

A HEADHUNT.AI PUBLICATION

AI Recruiting in Japan: The Handbook

Sixteen weeks of hands-off production. 17.2× return on credits. The full picture of AI-first recruiting in Japan, in one document.

17.2×

MEASURED ROI ON CREDITS · 16 WEEKS ·
ESAI DESK 2026

¥100,000 in → ¥1,720,788 out, in
expected revenue.

foreword A NOTE FROM KEN

A note from Ken.

In 2026, AI is everywhere. It is dominating the enterprise, the stock market, and people's pockets. ChatGPT was the fastest-growing app in history, and OpenAI is not even the largest AI company anymore. Google, Meta, and thousands of other companies have their own versions. Much like search ads destroyed newspapers and Netflix destroyed video rentals, AI is coming for the recruiting industry.

But unlike newspapers and video rental stores, you have a choice to stay competitive. A choice to actually win while your competitors struggle and fade away.

I founded ExecutiveSearch.AI over 8 years ago (back when getting a .AI domain meant faxing a form to the government of Anguilla, who owns it, and *universal translators* only existed in Star Trek). Our mission from day one was to apply AI to recruiting.

In my position, I am fortunate I get to see dozens of companies embracing AI. Seeing their Eureka! moments and the real day-to-day struggles with adoption, cost, compliance, and technical issues.

I did my 10,000 hours so the next guy only needs to do 100. That 100 starts here. In this handbook.

Starting with Headhunt.AI in 2026 is like a cheat code. It's built on top of algorithms and best practices that we painstakingly tested and refined over 8 long years.

I'm here to help recruiters in Japan reshape their work with AI.

You got this!

Ken Charles

CEO · EXECUTIVESEARCH.AI K.K.

P.S. Shoot me an email with any question you have at all. I'm here to help: ken@executivesearch.ai

inside THIS HANDBOOK

What follows, in six parts.

PART I **The state of Japan recruiting in 2026.***I* *The market reality. Why the existing operating model is breaking.***I.1** **The recruiter week, in thirty-minute blocks.***Where the hours actually go when you measure them.***I.2** **The decision that got harder, while hiring did not.***The 49% → 33% advance-rate drop, and why it matters.***I.3** **The cost layers stacked on top.***Per-placement-fee databases · seat-based platforms · named, not yet decomposed.*

PART II **What AI-first sourcing actually does.***II* *The mechanism. Not the marketing-copy version — the operating-model version.***II.1** **The unit of work is one qualified meeting.***¥107,676 per meeting · 39.625 meetings per placement · the unit economic atom.***II.2** **The funnel, end to end, and where the lifts compound.***+38% / +13.8% / +13.5% / +14% — and why the stack matters more than any one number.***II.3** **What aggregation, scoring, and scout mail each contribute.***Three distinct steps. Three separate lifts.***II.4** **The candidate database is the side effect, not the product.***Why every search reverses the 30% annual decay of your existing records.***II.5** **Two candidate pools, one ranked list.***Public-data pool plus Enterprise ATS pass-through, scored on the same axis.*

PART III **The production evidence.***III* *Where projection meets receipts. Comes early because evidence precedes evaluation.***III.1** **Sixteen weeks of hands-off operation.***123,675 candidates · 1260 meetings · 100% AI-sourced · zero human review.***III.2** **Daily replies, up 78%.***526 days of inbound logs, 10,932 replies, one variable changed.***III.3** **The cleanest A/B we can run on a working desk.***Same recruiters, same market, same engine. Different lists.*

inside THIS HANDBOOK · CONT.

PART **The math for your operation.****IV** *Persona-specific economics. Recruiter, principal, corporate budget-holder, worked scenarios.***IV.1** **If you hold the desk.**

What the lifts mean for take-home, by comp-share tier.

IV.2 **If you run the firm.**

The per-placement-fee database math · margin per placement, margin per year.

IV.3 **If you hold the budget.**

The LinkedIn seat math · what the renewal is actually buying, and the redirect.

IV.4 **Six versions of the same arithmetic.**

A worked scenario grid for the reader to apply to their own desk.

PART **The compliance frame.****V** *The chapter most foreign AI sourcing tools cannot pass.***V.1** **Two laws, one filing category.**

APPI and the 改正職業安定法, with the article numbers.

V.2 **Six of 1642 — the structural rarity.**

Why the 第4号 filing category is the structural-defensibility signal.

V.3 **Ten questions worth asking any vendor.**

A vendor-neutral self-audit the reader can use against anyone.

PART **The path forward.****VI** *Maturity model. Common objections. The test.***VI.1** **Where your operation sits, and where it's going.**

Stages one through four, and what the next 18 months look like.

VI.2 **The objections we hear most often.**

Surfaced, conceded where warranted, answered.

VI.3 **The test you can run this week.**

100 free credits. One JD. Two minutes. No card.

FRONT & CLOSING Foreword · Closing · AboutAPPENDICES A Methodology · B Glossary · C
Reading order

All 2026 numbers in this Handbook reflect ESAI Agency results unless otherwise noted. Recruiters, market segments, fees, and client mix held constant; the variable is manual sourcing vs. Headhunt.AI.

I.1 THE RECRUITER WEEK

The recruiter week, in thirty-minute blocks.

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. It is wrong. Audit a real calendar in thirty-minute blocks across five working days and the figure lands consistently between sixty and seventy percent.

That gap — between what people believe their week looks like and what their week actually is — is itself a finding. It explains why most agency leaders cannot reconcile their headcount, their fee structure, and their placement output. They are running the math on a calendar they do not have.

What the audit measures.

A Boolean search on LinkedIn for a typical mid-market role runs around 90 minutes. The output is a list of roughly 1200 names that need to be re-screened against unstated context the JD does not capture. A personalized scout mail takes 15 minutes to draft. The typical manual reply rate is six to nine percent. So 100 carefully written outreach messages return six to nine conversations. To produce 30 qualified meetings in a month, a recruiter sends 400–500 outreach messages, qualifies the replies, schedules around the no-shows, and writes up each conversation.

Add three hours of total recruiter time per meeting — prep, the conversation, write-up, follow-up — and the week stops being abstract. Sixty to seventy percent of it is sourcing and qualification work that does not differentiate the recruiter from the next recruiter holding the next desk.

I.1 CONT. · THE CONSTRAINT

This is the binding constraint. Not fees. Not market conditions. Not candidate quality. Recruiter hours per qualified meeting — and the share of those hours that any human recruiter must spend on work that is, by the time it reaches them, generic.

THE WEEK, AUDITED · PRODUCTION DATA · ESAI DESK

Where the recruiter week actually goes.

SELF-REPORTED

40–50%

MEASURED

60–70%

RECRUITER WEEK ON
SOURCING &
QUALIFICATION

LINKEDIN ALONE

26 hrs

RECRUITER TIME-TAX PER
WEEK, LINKEDIN-ONLY
WORKFLOWS

MANUAL REPLY RATE

6–9%

TYPICAL RESPONSE TO
CAREFULLY WRITTEN SCOUT
MAIL

Sources: ESAI desk audit; B01 §04 (recruiter-week share); B04 §02 (LinkedIn time-tax decomposition); B02 §02 (manual reply baselines).

Every section that follows returns to this number. A change to a sourcing tool, a workflow, a hire, a process — the right question is not whether the tool is novel or well-marketed or AI-enabled. The right question is whether it lowers the cost in recruiter hours of producing one qualified meeting. If it does, it is a margin lever. If it does not, it is theatre.

“The constraint is not market conditions. It is hours per qualified meeting.”

I.2 THE DECISION GAP

The decision that got harder, while hiring did not.

Mid-2025 marked a structural change in Japan corporate hiring. The decision to hire did not slow. The decision to advance a candidate from second interview to final did. Across twenty-five months of ESAI data — 3852 resumes sent, 385 second interviews, 165 finals, 74 placements — the second-to-final advance rate dropped from 49% to 33%.

Mann-Kendall non-parametric trend test on the rolling six-month conversion: $z = -2.42$, $p = 0.015$. The decline is real, not noise. It also did not affect every stage equally. Resume-to-first-interview is roughly flat. First-to-second is slightly down. Final-to-offer-accepted is up, from 40% to 58%. Companies that get a candidate to final close harder than they used to. Fewer get there.

The shape of the gap.

The aggregate effect on hiring volume is straightforward. From the same starting volume of resumes sent, 31% fewer placements emerge on the back end. Two point one offers accepted per hundred resumes has become one point five. The same recruiters with the same clients running the same processes are producing one fewer placement per fifty resumes — not because their work has degraded, but because the committee gating the second-to-final advance now requires more evidence than it used to.

I.2 CONT. · THE MECHANISM

CONVERSION TABLE · 25 MONTHS · ESAI · BEFORE/AFTER H2 2025

Where the funnel moved, and where it didn't.

RESUME → 1ST	1ST → 2ND	2ND → FINAL	FINAL → OFFER
29%	38%	49%	40%
24%	32%	33%	58%
-5 PTS	-6 PTS	-16 PTS · THE GAP	+18 PTS

Mann-Kendall non-parametric trend test on rolling 6-month 2nd → Final conversion: $z = -2.42$, $p = 0.015$. See Appendix A for the methodology.

The mechanism is not budget. Budget is in place. The mechanism is not headcount approval. Approvals are happening. The mechanism is the evidence a committee requires to commit. When the room is asked to advance a candidate to final, the question shifts from *do we want to hire* to *do we want to hire this person, against all the others we could be looking at instead*. The committee answers that question by demanding more — more comparable profiles, more written reasoning, more documented benchmark.

That demand is structural. It does not soften when the market does. The recruiters who keep up are the ones whose process produces a candidate file that already answers the committee's question — before the committee has to ask it. This is what universe-level scoring is for, and we will return to it in Part II.3.

“The decision gap is not about desire. It is about the evidence a committee requires to commit.”

I.3 THE COST LAYERS

Two layers stacked on top of the hours.

The hours diagnosis in Part I.1 is the binding cost on the recruiter's week. Layered on top of those hours, most Japan recruiting operations carry two more cost layers in the sourcing infrastructure — recurring fees for tools that automate parts of the work the hours were meant to do. This section names the layers so the persona-specific math in Part IV can refer to them by category.

The per-placement-fee database layer.

Per-placement-fee databases — BizReach, Recruit Direct Scout, doda X, AMBI, and similar agency-side scout platforms — operate on a three-layer cost structure most agency P&Ls fail to surface. A subscription that runs roughly ¥1.2M a year. A per-placement usage fee of ¥700K to ¥1.05M — typically 20% to 30% of the recruiting fee — due on every placement the database touched. And a tracked-contact contingent exposure that lasts well past the contractual term. Annual category spend on a mid-size agency desk runs ¥8M to ¥22M and above. Part IV.2 walks the firm-view math for this layer in detail, including the public-filing precedent.

The seat-based platform layer.

Seat-based recruiting platforms — LinkedIn Recruiter (RPS and Recruiter Corporate) is the canonical example in Japan — price on per-seat subscriptions rather than per-placement fees. A 6–10 user RPS seat lists at roughly ¥687,000 per year and ships a fixed InMail allotment plus access to the platform's search environment and candidate brand. Total annual category spend scales with seat count and tier. Part IV.3 walks the corporate-buyer math for this layer in detail, with the InMail utilization and seat-environment breakdown.

“The tools sit on top of the hours, not in place of them.”

II.1 THE UNIT OF WORK

The unit of work is one qualified meeting.

If Part I is the diagnosis, Part II is the mechanism. The right place to start is the unit of analysis. Not placements. Not resumes sent. Not InMails. The qualified meeting — the conversation a recruiter has with a candidate after both sides have decided the role is worth thirty minutes — is the unit that everything compounds against.

In our own 2026 production, the average placement fee was ¥4,266,675. The average number of qualified meetings required to produce one placement was 39.625. Divide and the unit value falls out cleanly: ¥107,676 of expected revenue per qualified meeting.

Why the meeting is the lens.

Once meetings are the unit, every other recruiting question becomes a question about meeting unit cost. Reply rate is a question about meetings per hundred outreaches. Interview pass rate is a question about meetings per placement. Senior biller take-home is a question about meetings per month per recruiter. Margin per desk is a question about meetings per desk per year. The whole operation, looked at in this unit, becomes legible.

II.1 CONT. · THE ARITHMETIC

The arithmetic to the 17.2× figure on the cover follows from the unit. ¥100,000 of credits — at production unit cost — buys about sixteen qualified meetings. Sixteen meetings at ¥107,676 of expected revenue per meeting is ¥1,720,788 of expected revenue. The 17.2× multiple is the same statement, divided. The number is not a projection. It is sixteen weeks of measured production, with same recruiters, same market, same client mix, same fee structure as prior periods. The variable that changed was the source of the list.

THE UNIT · ESAI DESK · 16 WEEKS · Q1 + APRIL 2026

One meeting, decomposed.

17.2×

MEASURED ROI ON CREDITS, 16 WEEKS OF PRODUCTION

~16

QUALIFIED MEETINGS PER ¥100,000 OF CREDITS

¥1,720,788

EXPECTED REVENUE PER ¥100,000 OF CREDITS

¥107,676

EXPECTED REVENUE PER QUALIFIED MEETING

¥4,266,675 average placement fee ÷ 39.625 qualified meetings per placement = ¥107,676 per meeting. Source: B08 §04.

There is a temptation, when looking at this row, to treat the 17.2× number as the headline and the rest as supporting detail. The opposite is true. The ¥107,676 figure is the load-bearing number. The 17.2× is what the load-bearing number returns when you multiply by sixteen and divide by a hundred thousand. Change the unit cost of a qualified meeting and every number on this page moves with it.

“The meeting is the unit. Everything else is in service of the meeting, or in the way of it.”

II.2 THE FUNNEL

The funnel, end to end, and where the lifts compound.

On the ESAI desk in Q1 2026, four numbers moved. Candidate meetings per recruiter went up 38%. Scout reply rate went up 13.8%. Interview pass rate went up 13.5%. Offer rate went up 14%. Each lift is small. The point is not any one of them. The point is that they sit on different multipliers across the funnel, and the multipliers compound.

What each lift means in isolation.

More candidates met means more pipeline at the top. A higher reply rate means less wasted outreach for the same recruiter time. A higher interview pass rate means the shortlist is tighter — clients trust it more, schedule more aggressively, push fewer candidates back. A higher offer rate means the work the recruiter put into qualifying the candidate is being rewarded by the client process. Each one is a clean operational improvement.

Read separately, they are unremarkable. Most AI sourcing vendors can point to a single-stage lift somewhere in the funnel. The interesting claim is what happens when the lifts stack.

II.2 CONT. · HOW THEY COMPOUND

FUNNEL LIFTS · PER RECRUITER · Q1 2026 · ESAI DESK

Same recruiters. Same market. Same clients. Same fees.

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

MORE CANDIDATE
MEETINGS PER
RECRUITER

SCOUT REPLY RATE

INTERVIEW PASS RATE

OFFER RATE

Comparison: same recruiters working with Headhunt.AI versus their own previous quarters on manual sourcing. Source: B01 §01, B02 §04, B08 §03.

The stack does the work. A 38% lift on meetings, multiplied by a 13.8% lift on reply rate, multiplied by a 13.5% lift on interview pass rate, multiplied by a 14% lift on offer rate — each on its own a single-digit-plus number — produces a doubled-to-tripled placement output for the same recruiter time. That is the mechanism behind the 17.2× return on the cover. None of the individual lifts is dramatic. The structural rearrangement is.

Across the same sixteen-week window, the absolute pipeline numbers tell the same story at scale. 123,675 candidates contacted, 3868 replies, 1260 qualified meetings, all of it 100% Headhunt.AI sourced. We will return to those numbers in Part III.1.

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

II.3 AGGREGATION, SCORING, SCOUT MAIL

What aggregation, scoring, and scout mail each contribute.

Headhunt.AI does three things in sequence. It aggregates a Japan-focused candidate universe of 4M+ profiles from public sources. It scores that entire universe against the role — not a pre-filtered shortlist — with an ESAI Score from zero to one hundred and explicit written reasoning per candidate. And it drafts a personalized scout mail per candidate, in business Japanese or English.

Each step has a separate measurable lift. Reading them as one product is a category error.

01 Aggregation: the universe, not a database.

The 4M+ Japan-focused profiles is a description of the population the scoring runs against. The data is public; the structure and currency of the index against that data is the asset. What it gives you that a per-placement-fee database does not is universe-level coverage at zero marginal cost per record surfaced. The per-placement layer disappears because the population is not gated behind a usage fee.

02 Scoring: the entire universe, before your recruiter sees a name.

Each candidate gets a 0–100 ESAI Score against the specific role, with a written paragraph explaining the score — tenure pattern, company tier, language signal, career trajectory, the explicit dimensions the score weighs. The contribution this makes over Boolean search is not “better matches.” It is that the recruiter never sees the longlist. The triage step that consumed half the recruiter week in Part I.1 stops happening.

03 Scout mail: personalized, drafted, ready to send.

A business-Japanese or business-English scout mail per candidate, referencing the candidate's actual current role and career signals — not template merge fields. The recruiter's fifteen-minute-per-message drafting load goes to zero, and the candidate reads a message that did not come from a list.

II.3 CONT. · THE STACK

The order matters. Most AI recruiting products start from a smaller, cheaper version of step three — scout mail generation — and bolt on lighter versions of steps one and two. The result is a tool that writes better outreach to roughly the same wrong people. The funnel lift sits at one stage and stops.

The Handbook's claim is the inverse. Step one — universe-level aggregation — and step two — universe-level scoring with written reasoning — are where the structural lift comes from. Step three is what makes the structural lift survive contact with the recruiter's calendar. Drop any one of the three and the math compresses.

What “universe-level scoring” actually means.

Not: a 50-candidate longlist scored 1–10. Universe-level scoring means the 4M+ profile population is evaluated against the role's specific criteria, the bottom 99% never surfaces, and the recruiter sees only the ranked top. The standing question — are there candidates worth contacting that our normal sourcing missed — is answered by construction. The answer is in the file before the recruiter has to ask.

This is also the layer that produces the evidence the committee in Part I.2 requires. A candidate file that opens with a written ESAI Score reasoning paragraph already addresses, in writing, the question the committee is going to ask. The decision gap closes from both ends — more candidates make it to second interview because the longlist itself is sharper, and more candidates survive second-to-final because the file documents why they should.

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

II.4 THE SIDE EFFECT

The candidate database is the side effect, not the product.

What you buy is qualified meetings. What you build, as a side effect, is a structured, current, AI-scored candidate database — and it is yours.

Annual ATS decay runs around 30%. People change jobs every two to three years; fifteen to thirty percent of any agency's existing records become stale per year through normal career movement alone. Most agencies absorb this as a fact of the business and write off the older half of their database as a graveyard. Headhunt.AI reverses the math without changing the recruiter's workflow: every search refreshes a slice of the database with current employment, current scoring, and current context. The decay does not stop. It is offset, per role, by the work that was already being done to fill the role.

DECAY REVERSED · ATS RECORDS, BY YEAR

Side effect compounds. The asset is real.

30%

ANNUAL ATS DECAY RATE,
UNTOUCHED RECORDS

46%

HIRES FROM COMPOUNDING
DATABASES · TODAY

26%

SAME METRIC · 2021
BASELINE

Job-change cadence: 2–3 years. Annual record decay range: 15–30%. Source: B06 §02, §06.

The compound that matters is not the decay reversal in any one quarter. It is what the asset looks like after eighteen months of normal-volume searches against it. The structured, current, AI-scored database is, by then, the operating moat — the thing a competitor cannot replicate by buying a subscription. The product is the meetings. The asset is the byproduct of the meetings. The asset is yours.

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

II.5 TWO POOLS, ONE LIST

Two candidate pools, one ranked list.

Headhunt.AI scores two pools and ranks them as one list. The first pool is the public-data universe — the 4M+ Japan-focused profiles aggregated from public sources, scored against the role. The second pool, for Enterprise customers only, is an ATS pass-through: the customer's own existing ATS records, scored against the same criteria, on the same axis.

The Enterprise pass-through is a database-level integration engagement, not a self-serve CSV upload. The customer's records sit inside their own system; the integration ingests, scores, and returns ranked output without the records leaving the customer's control surface. The framing matters because the alternative reading — *upload your data, we'll score it* — would carry compliance implications the actual architecture does not.

Why one ranking is the point.

A typical Enterprise hiring team holds two parallel longlist sources: the candidates the AI sourcing tool finds today, and the candidates already in the ATS from previous searches. Most workflows treat these as separate streams and rank them separately. The output is two lists the recruiter has to merge in their head. Headhunt.AI ranks them on the same axis — same ESAI Score, same written reasoning, same threshold — so the recruiter sees one ordered list with no source distinction except a small badge per row. The candidate already in the ATS from eighteen months ago is evaluated against the same role on the same terms as the candidate who has never been contacted before.

II.5 CONT. · THE OUTPUT

THE RANKED-LIST OUTPUT · PER CANDIDATE · BILINGUAL

What every row carries.

PER-CANDIDATE INTELLIGENCE

- Career history
- Education history
- ESAI Score (0–100) with written reasoning
- Fit summary paragraph
- Source-pool badge (public · ATS pass-through)

FIVE SOCIAL PROFILE LINKS

- LinkedIn
- X
- GitHub
- Facebook
- Instagram

PLUS

- Bilingual draft scout mail, ready to send

CSV export of the ranked list is universal and self-serve. ATS pass-through ingestion is Enterprise database-level integration only.

Two consequences follow. First, the recruiter spends zero time wondering whether the candidate they are about to contact has already been in conversation with the firm — the unified ranking surfaces it, and the per-candidate file shows the prior history. Second, the structural decay covered in Part II.4 reverses against the ATS pass-through pool the same way it reverses against the public pool: every search refreshes employment and scoring on whatever slice of the existing records overlaps with the role. The asset compounds across both pools.

“Two pools, one ranking. Your existing records score on the same axis as the universe.”

III.1 SIXTEEN WEEKS, HANDS-OFF

Sixteen weeks of hands-off operation.

January 1 to April 30, 2026. Sixteen consecutive weeks. One agency desk — ESAI's own — carrying its normal client load through Q1 and the first month of Q2. Across that window, 100% of candidates contacted were sourced by Headhunt.AI. 100% of scout mails sent were authored by Headhunt.AI. Zero human review of either, beyond the operational hand-offs every desk runs anyway.

The recruiters did not source. The recruiters did not draft. The recruiters did what the lifts in Part II.2 free them to do — they qualified replies, ran meetings, briefed clients, and closed. The sourcing layer ran underneath the desk the way the email server runs underneath the inbox.

What “hands-off” actually means here.

It does not mean “set it and forget it.” It means the recruiter's role moved up the stack. The desk still has a Candidate Manager team, still has Business Managers, still talks to clients daily. What the desk does not do anymore is the work that occupied 60–70% of the week in Part I.1 — Boolean searches, manual scoring, scout-mail drafting from scratch. That work is the work that produces no per-recruiter differentiation. It is the work the sourcing layer can do without the recruiter in the loop.

The receipts that follow are sixteen weeks of what happens to a working desk when that load is lifted off it.

III.1 CONT. · THE FUNNEL, END TO END

THE FUNNEL · 16 WEEKS · JAN 1 → APR 30, 2026 · ESAI DESK

123,675 contacted. 3868 replies. 1260 qualified meetings.**123,675**

CANDIDATES CONTACTED

3868

REPLIES

1260

QUALIFIED MEETINGS

100% of candidates sourced by Headhunt.AI. 100% of scout mails authored by Headhunt.AI. Zero human review of either. Meetings include both first-time candidates and ATS re-surfaces (Part II.5). Source: B08 §03.

Read the funnel left to right and the conversion rates fall out. 123,675 to 3868 is a 3.13% reply rate. The 6–9% manual baseline from Part I.1 is higher on a per-message basis, but every one of those manual messages took fifteen minutes of recruiter time to draft. Here the recruiter wrote zero of them. The comparison axis is not per-message — it is per-recruiter-hour. 3868 to 1260 is a 32.57% reply-to-meeting rate. That second number is the one that tends to surprise readers — a third of every reply turned into a scheduled qualified meeting on a working agency desk over four months.

The combined ratio is what compounds. 1260 qualified meetings against 123,675 outreach attempts is roughly one meeting per 98 messages — across a window long enough that ramp effects, seasonality, and the normal scatter of a working desk are all in the dataset. Not a sprint. Not a launch quarter. Q1 plus April, on the working desk.

Why the absolute numbers matter as much as the rates.

Rate improvements are easy to ship and easy to argue about. Absolute volume is harder to fake. The desk did not run 123,675 outreach attempts in sixteen weeks by being clever about the cohort. It ran them because the sourcing layer aggregates a 4M+ Japan-focused universe and scores it against the role before the recruiter sees a name — the contribution chain we set out in Part II.3. The cost per outreach attempt is set by the platform; the cost per qualified meeting is set by the funnel that follows.

III.1 CONT. · PROJECTED VS MEASURED

The credit math closes the section. At the unit-economics level: ¥100,000 of credits produces about 16 qualified meetings; at the desk’s 2026 average placement fee, those meetings carry ¥1,720,788 in expected revenue. That is the 17.2× number on the cover — measured across the sixteen-week window, with the placement-fee average at ¥4,266,675, the meetings-per-placement ratio at 39.625, and the per-meeting expected revenue at ¥107,676. The arithmetic is in Part II.1; the receipts behind it are the funnel above.

PROJECTED VS MEASURED · THE HEADLINE ROI ON CREDITS

What we said it would be. What it actually was.

15×

PROJECTED RETURN ON CREDITS · B01, 2025

17.2×

MEASURED RETURN ON CREDITS · Q1 + APRIL 2026

¥1,720,788

EXPECTED REVENUE PER ¥100,000 OF CREDITS

Projected figure is from B01 (the original anchor). Measured figure is from the sixteen-week production window. Per-meeting unit economics: ¥107,676 expected revenue. Source: B08 §01, §04, §05.

Two things are worth saying about the gap between 15× and 17.2×. First, the projection was honest — it was built from the funnel arithmetic of the desk as it then ran, with the lifts modeled at the low end of the range we have since measured. Second, the measured number is not the asymptote. It is sixteen weeks. The desk continues to run. Part III.2 looks at a separate 526-day reply dataset that says the structural lift is still compounding, not flattening.

“We projected 15×. The receipts say 17.2×. Same desk, same recruiters — different sourcing.”

III.2 DAILY REPLIES, UP 78%

Daily replies, up 78%.

A second dataset, separate from the credits funnel in Part III.1. 526 consecutive days of inbound reply logs on the ESAI desk. 10,932 total replies. One structural break: November 1, 2025, the day list-building moved from human recruiters to Headhunt.AI. Everything else — the recruiter team size, the scout-mail engine (Headhunt.AI's, in production since August 1, 2024), the outbound sending setup, the client mix — stayed the same across both periods.

The pre-period runs 345 days, from November 21, 2024 to October 31, 2025. The post-period runs 181 days, from November 1, 2025 through April 30, 2026. Daily replies rose from 16.4 to 29.1 — a +78% lift. The post-period also includes a weekend-send pause introduced in Q1 2026, which suppresses the daily average; on the cleanest like-for-like comparison (weekdays only), the lift widens to +97% (18.2 → 35.8). Net additional replies across the post-period: +4652 over what the pre-period average would have predicted.

Why this is a separate dataset, and why that matters.

The 17.2× in Part III.1 is a credits-to-revenue calculation. This is a list-quality measurement on a longer timeline. The two datasets do not double-count — they look at the same desk through different windows, with different metrics. If only one piece of evidence stood up, the case would be weaker. Two independent lines of evidence, on the same desk, pointing the same direction.


“One variable changed. Daily replies up 78%.
Weekday replies up 97%. The receipts run 526 days.”

III.2 CONT. · THE MONTHLY PICTURE

MONTHLY INBOUND REPLIES · 17 MONTHS, ONE DATASET, ONE STRUCTURAL BREAK


The shift is visible at the monthly level.

PRE-PERIOD · HUMAN-BUILT LISTS 345 DAYS · AVG 16.4 DAILY



DEPLOYED · NOV 2025

POST-PERIOD · HEADHUNT.AI LISTS 181 DAYS · AVG 29.1 DAILY



* Nov '24 covers 10 days (start of granular logging). Apr '26 includes a Golden Week send pause; pace-adjusted run-rate is about 1150. Source: B09 §04.

Consistency, not just the headline.

Averages can be carried by a single record month; distribution is harder to fake. 40+-reply days happened on 11 of the 345 pre-period days (3.2%). On the 181 post-period days, they happened on 56 (30.9%) — a 9.6x increase in production-grade days as a share of the calendar. The tail itself moved.

III.3 THE CLEANEST A/B

The cleanest A/B we can run on a working desk.

Working agencies cannot run randomized controlled trials. We do not have a control desk in a sealed room. We do not get to randomize role assignments to recruiters. What we can run, and what the dataset in Part III.2 is, is a same-everything-else comparison. One variable changed; the rest of the desk held.

That distinction matters. The receipts here are not a proof in the laboratory sense. They are a structural comparison on a working desk with strong design and strong execution — the kind of evidence a senior operator would weigh seriously when making a sourcing decision, and the kind a careful methodologist would caveat before generalizing.

What stayed the same across the 526-day window

Same recruiter team size, headcount held constant. Same scout-mail engine — Headhunt.AI's AI generator, adopted August 1, 2024, day one of the measurement window. Same outbound sending setup, in the same configuration. Same client mix and market: mid-market Japan recruiting across bilingual finance, IT, sales, commercial, HR, and marketing practice areas. The only structural change introduced on November 1, 2025 was the source of the target list — human recruiters and researchers built it by hand before; Headhunt.AI generated it after. The scout-mail engine wrote the messages either way; only the list changed.

Read the callout above slowly. Every line is a confound the design closes. The cleanest comparison a working agency can run is one where the desk operates as it always has on every dimension except the one being tested. That is what this dataset is.

III.3 CONT. · WHAT THE VARIABLE WAS

The one variable.

Before November 1, 2025, recruiters and researchers built target lists by hand — Boolean searches on LinkedIn, screens against a database subscription, manual triage of the longlist, prioritization against the JD. After November 1, 2025, the target list was the ranked output of Headhunt.AI: the 4M+ Japan-focused universe scored against the role, returned ranked with an ESAI Score per row. The scout-mail engine that wrote messages to those lists had been Headhunt.AI's since August 1, 2024 — day one of the measurement window. The change in November 2025 was not the messages. It was who the messages were going to.

Everything downstream of the list — who sends the message, when they send it, what they say in the follow-up, how they qualify the reply, how they prep the meeting — ran on the same desk it has run on since August 2024. That is the comparison. List source is the lever. The funnel that the list flows into is held constant.

What this is and is not.

This is the strongest comparison a working desk can produce. It is not lab-grade causal inference, and we say so. Appendix A walks through the methodology — dataset construction, time windowing, weekday filtering, what the +97% weekday number can and cannot be inferred from. Readers who want the statistical detail go there. Readers who want the operational read leave with this: when the entire 4M+ universe is being scored against your role, “better candidates” is not a marketing claim. It is a math claim about the size of the pool you are drawing from — and the size of the pool is what changed on November 1, 2025.

“This is not a lab. It is a working desk. Same recruiters, same market — one variable changed, and the variable was the list.”

IV.1 THE RECRUITER'S MATH

The working recruiter's take-home math.

If you hold a desk — if you are the recruiter whose name is on the placement — here is what the lifts in Part II.2 are worth to you, in your own pocket, on a normal year. The math is straightforward; the numbers are illustrative for the comp-share calculation; the underlying lifts are the same Q1 2026 figures the rest of the Handbook leans on.

Start from a baseline. A working biller on a Japan agency desk runs at roughly 0.75 placements per month — nine placements per year, give or take, on a focused role mix. The average placement fee runs around ¥4M (rounded — the production-data figure in Part III.1 is ¥4,266,675; we round here because the comp-tier math is illustrative, not a receipt). Comp share of the fee runs from 25% for junior-to-mid-level recruiters, through 35% for mid-to-senior billers, to 45% for senior billers and partners.

The lift, applied to your desk.

Compound the four Q1 2026 lifts — +38% on meetings, +13.8% on scout reply rate, +13.5% on interview pass rate, +14% on offer rate — and the math against your baseline produces about +0.5 additional placements per month. Six additional placements per year. On the same headcount, the same hours per week, the same client desk. The hours that were spent sourcing are now spent closing and pre-qualifying — which is where placements come from.

Six additional placements at ¥4M of fee is ¥24M of additional firm revenue passing through your name. Your share of that, depending on your tier, is where the take-home math lands.

IV.1 CONT. · THE TAKE-HOME TABLE

ADDITIONAL ANNUAL TAKE-HOME · +6 PLACEMENTS/YEAR · ¥4M AVERAGE FEE

Same desk, same fees, different sourcing.

JUNIOR TO MID-LEVEL

MID TO SENIOR BILLER

SENIOR BILLER / PARTNER

25% of fee

35% of fee

45% of fee

+¥6M

+¥8.4M

+¥10.8M

ADDITIONAL ANNUAL TAKE-HOME, PER YEAR

ADDITIONAL ANNUAL TAKE-HOME, PER YEAR

ADDITIONAL ANNUAL TAKE-HOME, PER YEAR

Baseline assumption: 0.75 placements/month, +0.5/month lift compounded across the funnel, ¥4M average placement fee (rounded). Source: BO2 §05. Production-data placement-fee average is ¥4,266,675 (Part III.1).

Read the right column carefully. Six to ten million yen, in your name, not the firm's — on the same fee structure, the same client relationships, the same forty hours the law says you are working anyway. The lift is not a different desk. It is the same desk with the sourcing hours redirected. That is the entire argument.

The number that does not appear in this table is the one that finishes the case. Senior billers leave agencies when their hours-out-to-placements-in ratio stops scaling. The sourcing tax is the binding constraint on personal scaling, not market conditions and not fee compression. Move the tax — you move the ratio — you move what a senior career on a desk is worth.

“Six to ten million yen per year, in your name — not the firm's. Same desk, same fees, redirected hours.

IV.2 THE PRINCIPAL'S MATH

The agency principal's per-firm math.

If you run the firm — if your name is on the company's lease and the company's P&L — here is what the same lifts mean across your bench. The recruiter math in Part IV.1 lands six to ten million yen per biller. The firm math is what happens when you stack twelve, twenty, thirty of those desks against the same sourcing architecture, and then redeploy the per-placement-fee database layer named in Part I.3 on top.

Capacity per recruiter moves from roughly 25–30 qualified meetings per month at the pre-Headhunt.AI baseline to 50+ on the same headcount. That doubling of meeting throughput, compounded through the funnel lifts in Part II.2, is what produces the per-recruiter placement uplift in Part IV.1. At the firm level, it produces two distinct revenue effects — more placements on the same bench, and recovered margin on each one.

The database layer, decomposed.

Per-placement-fee databases — BizReach, Recruit Direct Scout, doda X, AMBI, and similar agency-side scout platforms — carry the three-layer cost named in Part I.3: a subscription of roughly ¥1.2M per year, a per-placement usage fee of ¥700K to ¥1.05M at 20–30% of the recruiting fee, and a tracked-contact contingent exposure that can claim against placements made long after the firm exits the subscription. The total annual category spend on a mid-size 12-biller desk runs ¥8M to ¥22M, with the per-placement layer carrying the largest share by spend and the contingent layer carrying the largest share by surprise.

IV.2 CONT. · THE DISCLOSURE, AND THE REDIRECT

What the public filings show.

For Startups, Inc. (TSE: 7089) disclosed in its January 20, 2023 TDnet filing that it had restated cost of sales from FY3/2018 onward by ¥402M — ¥118M of unpaid principal plus ¥283M in contractual damages — relating to a per-placement-fee database dispute. That is roughly a 2.4× penalty multiple on the underlying principal. The disclosure is the documented cost of getting the per-placement layer wrong on a publicly traded Japanese recruiting firm. It is in the filings; it is not theoretical.

The unit comparison, per placement.

Walk through a single ¥10M-base placement at a 35% recruiting fee — the B05 reference case. The client pays the agency ¥3.5M. Under the per-placement database model, the agency pays the database vendor 25% of that fee: ¥875,000. Run the same placement through Headhunt.AI credit consumption and that cost lands at around ¥267,000. The delta is roughly ¥608,000 of recovered margin per placement, applied to every placement the firm closes. Not a lift on the top line. A redeployment from the database-vendor P&L to yours.

The same ¥875,000, redeployed once into AI sourcing credits at production unit cost, surfaces over 13,000 qualified candidates against a live role at a 25% fee structure — closer to 16,000 at 30%. The fee buys one record. The redirect buys the universe.

IV.2 CONT. · THE FIRM VIEW

A 12-BILLER DESK · MARGIN RECOVERY, PER YEAR

¥608K per placement. ~¥55M per year.**¥608K**RECOVERED MARGIN PER
PLACEMENT**~90**PLACEMENTS PER YEAR ·
12-BILLER DESK**~¥55M**RECOVERED MARGIN PER
YEAR, THIS LEVER ALONE

¥608K × ~90 placements/year ≈ ¥55M of margin recovery, on top of any revenue lift. The 90-placement figure is B05's anchor for a working 12-biller desk. Source: B05 §05.

Where the saved hours go on a firm view.

On a per-recruiter basis, the sourcing reallocation in Part II.3 frees hours that were doing list-building work that any sourcing tool can do. On a firm view, that reallocation runs through the researcher layer too. Where researchers spent roughly half their week on list building, that share collapses to closer to 5% — the residual is configuring searches and reviewing edge cases. The released capacity goes to candidate development, client-side qualification, and the harder-to-substitute research the firm actually pays its researchers for.

This researcher-reallocation effect is separable from the per-biller funnel lifts in Part IV.1. Under B01's framing, gross revenue per recruiter on a typical firm view roughly doubles when the list-building load comes off the bench — from ¥36M to ¥72M per recruiter per year. Senior-biller take-home at the 45% tier follows the same multiplier, from roughly ¥16.2M to ¥32.4M per year. The two effects — funnel lift and researcher reallocation — compound at the firm level.

“Same headcount. Twice the meetings per recruiter. ¥55M of recovered margin a year, on a 12-biller desk — before the revenue lift.

IV.3 THE CORPORATE BUYER'S MATH

The corporate budget-holder's LinkedIn redirect math.

If you hold the budget — if your name is on the procurement form when the LinkedIn RPS renewal comes around — here is the math the renewal cycle does not usually surface. The line is treated as a fixed cost. It is not a fixed cost. It is a sourcing decision with two answers, and only one of them scales the way the function needs it to scale.

Take the seat math first. At the 6–10 user RPS tier, one seat lists at roughly ¥687,000 per year — about \$4375 at the ¥157 per dollar FX rate. Each seat ships 100 InMails a month, 1200 across the year. A typical Japan in-house team uses around 660 of those — about 55% utilization — with the remainder expiring unsent. The recruiter using the seat spends roughly twenty-six hours a week inside the LinkedIn interface, by audit. Most of that time is not InMail composition. It is search environment, Boolean refinement, and longlist triage — the search-environment layer named at category level in Part I.3.

What the seat is actually paying for.

Three components are stacked into the ¥687,000. The InMail allotment, which the team uses about 55% of. The search environment, which the team uses heavily and which sits inside a third-party walled garden. The brand layer — the comfort of recruiting on the platform every candidate already has an account on. Most procurement reviews treat the seat as one thing and renew on auto-pilot. The decomposition makes the trade explicit.

IV.3 CONT. · THE REDIRECT

SAME ¥687,000 PER SEAT PER YEAR · TWO DESTINATIONS

What the yen buys, on each side.

LINKEDIN RPS · 6-10 USER TIER

1200

INMAIL ALLOTMENT PER SEAT PER YEAR

~660

TYPICAL ANNUAL USAGE (ABOUT 55%)

~26 hours/week of recruiter time inside the interface. Search environment is the heaviest layer.

HEADHUNT.AI · SAME ¥687,000

10,775

RANKED QUALIFIED-CANDIDATE MATCHES, SCORED AGAINST YOUR ROLE

¥63.75

COST PER RANKED MATCH, DRAFT SCOUT MAIL INCLUDED

Same recruiter, same hours – +38% more candidate meetings per recruiter, per the Q1 2026 production data.

LinkedIn pricing assumes the 6-10 user RPS tier. Headhunt.AI unit cost is current Q1 2026 production rate. FX: ¥157 per dollar. Source: B04 §03, §05.

The honest read of this comparison is not “drop LinkedIn.” It is “know what each yen is buying.” If the team uses about half its InMail allotment, the seat is largely paying for the search environment and the brand layer — not the outreach. That may be the right trade for some teams. For others, especially those scaling the sourcing function rather than maintaining it, the redirected yen produces more matches, more meetings, and a per-match unit economic that is internally legible to procurement in a way the seat is not.

IV.3 CONT. · BUILT IN JAPAN, FILED IN JAPAN

What gets unlocked at the team scale.

On the team view, the line item compounds. A five-recruiter Japan in-house team running at five RPS seats spends roughly ¥3.4M per year on the seats alone. Redirecting that same yen to Headhunt.AI credits produces about 54,000 ranked qualified matches across the team — candidates scored against the team's actual open requisitions, returned ranked, with draft scout mail per row. The lift on candidate meetings per recruiter on the desk, per Part II.2 and the production-data Part III.1, is +38% on the same hours.

Per Part IV.1 and IV.2, the downstream effect on placements per recruiter compounds the same way it does for an agency biller — the lever is hours per qualified meeting, and the lever does not care whether the meetings end in an agency placement or an in-house hire. The corporate function gets the same compounding through a different revenue ledger.

The compliance layer that closes the case.

There is one more component the renewal decision does not usually surface. The Japanese regulatory regime for AI candidate aggregation is narrower than most procurement frameworks assume. Headhunt.AI is one of six services filed under the 第4号 category of the 特定募集情報等提供事業 — the narrow category of regulated AI candidate-aggregation services — out of 1642 total filed services. That filing matters for the same reason the Rikunabi precedent matters: Japanese regulators have demonstrated they will act on candidate-data handling, and the precedent is enforceable. Part V walks through the compliance frame in detail; the ten-question vendor self-audit in Part V.3 is the version any buyer can run against any vendor, not just Headhunt.AI.

For the renewal review specifically: the question stops being whether the line is justified and starts being which line, at this scale, is structured for the regulatory regime the function will be operating under for the rest of the decade.

“The renewal is not a sourcing question. It is a unit-economics question — with two answers, and only one of them scales.

IV.4 YOUR DESK, YOUR NUMBERS

Your desk, your numbers, your projection.

The three sections above are the math at three reader scales — the recruiter, the principal, the corporate buyer. This section is the same math, applied to six concrete desk configurations. It is a worksheet, not a calculator. The Handbook is a PDF; the arithmetic is yours to run on your own desk. What this page does is hand you the structure to run it.

Every number in the grid on the next page is derived from the same canonical inputs the rest of the Handbook leans on. Nothing is invented. The inputs are listed below for cross-check. Carry your own placement fee, your own seat count, your own bench size; the structure stays the same.

CANONICAL INPUTS

What the grid on the next page is built from.

Average placement fee (illustrative)	¥4M
Per-meeting expected revenue (measured)	¥107,676
Margin recovered per placement vs database route	¥608,000
LinkedIn RPS seat · 6–10 user tier	¥687,000/year
Per-recruiter placement lift (compounded Q1 2026 lifts)	+0.5/month
Per-recruiter capacity post-Headhunt.AI	50+ meetings/month
Pre-Headhunt.AI baseline	25–30 meetings/month

HEADLINE ROI ON CREDITS

17.2×

The grid that follows is not a forecast. It is a structural projection — what the arithmetic above produces when you apply it to the scale of desk you actually run. The numbers are directionally sound; the precision depends entirely on how cleanly your own inputs map onto these baselines.

IV.4 CONT. · SIX SCENARIOS

SIX DESKS · SAME ARITHMETIC · DIFFERENT SCALES

The grid is the worksheet.

DESK · PERSONA	CURRENT SOURCING SPEND	HEADHUNT.AI REDIRECT	YEAR 1 SWING (STRUCTURAL)
1-biller boutique SOLO PRINCIPAL · PART IV.1	Database stack ~¥1.2M/year	Self-serve credits (B02 100-credit entry)	+6 placements +¥10.8M take-home (45% tier)
3-biller team PARTNER-LED FIRM · PART IV.2	Database + LinkedIn ~¥5M/year	Credits + 1-2 RPS seats retained	+18 placements ~¥14M margin recovered + revenue lift
12-biller mid-size PRINCIPAL-RUN FIRM · PART IV.2	Database + 5 RPS seats ~¥13M/year	Credits at scale + 1-2 RPS seats	+72 placements ~¥55M margin recovered + revenue lift
30-biller multi-vertical MULTI-DESK PRINCIPAL · PART IV.2	Enterprise database ~¥32M+/year	Enterprise credits + ATS pass-through	+180 placements ~¥135M margin recovered + revenue lift
5-recruiter in-house team CORPORATE TALENT LEADER · PART IV.3	5 RPS seats ~¥3.4M/year	Redirect to credits retain 1 RPS seat	~54,000 ranked matches +38% candidate meetings, same hours
15-recruiter enterprise HEAD OF TA · PART IV.3	15 RPS seats + ATS ~¥10.3M/year	Enterprise credits + ATS pass-through (Part II.5)	~162,000 ranked matches +38% candidate meetings · audit-ready

Agency rows: “Year 1 swing” shows incremental placements (+6/yr per biller) plus margin recovered at ¥608K per baseline placement (B05 anchor). Revenue from incremental placements compounds on top, separately. Corporate rows show match volume and same-hours meeting lift. Treat as directional. Source: B04 §06, B05 §05/§09.

“The arithmetic is yours to run. The numbers in this Handbook are the receipts.”

V.1 TWO LAWS, NOT ONE

Two laws govern AI candidate aggregation in Japan, not one.

Compliance discussions about AI recruiting in Japan often collapse into a single conversation about data privacy. The actual legal picture is two parallel regimes, each with its own regulator, registration, penalty structure, and audit cadence. A platform can comply with one and remain unlawful under the other.

APPI — the 個人情報保護法, or Act on the Protection of Personal Information — governs every record about an identifiable Japan-resident individual, regardless of where the operator is incorporated or where the server sits. The 職業安定法, the Employment Security Act as amended October 2022, governs the act of providing candidate information to employers — which is what every AI sourcing platform does at the moment it returns a ranked candidate list to a paying client.

LAW 1 · DATA PRIVACY

個人情報保護法

Act on the Protection of Personal Information (APPI)

REGULATOR	PPC · 個人情報保護委員会
SCOPE	All personal data handling
REGISTRATION	Not required
REACH	Extraterritorial · Art. 171
MAX PENALTY	¥100M (corporate)

LAW 2 · RECRUITING OPERATIONS

職業安定法

Employment Security Act (as amended October 2022)

REGULATOR	MHLW · 厚生労働省
SCOPE	Recruiting information provision
REGISTRATION	届出 required · 4 categories
REACH	Japan-resident candidates
MAX PENALTY	6 mo. / ¥300K · Art. 65(7)

Both laws apply at the same time. A platform incorporated in California, scraping LinkedIn data, processing on US-based AI infrastructure, and selling subscriptions to Tokyo recruiters needs

V.1 CONT. · WHERE MOST PLATFORMS BREAK

Acquisition is where most foreign platforms break first.

APPI Article 20 requires data to be acquired by fair means. Public visibility on the source site does not by itself make the means fair. If the source platform's terms expressly prohibit automated collection — LinkedIn User Agreement Section 8.2 is the canonical example, prohibiting crawlers, bots, scrapers, and automated extraction — then scraping does not sit within what Article 20 contemplates, regardless of whether the data is technically reachable to a logged-out browser. The 2022 ESA amendment then added a second layer: operating an unfiled crawler-type service is a criminal offense under Article 65(7) of the 職業安定法. Two laws, one acquisition decision.

The foreign-processor problem, in one paragraph.

“Our servers are in Japan” is a common defense from foreign AI platforms. The PPC has answered it directly in its General Guidelines Q&A (Q10-25): when a domestic operator entrusts the handling of personal data to a foreign third party, the operator must implement security measures based on an understanding of the personal-information-protection regime in the foreign country — even where the personal data itself is stored on servers located within Japan. The implication for AI recruiting: if your model runs on US-based OpenAI, Anthropic, or Google infrastructure, you are processing personal data through a foreign entity even if the storage layer is in Tokyo. APPI Article 23 外的環境の把握 (external-environment understanding) attaches, and Article 28 cross-border analysis attaches with it. Only the EU and the UK are on Japan's adequacy whitelist; the United States is not.

“APPI is the law most platforms address. The Employment Security Act is the law most platforms cannot.

V.2 SIX OF 1642

Six of 1642 services are filed in the category that matters.

As of June 1, 2025 — the most recent published MHLW aggregate, released March 2026 — 1283 entities had filed as 特定募集情報等提供事業者 (Specified Recruitment Information Providers) in Japan. Among them, 1642 distinct services were registered. The four legal categories are not evenly populated. One of them is almost empty.

MHLW STATUS REPORT · AS OF JUNE 1, 2025 · PUBLISHED MARCH 2026

Filed services by 号 classification.

CATEGORY	DESCRIPTION	SERVICES
第1号	Job postings at employer request	1502
第2号	Job postings without employer request (aggregators)	132
第3号	Candidate information at candidate request	623
第4号	Candidate information collected without request	6

A single entity can file in multiple 号 categories. Total services exceed total entities for that reason. Source: B07 §01, drawn from MHLW 特定募集情報等提供事業概況報告書 aggregate, R7年6月1日 status, R8年3月 publication.

第4号 — candidate information collected without candidate request — is the category most AI sourcing platforms structurally need. Crawler-type services, public-data aggregators, universe-scale scoring layers, anything that builds a candidate pool without the candidate initiating the relationship. Six services are filed there. Headhunt.AI operates one of them.

V.2 CONT. · WHY THE RARE FILING IS RARE

Why 第4号 is the structurally hardest filing to satisfy.

The 2022 ESA amendment did three things at once. It expanded the definition of 募集情報等提供 to include crawler-type platforms collecting candidate data without per-candidate request. It created a notification requirement — 届出 — that must be filed before operation begins. And it attached new ongoing obligations on filed providers: an annual 概況報告書 status report, 的確な表示 (accurate display) of recruiting information, a documented complaints-handling window, and disclosure of the principal factors used to rank search results, per Article 43-6. The last item is a direct nod to AI scoring. The ranking-disclosure rule does not require the algorithm code to be public — those weights remain 営業秘密 — but the principal factors used in scoring must be publicly disclosed.

For a platform whose data supply chain begins with scraping, none of these obligations can be satisfied without re-architecting the upstream. The filing posture and the operating model travel together. That is why the 第4号 count stays at 6 in a registry that grew to 1642 — the bar is operational, not procedural.

Above the floor — 優良認定, the audited certification ceiling.

Beyond the basic 届出 filing, MHLW commissions a voluntary audited certification — the 優良募集情報等提供事業者認定制度. The audit covers seven categories: legal compliance, accurate display, personal-information handling, information disclosure, advertiser-side vetting (審査), complaints handling, and other governance items. As of April 1, 2026, 42 entities are certified across all four 号 categories. Among the 6 entities in 第4号, exactly one currently holds the certification. Certification term is 3 years, renewable. The audit is the closest thing Japan has to a publicly visible compliance benchmark that does not depend on the vendor's own marketing.

“Six of 1642. The rarest filing category is the one most AI sourcing platforms quietly need.”

V.3 A SELF-AUDIT ANY BUYER CAN RUN

A self-audit any buyer can run today.

Ten questions a procurement team, legal counsel, or talent leader can put to any AI sourcing vendor on their shortlist. The checklist is vendor-neutral — it is the same set of questions you should ask Headhunt.AI, the same set you should ask any competitor, and the same set you should be able to answer about your existing tools. Score one point per question the vendor can support with a document delivered within 48 hours. Anything else is a no.

The foundational five — sourcing posture and filings.

01 Is the vendor filed as a 特定募集情報等提供事業者 with MHLW?

Ask for the 届出受理番号 (notification receipt number) and the 号 classification. 第4号 is the answer for any platform that aggregates candidate data without per-candidate request. See Part V.2.

02 Does the vendor's privacy policy specifically disclose AI scoring, candidate matching, presentation to clients, and scout-mail generation as purposes of use?

APPI Articles 17, 18, and 21 require purpose to be specified, publicly disclosed, and not exceeded. Generic "to improve our services" language does not cover scoring. The Rikunabi case is the operative precedent — see question 09.

03 For each candidate record, can the vendor produce the data source and acquisition path?

APPI Article 20 requires lawful means. Public visibility is not by itself fair means. If the source's terms of service prohibit automated collection — LinkedIn User Agreement Section 8.2 is the canonical example — scraping is outside Article 20.

04 Where is the AI inference layer actually running, and on whose infrastructure?

"Our servers are in Japan" does not answer the question if inference happens elsewhere. The PPC Q&A Q10-25 is explicit: 外的環境の把握 (Article 23) attaches to foreign processors even when storage is local.

05 Are the principal ranking factors publicly disclosed in compliance with ESA Article 43-6?

The factors are public-disclosure items; the algorithm code and the calculation procedure can remain proprietary as 営業秘密. The MHLW Q&A on the 2022 amendment is explicit on what must and must not be disclosed.

V.3 CONT. · THE OTHER FIVE

Five more — candidate posture, regulator history, your own exposure.

06 Has the vendor performed an APPI Article 23 外的環境の把握 assessment for each foreign jurisdiction it processes in, and is the assessment documented?

The United States is not on Japan's adequacy whitelist; only the EU and the UK are. Cross-border consent or Article 28 appropriate-safeguards paths apply, plus Article 23 security obligations on top.

07 Has the vendor received a 勧告 (corrective recommendation) or comparable administrative guidance from the PPC or MHLW?

The Rikunabi precedent (Recruit Career, 2019) established that AI prediction and scoring of candidate data must be specifically disclosed and that hashing fails when the receiving party can re-identify. Public record disclosures and any open guidance items should be on the table during procurement.

08 For email outreach functionality, does the vendor obtain opt-in consent from each candidate before commercial messaging?

Required by 特定電子メール法 (the Act on Regulation of Transmission of Specified Electronic Mail). Opt-in record-keeping should be reproducible per recipient.

09 Does the buyer's organisation understand its own derivative liability under APPI Article 30?

The recipient of personal data has a confirmation duty: verify the provider's identity, the acquisition path, and the lawful basis. A non-compliant upstream platform transfers a piece of the compliance problem to the recipient. "We just bought the data" is not the defence Article 30 anticipates.

10 Can the vendor produce the 届出受理通知 and the data-source documentation, in writing, within 48 hours of request?

A platform that cannot produce both, on a Japanese client's request, in a few days, is a platform whose risk has been transferred to the client. This is the closing-condition test for procurement.

“The buyer's liability is real. The audit is yours to run. The answers should not be hard to get.”

VI.1 WHERE YOU SIT

Where you sit today, and where you sit in eighteen months.

Most Japan recruiting operations — agencies and in-house teams alike — fall on a 5-stage progression in how they produce candidate meetings. The stages are not subtle. Where the operation sits today determines what its meetings-per-recruiter number will look like in eighteen months — and what its competitive position will look like at the same point.

MATURITY MODEL · 5 STAGES · INDICATIVE MEETINGS PER RECRUITER PER MONTH

The ladder, with the numbers attached.

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

Meetings-per-month figures are indicative, drawn from peer-agency conversations and ESAI's own desk at each stage during 2018–2026. Stage 4 is the current-state ESAI desk position — the 16-week receipts in Part III.1 are produced from this stage. Source: B01 §08.

VI.1 CONT. · THE CURVE AND THE OVERRIDE

The curve is steeper than it looks.

Eighteen months from Stage 1 to Stage 3 is realistic on a working desk. Eighteen months to Stage 4 is aggressive but achievable on a single desk — the desk producing the 16-week receipts in Part III.1 reached Stage 4 over roughly that timeframe, with a measurement window that opened in August 2024 and a hands-off cutover that completed on November 1, 2025. Not every desk needs Stage 4. Some niches reward Stage 3 indefinitely, particularly very-senior-end search where the override (below) does more work than the base rate. But the slope from Stage 1 to Stage 3 is steeper than most operators expect, because the contributions in Part II.3 compound rather than add.

Hands-off is the floor, not the ceiling.

The 17.2× in Part III.1 is what an agency gets without trying. No human reviews any candidate before contact; no human edits any scout mail before it sends. That is the base rate. The ceiling is higher. Every scout mail in Headhunt.AI is fully editable before sending. Every candidate in the ranked list can be excluded with one click. Recruiters who selectively override — improving message tone for a senior candidate, removing a known wrong-fit profile — produce results above the base rate, particularly in role pockets where domain knowledge concentrates. The base rate gives back the hours; the override is where the recruiter's judgment concentrates on the work that pays.

“The curve is steeper than it looks. Sixteen weeks of hands-off is not the long version. It is the proof of where the curve ends.”

VI.2 THE OBJECTIONS, CONSOLIDATED

The objections we hear, consolidated.

Seven objections, surfaced consistently across the past twelve months in conversations with peer agency principals, corporate talent leaders, and procurement teams. Each gets the concession before the rebuttal. The aim is not to win on every line — the aim is to identify the parts of the pushback that are legitimate and the parts that are pattern-matching to old AI tools.

“We tried an AI sourcing tool before. It did not work.”

Fair pushback. Most operators who have been around long enough to evaluate AI tools have seen at least one that did not work, and most often the one they saw was scoring-only — a longlist filter bolted onto a manual outreach flow. The lift from a scoring-only tool is real but bounded. The lifts in Part II.2 compound across the funnel because three distinct components run together: universe-scale aggregation, scoring against the role, and AI scout-mail generation. A head-to-head test on the desk’s hardest current role is the cleanest way to separate “this particular tool, then” from “this class of system, now.” That test is in Part VI.3.

“Japan is different. Global AI tools do not translate.”

Largely correct, and the reason Headhunt.AI is built on a Japan-focused profile universe with bilingual scoring rather than a North-American model retrained on translated data. AI scoring works decisively for roles where the qualifying signals are profile-readable — bilingual finance, IT, sales, commercial, HR, marketing, GTM, operations, supply chain, product, legal, and most engineering subfields, which together cover roughly 80–90% of mid-market billable volume on a typical desk. For the narrow niches where the qualifying signal is not readable from a profile, the override in Part VI.1 is the operating answer.

“This is a database, not a recruiter.”

Correct, and intentionally so. Headhunt.AI is the sourcing layer underneath the recruiter, not a replacement for the recruiter. Qualification, meeting strategy, client positioning, and closing are where the fee comes from, and those parts stay human. What changes is what the recruiter spends the rest of the week on. The work the database removes — 60–70% of the week per Part I.1 — is the work that produces no per-recruiter differentiation. That redirect is the entire argument.

VI.2 CONT. · FOUR MORE

“The receipts are from your own desk, not ours.”

Fair, and the entire reason the test in Part VI.3 exists. We can show 16 weeks of hands-off production on a working desk (Part III.1), 526 days of inbound reply logs across a structural change (Part III.2), and a same-everything-else comparison design (Part III.3). What we cannot show is what the same system does on your specific roles, in your specific market segment, with your specific candidate definition. That is what 100 free credits, one JD, and two minutes are for.

“Compliance risk is too high for our company to pilot AI sourcing in Japan.”

Inverted, in our reading. Part V is the operative chapter: as of June 2025, 6 of 1642 filed services sit in the 第4号 category that AI candidate aggregation legally requires. The compliance risk lives with the foreign AI sourcing tools that operate without that filing — and APPI Article 30 transfers a piece of the risk to the Japanese buyer. For a TSE-listed enterprise, the safer compliance posture is the one with the filing, the documented data sources, and the published ranking factors. The ten-question self-audit in Part V.3 is vendor-neutral by design; the same questions you ask Headhunt.AI you should ask any competitor.

“Our recruiters will not use it.”

True at first, and the fix is not training. The fix is transparency. Recruiters trust scoring when they can see what the score is built from — tenure pattern, company history, role fit on specific dimensions — not just the number. Headhunt.AI exposes the principal factors per candidate in the ranked list, in compliance with ESA Article 43-6 and in service of recruiter adoption. Once the reasoning is visible, scoring stops being magic and starts being a screening tool the recruiter can argue with.

“We already have LinkedIn and BizReach. We do not need another layer.”

Fair pushback on cost, less fair on capability. The decomposition in Part IV.3 walks through what ¥687,000 a year of LinkedIn Recruiter actually buys for a Japan team — 660 used InMails out of 1200 allotted, plus a search environment plus a brand layer. The same ¥687,000 redirected to Headhunt.AI credits buys roughly 10,775 ranked qualified candidates at ¥63.75 per match. The BizReach-class decomposition in Part IV.2 makes the parallel case for per-placement-fee databases. The two tools are not substitutes on the same axis; the comparison is what each does per unit of redirected spend. On our own desk, the LinkedIn bill is down 90% post-redirect.

“Fair pushback. The point is not that you will believe us. The point is what the test shows on your own desk.”

VI.3 THE TEST

The test you can run this week.

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

01 Sign up at headhunt.ai.

No card. No demo call. No sales contact unless you ask for one. The 100 free credits land in your account immediately.

02 Pick one open role on your desk.

The honest test is your hardest current role — the one your team has been working for months. Mid-market and contingent roles in segments where AI scoring works well (Part VI.2) produce the cleanest signal.

03 Paste the JD.

The platform returns up to 1000 ranked candidates from the 4M+ Japan-focused universe in 1–2 minutes. The first 100 candidates above the score-50 threshold come from your free credits. One credit equals one qualified candidate at score 50+.

04 Read the list. Ask the binary question.

Are any of these candidates ones your team has not 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 the team. If no, walk away. Cost: zero. Time: two minutes.

The 100 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 the first list is full of candidates your team has already worked, you have your answer and you walk away. If it is not, you have a different answer and the rest of the conversation gets simpler.

“That is the offer.”

closing THE HONEST TAKE

The honest take.

This Handbook has made one argument across six parts. The market diagnosis in Part I is that the binding constraint on a Japan recruiting operation in 2026 is not fees, not candidate quality, not market conditions — it is recruiter hours per qualified meeting. The mechanism in Part II is that AI-first sourcing redirects 60–70% of those hours away from the work that produces no differentiation. The production evidence in Part III is sixteen weeks of hands-off operation, 123,675 candidates contacted, 1260 qualified meetings, and a 17.2× return on credits measured rather than projected. The economics in Part IV are persona-specific. The compliance frame in Part V is the chapter most foreign AI sourcing tools cannot pass. The path forward in Part VI is a single test that takes two minutes and costs nothing.

None of that is a forecast. It is a description of what is already running on one desk, written for the operator who is deciding whether to test it on another.

Not every reader will act. Some will close this document and continue with the same desk, the same tools, the same calendar split between sourcing and closing. That is a choice, and there are practice areas and seniority bands where the choice is defensible — at the very senior end of search, in narrow technical niches where the qualifying signal is invisible to AI, on small desks where the override does more work than the base rate. For most desks, in most segments, in most of the Japan mid-market, the math the Handbook has been walking through points in one direction.

The hardest read is the one for the operator who recognizes the diagnosis, recognizes the mechanism, and recognizes the math — and still does not act. The cost of doing nothing in 2026 looks like the present, which is precisely what makes it feel safe. The present is moving.

closing CONT. · THE REMINDER

The test in Part VI.3 takes two minutes and costs nothing. Whatever the result, you walk away with a real answer for your own desk rather than a position on someone else's.

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

About Headhunt.AI.

Headhunt.AI is built and operated by ExecutiveSearch.AI K.K., founded in 2018 as Japan's first AI-first recruiting firm and a subsidiary of Monstarlab Inc. (TSE: 5255) since 2023. The platform is the operational layer that runs underneath our own desk. Every receipt in this Handbook is from that desk.

We operate a 4M+ Japan-focused candidate profile universe, bilingual in English and Japanese, with the 第4号 特定募集情報等提供事業 filing (届出受理番号 · 発行待ち / pending issuance — see Part V.2 for the category framing). Since August 2024 the desk has run on a 100% AI-driven outbound recruiting model, with the November 2025 cutover to fully Headhunt.AI-generated target lists producing the sixteen-week receipts in Part III.1 and the 526-day reply log in Part III.2.

Headhunt.AI

THE TEST YOU CAN RUN THIS WEEK

100 free credits. One JD. Two minutes. No card.

Sign up at headhunt.ai, paste the JD for the role your team has been working hardest, and read the list. If the candidates are new to your team, you have your answer. If not, walk away. Cost: zero.

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

Methodology.

The Handbook draws on three datasets, each constructed for a different argument and each with its own methodological notes. This appendix walks through what each dataset is, how it was built, and what it does and does not establish. The presentation is consultative — the caveats are part of the work, not a defensive footer.

A.1 · The decision-gap trend (Part I.2).

Data covers March 2024 to March 2026 — 25 months — across corporate clients in Japan, drawn from ESAI's contingent-search desk. Full sample: 3852 resumes sent, 1030 first interviews, 385 second interviews, 165 final interviews, 74 placements. Stage-to-stage progression takes roughly two weeks; for monthly aggregation we apply a +1 month lag (second interviews in month M against finals in month M+1).

The primary trend test is a Mann-Kendall non-parametric test on the rolling six-month 2nd-to-Final conversion rate. Result: $z = -2.42$, $p = 0.015$. The test was chosen because it does not assume normality and is resistant to outliers — desirable properties for a small-sample monthly series with the kind of dispersion typical of agency placement data. The same test was applied to every other funnel stage with the same methodology. Only 2nd-to-Final shows a significant declining trend; Final-to-Offer shows an improving trend but does not reach significance at $p < 0.05$.

Breakpoint analysis tested every possible split in the monthly series. The strongest break is at July 2025 (before: 52% average; after: 37% average; Welch's $t = 1.32$, $p = 0.20$). The breakpoint does not reach individual significance at $p < 0.05$, but the overall trend does.

Caveats.

25 months is a modest sample. December 2025 was an unusually strong month (8 placements). Team capacity changes in late 2025 may partially contribute to the trend signal. Placement distribution is overdispersed (variance/mean = 1.82, above the Poisson expectation of 1.0), which is why quarterly or longer windows are the minimum reliable forecasting unit. The analysis locates the symptom; the AI sourcing thesis in Parts II and III proposes a structural answer to the cause.

A.2 · The 526-day reply-log comparison (Parts III.2 and III.3).

Two periods of inbound-reply logging on the ESAI desk. Pre-period: November 21, 2024 through October 31, 2025 — 345 days, daily inbound average 16.4 replies. Post-period: November 1, 2025 through April 30, 2026 — 181 days, daily inbound average 29.1 replies. Total measured window: 526 days, 10,932 logged inbound replies. The pre/post boundary is the November 1, 2025 cutover to Headhunt.AI-generated target lists.

The headline daily lift is +78% ($29.1 \div 16.4 - 1$). The cleaner like-for-like comparison restricts to weekdays only — the post period also includes a weekend-send pause introduced in Q1 2026 that depresses weekend daily averages without changing the underlying production rate. On weekdays only, daily averages move from 18.2 to 35.8, a +97% lift. The Handbook reports both numbers; the +97% weekday figure is the cleaner inference about the structural change.

What was held constant across the 526 days.

Same recruiter team size, headcount held constant across both periods. Same scout-mail engine — Headhunt.AI's AI generator, adopted August 1, 2024 — was the sending engine on day one of the measurement window. Same outbound sending setup, in the same configuration. Same client mix and market: mid-market Japan recruiting across bilingual finance, IT, sales, commercial, HR, and marketing practice areas. The only structural change introduced on November 1, 2025 was the source of the target list — recruiters and researchers built it by hand before; Headhunt.AI generated it after. The scout-mail engine wrote the messages either way.

Caveats.

This is a same-everything-else comparison on a working desk, not a randomized controlled trial. The evidentiary class is structural-comparison, not laboratory-causal inference. The 17-month window mitigates seasonality and ramp confounds — each period spans multiple quarters and includes both peak and trough months — but does not eliminate them. The November 2025 boundary is a structural break rather than a random assignment. The strongest defensible read is that on a working agency desk with the listed dimensions held constant, switching list source from manual to Headhunt.AI produced the measured lift. Stronger causal claims belong in a controlled environment that working agencies cannot run.

A.3 · The sixteen-week production window (Part III.1).

January 1, 2026 through April 30, 2026 — 16 consecutive weeks on the ESAI desk, covering Q1 2026 plus the first month of Q2. Across that window: 123,675 candidates contacted, 3868 replies, 1260 qualified meetings. 100% of candidates contacted were sourced by Headhunt.AI. 100% of scout mails sent were authored by Headhunt.AI. Zero human review of either before send. Conversion rates: 3.13% reply rate; 32.57% reply-to-meeting rate; one meeting per 98 outreach attempts on the combined ratio.

What hands-off means in practice.

The desk did not stop having a Candidate Manager team. It did not stop having Business Managers. It did not stop talking to clients. What it stopped doing was the work that occupied 60–70% of the recruiter week in Part I.1 — Boolean searches, manual scoring, scout-mail drafting from scratch. That work is the work that produces no per-recruiter differentiation. It is the work the sourcing layer can do without the recruiter in the loop. Hands-off describes the sourcing layer, not the desk.

The override (referenced in Part VI.1).

Every scout mail in Headhunt.AI is fully editable before sending. Every candidate in the ranked list can be excluded with one click. The 17.2× figure is what the autonomous system produces when no human reviews any candidate before contact and no human edits any scout mail before it sends — the floor of the system, deliberately measured without the override engaged so any agency adopting Headhunt.AI knows what it gets without trying. Recruiters who selectively override produce results above the base rate, particularly in narrower role pockets where domain knowledge concentrates.

Caveats.

16 weeks is one window. The 526-day reply log in A.2 is the longer comparison. The 17.2× figure is the unit-economics ratio on credits (¥1,720,788 in expected revenue per ¥100,000 of credits), not gross margin and not net P&L — Part II.1 carries the per-meeting unit-of-work decomposition. Average placement fee for the 2026 ESAI desk is ¥4,266,675 across the practice areas listed in A.2; readers operating at materially different fee structures should re-run the unit math for their own desk. The Handbook's position is that the methodology travels; the specific ratios do not, by default.

App. B GLOSSARY

Glossary.

Terms used in the Handbook, with the section in which each is introduced or most fully treated. Definitions are Handbook-native, drawn from across all nine briefings but rewritten for the consolidated context.

Credit PART II.1

The unit of consumption on Headhunt.AI. One credit equals one qualified Japan candidate matching the role at ESAI Score 50 or higher. Credits do not expire.

ESAI Score PART II.3

The 0–100 ranking number Headhunt.AI assigns each candidate against a specific role. Principal factors are publicly disclosed per ESA Article 43-6; algorithm weights remain proprietary as 営業秘密.

ATS pass-through PART II.5

The Enterprise-tier feature that scores an existing ATS's candidate pool on the same axis as the public-data pool. Database-level integration only — never self-serve upload.

届出受理番号 PART V.1

The notification receipt number issued by MHLW after a 特定募集情報等提供事業 filing is accepted. The primary procurement evidence that a vendor has filed.

APPI PART V.1

個人情報保護法, the Act on the Protection of Personal Information. Japan's data-privacy law, extraterritorial under Article 171.

Hands-off operation PART III.1

Sourcing-layer operation in which no human reviews any candidate before contact and no human edits any scout mail before send. The 16-week ESAI desk window in Part III.1 ran in this mode.

Scout mail PART II.3

An outreach email authored by Headhunt.AI's scout-mail generator and sent through the agency's own infrastructure. Fully editable before send; the override (App. A.3) operates on this layer.

Qualified meeting PART II.1

A first-time scheduled meeting with a candidate who has cleared the agency's qualification screen. The unit of work the Handbook returns to throughout.

第4号 PART V.2

The MHLW 特定募集情報等提供 filing category for candidate information collected without candidate request. The category that AI candidate aggregation legally requires; 6 of 1642 filed services sit here.

職業安定法 PART V.1

Employment Security Act. The recruiting-operations law in Japan, as amended October 2022 to include crawler-type candidate platforms.

Mann-Kendall APP. A.1

A non-parametric trend test that does not assume normality. Used in Part I.2 / Appendix A.1 on the rolling 6-month 2nd-to-Final conversion ($z = -2.42$, $p = 0.015$).

Override PART VI.1

Manual recruiter intervention on top of the hands-off base rate — editing a scout mail, excluding a candidate. The base rate is the floor; the override is where domain judgment concentrates.

App. B CONT. · GLOSSARY

Universe-level scoring PART II.3

Scoring run against the full Japan-focused candidate universe (4M+ profiles) rather than against a hand-built longlist. The change in scoring scope is the structural lever Part III.3 measures.

RPS · Recruiter Corporate PART IV.3

LinkedIn Recruiter Professional Services / Recruiter Corporate. The flagship seat-based product; ~¥687,000 per seat per year for a 6–10 user tier. Decomposed in Part IV.3.

Contingent exposure PART IV.2

The continued per-placement fee that a per-placement-fee database can claim for any historical contact between the firm and a candidate, even after the firm exits the subscription.

Decision gap PART I.2

The structural 2nd-to-Final advance-rate drop from 49% to 33% measured across 25 months on the ESAI corporate desk. Hiring did not slow; the committee evidence bar rose.

Rikunabi precedent PART V.3

The 2019 PPC 勧告 actions against Recruit Career establishing that AI prediction/scoring of candidate data must be specifically disclosed and that hashing fails when the receiving party can re-identify.

Two-pool ranking PART II.5

The single ranked list produced by scoring the public-data pool and the Enterprise ATS pool on the same axis. The reader sees one list; two pools fed it.

InMail PART IV.3

LinkedIn's in-platform messaging credit. An RPS seat carries 1200 InMails per year, of which typical Japan team usage burns about 660.

Per-placement-fee database PART I.3

BizReach-class candidate database charging a three-layer cost: subscription, per-placement usage fee (20–30% of recruiting fee), and tracked-contact contingent exposure. Category named in Part I.3; firm-view math in Part IV.2.

優良認定 PART V.2

The MHLW-commissioned voluntary audited certification 優良募集情報等提供事業者認定制度. 42 entities certified as of April 2026; 1 of the 6 第4号 entities currently holds it.

外的環境の把握 PART V.1

APPI Article 23 external-environment understanding obligation — attaches to overseas processors regardless of where the storage layer sits. The PPC Q&A Q10-25 is the operative reference.

App. C READING ORDER

Reading order for the briefing series.

The Handbook is the consolidated thesis. The nine briefings are the persona-specific source documents — each tells the same story from one reader's seat, with the analysis and disclosures the Handbook compresses. The map below names each briefing, its primary reader, and the section a new reader should open first if drilling into one briefing rather than starting from the cover.

BRIEFING	TITLE	READER	START AT
B01	100,000 Yen Per Meeting	<i>Agency principals</i>	§08 · The maturity model
B02	Less Sourcing, More Closing	<i>Working recruiters</i>	§09 · How the 100 credits work
B03	The Decision Gap	<i>Hiring leaders</i>	§02 · The 49% → 33% drop
B04	Reduce Your LinkedIn Recruiter Spend by Half	<i>Corporate budget-holders</i>	§05 · The decomposed ¥687K
B05	The Database Tax	<i>Agency principals</i>	§03 · Three layers of cost
B06	Enrich Your ATS	<i>Agency principals</i>	§08 · Zero LinkedIn TOS exposure
B07	Is Your AI & Sourcing Stack Illegal in Japan?	<i>Compliance + legal</i>	§10 · The self-audit
B08	Trusting the AI	<i>Agency principals</i>	§01 · The 17.2× number
B09	+78% Replies	<i>Skeptical operators</i>	§02 · What stayed the same

All nine briefings, EN and JP editions, are available at headhunt.ai/insights. Each briefing carries the same editorial conventions as this Handbook: spaced em-dashes, full Fraunces section titles, the same REMINDER and honest-take closing pattern.