sending. The platform drafts the first 80%; you finish the last 20%. It is not “send and forget.” It is “draft instead of stare at a blank cursor for 15 minutes.”

“What happens to my candidate notes and data?”

Your data stays your data. Workspace isolation at the tenant level. Never shared with other agencies. Never used to train models for anyone outside your workspace. Every search you run builds your own candidate database, exportable on demand. We are not trying to lock you in — we are trying to be the engine that builds your asset.

“Won’t AI sourcing commodify recruiting?”

The opposite. Scout mails, longlist building, and triage are the easy, low-leverage parts of recruiting. Candidate qualification, meeting strategy, client positioning, and closing are where your fee comes from — and those parts stay human. The recruiters who use AI well don’t get replaced. They get to spend more of their week on the work that actually pays.

09 THE MECHANIC

How the 10 free credits work.

End to end.

01 Sign up at headhunt.ai.

No card. No demo call. No sales contact unless you ask for one.

02 Ten free credits land in your account immediately.

1 credit = 1 qualified Japan candidate matching your role at score 50+.

03 Pick one open role on your desk.

The honest test is your hardest role — the one your team has been working for months. Mid-market and contingent roles in segments where AI scoring works well (bilingual finance, IT, sales, commercial, HR, marketing, GTM, ops, supply chain, most engineering) produce the cleanest signal.

04 Paste the JD.

The platform returns up to 1000 ranked candidates from the 4M+ Japan database in 1–2 minutes. The first ten candidates above the score-50 threshold come from your free credits.

05 Read the list. Ask the binary question.

Are any of these candidates ones I haven't already worked? If yes, the platform is finding people your current process is missing. From there, top up credits, scale to more roles, or onboard your team. If no, walk away. Cost: zero. Time: 2 minutes.

10 THE BOTTOM LINE

The honest take.

Most pitches you read about AI recruiting tools are written to make you sign before you have tested anything. This one is written the other way around.

We built Headhunt.AI on our own desk because the existing tools didn't work and we needed something that did. We have been running on it since 2018. The Q1 2026 production numbers in this briefing are from our own recruiters working real roles. The personal compensation math holds at typical Japan agency comp shares. The 10 free credits exist because the only honest way to know if any of this works on your desk is to test it on your desk.

If you sign up and the first list is full of candidates your team has already worked, you have your answer and you walk away. If it isn't, you have a different answer and the rest of the conversation gets simpler.

REMINDER

*These systems are the worst they will ever be today. The pace of improvement in AI is not linear — **invest now to stay ahead of your competition, or fall behind.***

“That is the offer.

about HEADHUNT.AI · FOR RECRUITERS

Built by an agency. For recruiters.

Headhunt.AI is the operating platform built and run by ExecutiveSearch.AI K.K. — a Tokyo software company born out of an agency, operating an AI-first model since 2018. We are not a vendor that wandered into the recruiting industry. We built the platform we needed on our own desks, ran it there for eight years, and now extend it to recruiters at peer Japan agencies who want the same lift on their own roles.

The 4M+ profile Japan database, the ESAI Score, the bilingual scout-mail engine, and the Q1 2026 production numbers in this guide are all from the same platform that runs our own desk today.

Headhunt.AI

START WITH 10 FREE CREDITS

**One JD. One ranked list.
2 minutes. No card.**

Sign up. Paste a real role. Read the list. If the candidates are new to you, you have your answer — and you got it for zero cost in two minutes.

SALES@EXECUTIVESEARCH.AI · HEADHUNT.AI · TOKYO, JAPAN