

BRIEFING 01 · AGENCY-TO-AGENCY

100,000 Yen Per Meeting.

How AI sourcing turns ¥100,000 of credits into ¥1,500,000 of expected revenue. A 15× return — and how to test it on your own desk for just ¥75,000.

15×

EXPECTED RETURN ON CREDITS

¥100,000 in → ¥1,500,000 out in
expected revenue.

inside THIS BRIEFING

What follows, in eleven sections.

A brief written for two audiences: agency principals and partners running Japan recruiting desks. Built by an agency, tested daily on our own desks since 2018, and offered to peer agencies on the same terms we use ourselves.

01 The 15× ROI math.
¥100,000 of credits · 15 qualified meetings · ¥1,500,000 expected revenue.

02 Where the fees went.
Why per-placement fees aren't your problem — and what actually is.

03 The unit economic atom.
Why ¥100,000 of expected revenue per meeting is the right anchor.

04 The bottleneck.
Why your recruiters' week, not your fees, is the binding constraint.

05 What changes.
What AI sourcing does to the math, step by step.

06 The researcher multiplier.
Why agencies with researchers see compounding lift on the same headcount.

07 Building for the future.
The candidate database that compounds as AI improves.

08 The maturity model.
Where your agency sits today, and where it sits in 18 months.

09 Common pushback.
Seven objections, answered honestly.

10 A test you can run this week.
One week, one open role, one ranked list — proof on a single desk.

11 Self-audit.
Seven questions to ask inside your agency today.

01 THE BIG NUMBER

15× ROI on AI sourcing credits.

Spend ¥100,000 on Headhunt.AI credits. Expect 15 qualified candidate meetings in return. Each meeting is worth about ¥100,000 in expected revenue. That works out to ¥1,500,000 in expected revenue from a ¥100,000 spend. A 15× return.

That is the headline. Now the proof.

We use this on our own desks every day.

Headhunt.AI is the platform we built for our own agency. ExecutiveSearch.AI K.K. has been running on it since 2018. In Q1 2026, our recruiters showed clear lifts at every stage of the funnel — comparing the same recruiters working with Headhunt.AI to their own 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.

01 CONT. · WHY THE LIFT COMPOUNDS

These are not lab numbers. They are actual production results from our recruiters in Q1 2026. We use the platform every day. The lift is real, and it shows up at every step of the recruiting process — not just at the top.

Why this matters more than just “more candidates.”

Most AI recruiting tools claim better candidate lists. The hard part is what happens after. A bigger list does not help if your reply rate stays the same and your interview pass rate stays the same. You just send more messages and waste more time.

Headhunt.AI lifts every stage. More candidates met means more pipeline. Higher reply rates mean less wasted outreach. Higher interview pass rates mean better fit. Higher offer rates mean clients trust your shortlists. Each lift on its own is small. Stacked together across the funnel, they double or triple your placement output for the same recruiter time.

The rest of this guide explains why the math works, what it does to your unit economics, and how to test it on your desk for ¥75,000.

“Each lift is small. Stacked across the funnel, they double or triple placement output — for the same recruiter time.”

02 WHERE THE FEES WENT

Per-placement fees aren't your problem.

There is a story going around about Japan recruiting agencies. It says fees are collapsing. Margins are gone. AI is the only way out. Half of that is right. The fee-collapsing part is wrong. And getting it wrong leads to the wrong fix.

Per-placement fees are actually up.

Most agencies still charge 30–35% of annual salary for placements. That has not changed. And salaries in white-collar Japan have been rising for three years. So the average fee per placement is higher today than it was. A ¥10M base salary that used to give you ¥3M now gives you closer to ¥4M. Revenue per placement is up, not down.

So if revenue per placement is up, where is the squeeze?

The total contingent fee pool is shrinking.

Three things have happened to the volume of contingent fees available to agencies:

- **RPO contracts have eaten large-client volume.** When a Fortune 500 Japan office signs a multi-year RPO with one provider, the 40–80 contingent placements per year that used to flow to multiple agencies disappear from the pool.
- **In-house TA teams have absorbed the easier roles.** Mid-size and large companies have built real internal recruiting functions, often hiring senior agency consultants to lead them.
- **Licensed firms in Japan have surged since COVID.** Per MHLW data, licensed paid placement establishments grew from 22,977 in FY2019 to 30,113 by FY2023 — over 7,000 new agencies in four years, a 31% jump. Agency bankruptcies in 2024 hit a record high (~5× the pre-COVID rate, per Tokyo Shoko Research).

02 CONT. · THE DATA PICTURE

So the picture is this: each placement pays better, but there are fewer placements going to agencies. More firms are competing for what remains. And the placements that do happen take more time per role. Per-recruiter revenue is flat or down — even though per-placement fees are higher.

BY THE NUMBERS · 2019 → TODAY

The numbers that moved are volume and competition — not fees.

PER-PLACEMENT FEE

¥3M

¥4M

PLACEMENTS / RECRUITER
/ MO

1.5–3

0.75–1.5

LICENSED JAPAN AGENCIES

22,977

30,113

Each placement pays better. There are fewer placements going to agencies. More firms compete for what remains.

That distinction matters. If fees were the problem, the answer would be “compete on value” or “move upmarket.” The problem is volume per recruiter. The answer is to fix recruiter capacity. Different problems. Different solutions.

“If fees were the problem, the answer would be value or upmarket. The problem is volume per recruiter. The answer is capacity.”

03 THE UNIT

¥100,000 of expected revenue per meeting.

If you want to understand what is happening to your agency, the right unit to think in is not placements per month and not fees per role. It is candidate meetings.

Most recruiters need about 40 qualified candidate meetings to make 1 placement. That ratio shifts a little by practice area — 35 in tighter markets, 45–50 in noisier ones — but 40 is a solid working number based on our research across our own desk and benchmarks against multiple Japan agencies.

At a ¥4M average placement fee, that ratio gives you a clean number:

THE UNIT ECONOMIC ATOM

$$\begin{array}{ccc} \text{¥4M} & \div & 40 \\ \text{AVG PLACEMENT FEE} & & \text{QUALIFIED MEETINGS TO PLACEMENT} \\ & & = \\ & & \text{¥100,000} \\ & & \text{EXPECTED REVENUE PER MEETING} \end{array}$$

Every part of your sourcing operation feeds into this number. Reply rate, sourcing quality, and conversion all roll up to one question: what does it cost to produce a qualified meeting?

The right way to evaluate any change to your sourcing operation — a tool, a process, a hire, a workflow — is to ask: **does this lower the cost of getting to a qualified candidate meeting?** If yes, it is a margin lever. If no, it is not.

“¥100,000 of expected revenue per qualified candidate meeting is the lens for everything that follows.”

04 THE BOTTLENECK

Your recruiters are stuck on sourcing.

If meetings are the unit and revenue per meeting is fixed at ¥100,000, then how many placements your recruiter makes per month comes down to one thing: how many qualified meetings they can run. Walk through what eats their week.

The meetings themselves take 3 hours each.

A typical 60-minute candidate meeting takes about 3 hours of total recruiter time. Profile review and prep before. The meeting itself. Writeup, calendar follow-up, and email after. A recruiter running 30 meetings in a month spends 90 hours on meetings alone — before any sourcing work, client work, or closing activity.

The sourcing to fill those 30 meetings takes another 60–90 hours.

To get 30 qualified meetings at typical reply rates, a recruiter needs to send roughly 400 to 500 outreach messages. Qualify replies. Schedule meetings. Handle no-shows. Another 60–90 hours of work — most of it low-leverage.

The math does not fit a 40-hour week.

90 hours of meetings + 75 hours of sourcing + client work + admin + closing easily reaches 200 hours of work per month — for under 1 placement output. Real recruiters work about 160 hours per month, not 200. That is why most desks are stuck at 0.75–1.5 placements per month — even though the fee structure could support more if the capacity were freed.

04 CONT. · THE LESSON

Where recruiter time actually goes.

Sourcing and qualification consume 60–70% of the average recruiter week. Most agency leaders will guess 40–50%. When you actually instrument it across multiple Japan desks, the number is consistently 60–70%. The gap between what people believe and what people measure is itself a finding.

Here is the lesson. If you take the sourcing and qualification work off your recruiter without cutting their meeting volume, you free up real hours. Those hours go to more meetings. More meetings, at ¥100,000 each, mean more placements at the same fee structure.

“The lever is not fees. It is hours per qualified meeting.”

A note on the alternate model. This walkthrough assumed solo recruiters. The researcher/associate model — covered in Section 06 — has a different bottleneck but reaches the same conclusion.

05 WHAT CHANGES

What AI sourcing actually does to the math.

There is a lot of confused vendor messaging about AI in recruiting. Most of it focuses on writing better outreach copy. That is among the least valuable applications. The candidate inbox saturation problem is not a copy problem. It is a targeting problem. Better copy sent to the wrong 100 candidates will not beat decent copy sent to the right 20.

Here is what AI-first sourcing actually does, step by step.

01 It scores the entire candidate universe before your recruiter sees a name.

Headhunt.AI scores 4M+ Japan-focused profiles against your role's specific criteria. Not keyword matching. Real fit on role, company tier, tenure pattern, language signal, career trajectory. Each candidate gets a 0–100 ESAI Score with a written reason. Your recruiter never sees a longlist of 1000 noisy candidates.

02 It routes only the top candidates to your recruiter.

Of a 4M-profile universe, a typical mid-market role surfaces 200 to 800 candidates worth contacting. Your recruiter only sees the top scorers. The 60–70% of recruiter time that used to go to triage disappears. Reply rates run higher because the targeting is sharper.

03 It writes personalized scout messages in business Japanese or English.

Headhunt.AI generates scout mails referencing each candidate's actual profile, current role, and visible career signals — not template merge fields. Native Japanese with proper keigo, or clean business English. The 6 hours per week your recruiter spends drafting outreach goes to zero.

04 Your recruiter's time shifts from sourcing to closing.

The capacity that was locked up in producing meetings goes back to running them. A recruiter capped at 25–30 meetings per month can now run 50 or more, because the supporting time is no longer the bottleneck. Same fee. Same conversion. More meetings.

06 THE RESEARCHER MULTIPLIER

If you have researchers, the math compounds.

Section 04 assumed a billing recruiter who does their own sourcing. That is common at smaller agencies and contingent shops. It is not common at all agencies. A meaningful share of Japan agency revenue — search firms, retained boutiques, larger contingent agencies, and shops with offshore researchers — uses a different model. Senior billing recruiters who close. Researchers who handle the front of the funnel.

If your agency runs that model, the math is even better. Two things shift at the same time. Senior recruiters stop being bottlenecked by associate availability. And one researcher can support 2–3× the billing volume — without adding headcount. Both shifts feed straight into senior billing recruiter pay.

THE RESEARCHER REALLOCATION · WHERE THE WEEK GOES

Same headcount. Same hours. Very different output.

BEFORE · MANUAL



↓ 50% → 5%

↑ 30% → 70%

AFTER · AI-NATIVE



Universe-level scoring kills the candidate identification step. The 50% of researcher time that used to go to manual list-building drops to about 5% — checking the AI’s output and giving feedback for calibration. The capacity that frees up moves into actual outreach.

The senior biller's take-home doubles.

This is where the math really matters for senior billing recruiters specifically. Most agency comp models pay senior billers a percentage of placement fees they generate — usually 35 to 50% at this level, especially at search firms and retained boutiques. Walk through 45%, a common figure.

Two effects show up at the same time. The firm sees its gross revenue per recruiter double. The senior recruiter sees their take-home double.

BOX 01 · FIRM VIEW

Gross revenue per recruiter

BEFORE · ~30 meetings/month	$0.75 \times \text{¥}4\text{M}$
Monthly gross	¥3.0M
Annualized	¥36M / yr
AFTER · ~60 meetings/month	$1.5 \times \text{¥}4\text{M}$
Monthly gross	¥6.0M
Annualized	¥72M / yr

ANNUAL DELTA · PER RECRUITER **+** **¥36M**

BOX 02 · RECRUITER VIEW

Senior biller take-home @ 45% commission

BEFORE · 0.75 placements/mo	$\text{¥}3.0\text{M} \times 45\%$
Monthly take-home	¥1.35M
Annualized	¥16.2M / yr
AFTER · 1.5 placements/mo	$\text{¥}6.0\text{M} \times 45\%$
Monthly take-home	¥2.7M
Annualized	¥32.4M / yr

ANNUAL DELTA · PER RECRUITER **+** **¥16.2M**

Firm gross doubles. Senior take-home doubles. Same fee structure, same headcount. The lever is meeting volume.

07 BUILDING FOR THE FUTURE

The asset you're really building.

The most underrated asset in any agency is its candidate database. Not the tools. Not the AI. The names, profiles, and contact context you can come back to next quarter or next year when a new role lands. Build that database well and every new search compounds. Build it badly and you start from scratch on every role.

Most agency databases are graveyards.

Three-year-old records. No current company info. No current title. No fit context. Your ATS might hold 30,000 entries, but only 500 are usable for any given current role. The rest are dead weight.

AI is moving fast. Your data is what compounds.

The AI scoring tools available in 2027 will be materially better than today. The same will be true in 2028 and 2029. Each generation produces more accurate matches. But all of those gains depend on the same input: structured, current, well-organized candidate data. The agencies that own a clean database take advantage of every AI improvement as it lands. The agencies that do not rebuild from scratch each time.

Built around export from day one.

- Push ranked candidates straight into your existing ATS — Salesforce-based systems, Bullhorn, Zoho, or equivalent.
- Export to LinkedIn Recruiter CSV ready for direct import into your Projects.

“You don't just buy meetings. You build a structured, current, AI-scored candidate database — and it's yours.”

08 THE MATURITY MODEL

Where your agency sits today.

Most Japan agencies fall on a 5-stage progression in how they produce candidate meetings. The stages are not subtle. Where you sit today determines what your meetings-per-recruiter number will look like in 18 months.

STAGE	DEFINING FEATURE	MEETINGS / MO	TRAJECTORY
0	Manual · Boolean search, no scoring	~10	Capacity-capped
1	Tool-assisted · LinkedIn Recruiter / BizReach filters	~25	Marginal speedup
2	AI-assisted · AI outreach drafting, light scoring	~30	Modest improvement
3	AI-native · Universe scoring, top-decile routing	~50	Step change
4	AI-first ops · Business model rebuilt around it	60+	New baseline

STAGE 2 → STAGE 3

Not a tooling upgrade. **A workflow rebuild.**
 "Same workflow, faster" vs "different workflow, different cost structure."

STAGE 1 → STAGE 2

Standard platform features compress some manual work. Real improvement, but bounded. The change is *tactical*.

09 COMMON PUSHBACK

Common objections we hear.

Seven questions we have heard from peer agency principals across the last twelve months. Each gets a direct answer, not a deflection.

“There are cheaper AI tools. LinkedIn has AI built in. Why pay for this?”

Fair pushback. Most modern ATSs have some AI scoring. Many are cheaper than Headhunt.AI. They improve one part of the funnel — usually the longlist. They do not improve what comes after. The lift has to compound across the funnel, not sit at one stage. Cheaper AI tools win on price; they lose on funnel compound.

“My recruiters' calendars are already full. They can't fit more meetings.”

Usually wrong, but for an interesting reason. Most agency leaders genuinely believe their senior recruiters are full. When you actually instrument the calendars, the audit consistently shows 60–70% of the week on sourcing-related work — not on meetings. That is the work AI absorbs. Once it goes, the calendar opens up significantly.

“Our recruiters won't trust AI-scored candidates.”

True at first. The fix is not training. It is transparency. Recruiters trust scoring when they can see *why* a candidate scored well — the actual evidence, not just the number. Headhunt.AI shows the reasoning behind every ESAI Score: tenure pattern, company history, role fit on specific dimensions. That is what makes scoring usable rather than just visible.

“Our clients want human-curated shortlists.”

Clients want shortlists they trust. Whether the curation is human or AI-assisted matters less to them than the hit rate. A recruiter presenting five candidates with explicit AI-derived match reasoning, plus their own commentary, is a stronger product than a recruiter-only shortlist with no scoring framework.

“AI can’t understand niche skill sets. Our specialty is too specific.”

Sometimes true. AI scoring genuinely struggles with very narrow technical specialties where expertise is invisible from a public profile — certain hardware engineering subfields, some compliance niches. For everything else — bilingual finance, mid-tier IT, sales, ops, GTM, commercial, supply chain, product, marketing, HR, legal, most engineering — AI scoring outperforms manual screening on profile-readable signals. For most agencies, 80–90% of billable role volume is in segments where the math works.

“We tried an AI tool already and it didn’t work.”

Fair pushback. The cleanest way to handle it is a head-to-head test rather than a debate. The test in Section 10 is built around exactly this kind of side-by-side comparison. If our list is not visibly better than the alternative on your roles, do not buy more credits. The test costs a single recruiter-day to find out.

“Why wouldn’t every agency just build this themselves?”

They could. Building a Japan-specific scoring model, maintaining a 4M-profile database, running outreach infrastructure, and calibrating scoring against placement outcomes is roughly an 18-to-24-month engineering project. The pitch is straightforward: license what we already built rather than rebuild it. The math works at agency scale starting on day one.

10 THE TEST

A test you can run this week.

Everything in this guide is theory until you put it in front of your own recruiters on your own roles. Here is the simplest way to do that.

THE TEST, IN ONE LINE

500 ranked, scored Japan candidates against one of your real open roles. No contract. No setup. No integration.

How it works.

01 Buy a ¥75,000 credit pack.

500 credits = up to 500 qualified candidate matches against your search criteria, scoring 50+ on the ESAI Score. No subscription. No annual commitment. Credits never expire.

02 Pick one open role and paste the JD.

Mid-market, contingent, in a segment where AI scoring works well — bilingual finance, IT, sales, commercial, HR, marketing, or similar. Headhunt.AI returns up to 1000 ranked candidates from our 4M-profile Japan database in 1–2 minutes.

03 Show the list to the recruiters who work that segment.

Ask one question: *“Are there candidates on this list you haven’t already seen through your normal sourcing?”* If yes — even a handful — Headhunt.AI is finding people your current process is missing. That is your proof of concept.

11 SELF-AUDIT

Seven questions worth asking inside your agency.

The right starting question is not “should we adopt AI.” It is a set of specific operational diagnostics. The point is not to count check marks. It is to find which of these you can answer with real numbers and which you can only answer in generalities. The latter list is where the work is.

- Do you know your **meetings-per-placement ratio, by practice area**, measured rather than estimated?
- Do you know what percentage of your recruiters' time goes to **sourcing and qualification** — measured, not guessed? Most quote 40–50%. Audits show 60–70%.
- Have you measured your **reply rate by candidate quality decile**, or only on overall blended averages?
- What is your **average time-to-first-qualified-CV, by practice area**, over the last 12 months? If this number is rising, you are losing competitive ground.
- Do you know what fraction of your **historical contingent fee pool has migrated to RPO contracts and in-house TA** over the last three years?
- Have you run a **structured AI sourcing test** in the last 12 months — or is your view based on vendor demos and conversations with peers?
- If a peer agency announced **80 qualified meetings per recruiter per month** at your fee tier tomorrow, what is your concrete response?

SCORE INTERPRETATION

6–7 specifics: Stage 3 or close. Remaining gaps are calibration, not redesign.

4–5 specifics: Stage 2 with Stage 3 ambitions. The fastest path is the one-week test on a single recruiter.

2–3 specifics: Stage 1 with some tool adoption. The bottleneck is operating model.

0–1 specifics: Structurally exposed. The desk you run today will not survive 18 months of further volume migration.

12 WHERE THIS LEADS

The honest take.

The Japan recruiting agency market is in a transition that will be obvious in retrospect. Per-placement fees are not the problem. Volume migration to RPO and in-house TA, plus recruiter capacity constraints on the volume that remains, is the problem.

The agencies that solve the capacity problem will grow out of this transition with stronger margins than they entered it with. The agencies that don't will operate against a permanently smaller fee pool with the same cost structure they have today.

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.***

“This is uncomfortable to read. It is more uncomfortable to act on. Doing nothing is a decision, the same as any other. It just looks more like the present, which makes it feel safer than it is.

about HEADHUNT.AI · FOR AGENCIES

Built by an agency. For agencies.

Headhunt.AI is the operating platform built and run by ExecutiveSearch.AI K.K. — a Tokyo recruiting agency operating an AI-first model since 2018. We are not a vendor that wandered into recruiting. We are an agency that built the platform we needed, ran it on our own desks for eight years, and are now extending it to peer Japan agencies.

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

Headhunt.AI

START WITH THE ¥75,000 TEST

500 candidate matches. One conversation. No contract.

Buy a credit pack. Run one search. Show the list. Ask one question. If the candidates on the list are new to your team, you have your answer — and you got it for less than the cost of a single recruiter-day.

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