BRIEFING 05 · AGENCY-TO-AGENCY

The Database Tax.

Per-placement-fee databases became one of Japan's dominant agency sourcing channels. They no longer have to be the default — one placement fee, redeployed through Headhunt.AI, surfaces 13,000+ qualified candidates.

13,000+

CANDIDATES PER PLACEMENT FEE

One placement fee redeployed → 13,000+ ranked candidates on Headhunt.AI.

Built by an agency. Tested daily on our own desks since 2018.

inside THIS BRIEFING

What follows, in thirteen sections.

A briefing for agency principals running Japan recruiting desks. The argument is one sentence: **find every qualified candidate in the known universe before you ever pay a per-placement database fee.** Same yen, far more candidates, zero new contractual exposure.

01 **The big number.**
One placement fee redeployed surfaces 13,000+ qualified candidates on Headhunt.AI.

02 **Where the agency fee pool went.**
The scale of the per-placement-fee category, in public numbers.

03 **Three layers of cost.**
Subscription, per-placement fee, and the third layer — the one most agencies don't price.

04 **What one disclosure surfaced.**
¥402M. The one publicly documented cost of getting database-fee accounting wrong.

05 **The unit comparison.**
¥608K more margin per placement. ~¥55M per year on a 12-biller desk.

06 **What changed, and what didn't.**
Public profile data isn't new. The throughput cost of working it is what collapsed.

07 **Flag the exposure before the placement.**
ATS integration shows which of your candidates still carry tracked-contact risk.

08 **The new operating model.**
AI-first on every search. Fall back to per-placement databases only on the residual.

09 **Where the saved spend goes.**
Recruiter pay, team growth, and an asset you actually own.

10 **Common pushback.**
Seven objections from peer principals, answered honestly.

11 **The test.**
¥75,000. One exhausted role. One ranked list. One simple verdict.

12 **Self-audit.**
Seven diagnostic questions on database concentration risk.

13 **The honest take.**
Where this lands. Closing exhortation.

BizReach is named only where the underlying claim is sourced to public Visional TSE filings or BizReach's published pricing. All other framing is category-level.

01 THE BIG NUMBER

One placement fee. 13,000+ candidates.

A typical Japan agency placement made through one of the major per-placement-fee databases — BizReach, Recruit Direct Scout, doda X, AMBI, and similar agency-side scout platforms — carries a database usage fee of ¥700K–¥1.05M, 20–30% of the agency's recruiting fee. That same yen, deployed through Headhunt.AI credits at the enterprise rate of ¥63.75 per qualified match, surfaces between 10,980 and 16,470 qualified candidates from a 4M+ Japan-focused database — ranked, scored, ready to evaluate.

That is a five-figure spread on the most expensive cost line in your sourcing operation. More importantly: **those 13,000+ candidates have not been touched on a per-placement-fee platform.** Outreach goes out cold through your own LinkedIn, your CRM, your existing infrastructure. No tracked-contact record. No new contingent fee exposure.

UNIT COMPARISON · ¥10M BASE PLACEMENT · 25% PER-PLACEMENT RATE

¥875,000 placement fee = 13,725 qualified candidates surfaced.

¥875K	¥63.75	13,725
PER-PLACEMENT FEE AT 25% OF RECRUITING FEE	PER QUALIFIED MATCH · HEADHUNT.AI ENTERPRISE TIER	QUALIFIED CANDIDATES SURFACED FOR ONE PLACEMENT FEE

Math: ¥875,000 ÷ ¥63.75 per match = 13,725 candidates. At 30%, the spread rises to 16,470. At 20%, it is 10,980. The floor is 10,000+ — case-independent.

> Put Headhunt.AI in front of every search. Watch per-placement spend fall.

02 WHERE THE FEES WENT

Where the agency fee pool went.

The dominant Japan recruiting database is BizReach, operated by Visional Inc. (TSE: 4194). Visional's own TSE filings describe the scale of the per-placement-fee category at the platform side. The numbers below come from the company's March 2026 disclosure of H1 FY7/2026 results — read them as scale, not characterization.

VISIONAL INC. · TSE: 4194 · H1 FY7/2026 DISCLOSURE

The platform side of the ecosystem.

¥46.6B	+26.2%	82%	¥99.2B
H1 FY7/2026 CONSOLIDATED REVENUE	YEAR-OVER-YEAR GROWTH	SHARE FROM BIZREACH SEGMENT	FULL-YEAR FY7/2026 FORECAST

Source: Visional H1 FY7/2026 results, TSE filing dated 17 March 2026. BizReach segment H1 revenue ¥38.3B (+19.2% YoY), 42.7% adjusted operating margin before corporate allocation. Standard headhunter plan: ¥600K per six months, plus per-placement performance fees. BizReach's public materials report 41,800+ direct employer subscribers and 9,700+ active headhunter accounts.

The platform side compounded for over a decade on a structurally similar input: agency and corporate fees collected on placements of candidates the platforms themselves do not employ, train, or place. The candidates are public. The agencies do the work. Each agency placement made through these platforms is, by definition, a placement that *could* have been made another way. Until very recently, the throughput cost of working those alternatives made the database channel the rational primary path for many agencies. Section 06 explains what changed.

> **"The platform's margin is not the issue. The structural cut from your P&L is.**

03 THREE LAYERS OF COST

Three layers of cost, not one.

Most agency principals price their per-placement-fee database costs as two line items: the subscription fee and the per-placement usage fee. There is a third layer — and it is the one that compounds quietly across thousands of outreach actions over a contract year.

01 **Subscription** — **~¥1.2M / yr**

The standard headhunter plan on the dominant Japan platform is ¥600,000 per six months, paid lump-sum in advance regardless of placements. Smaller platforms run lower; enterprise tiers run higher.

02 **Per-placement usage fee** — **~¥700K–1.05M**

20–30% of the agency's recruiting fee, paid to the operator on every placement of a candidate sourced through the platform. Roughly ¥700K–¥1.05M on a typical ¥10M base role at a 35% agency fee.

03 **Tracked-contact contingent exposure** — **Every scout mail sent**

Most standard Japan recruiting database agreements include audit and attribution provisions under which a placement traceable to outreach activity on the platform may incur per-placement fees, regardless of the channel through which the candidate ultimately responded. Terms vary by operator; the category-level structure is consistent.

Σ **Annual category spend, mid-size desk** — **¥8M–22M+**

Subscription + ~10–20 placements/year × per-placement fee, plus residual exposure on under-attributed placements.

The first two layers are predictable. The third — the ongoing claim against placements that ever trace back to a tracked outreach — is not. **That third layer does not exist on Headhunt.AI.**

> Sourcing cost is paid once. Or paid every time the candidate moves.

04 WHAT ONE DISCLOSURE SURFACED

The cost of getting the contract math wrong.

On 20 January 2023, For Startups, Inc. (TSE: 7089) — a publicly listed Japan recruiting agency — disclosed a retroactive restatement covering FY3/2018 onward, recognizing ¥402 million in additional cost of sales owed to *multiple* human resource database operators. ¥118M was unpaid principal. ¥283M was contractual damages — a 2.4× penalty.

The cause, in the company's own words, was a 規約の誤認 — a misinterpretation of fee-attribution terms in cases where a single career-changer was simultaneously registered on multiple databases. Under several operators' contracts, fees could be owed to *each* operator the candidate had touched, not only to the one through which the placement was finalized. The company resolved it through discussions with various database operators.

FOR STARTUPS, INC. · TSE: 7089 · 20 JANUARY 2023 · TIMELY DISCLOSURE

The numbers attached to one episode.

¥402M
RECOGNIZED IN COST OF SALES, RETROACTIVE

¥118M
UNPAID PRINCIPAL

2.4×
CONTRACTUAL PENALTY MULTIPLIER

Source: For Startups, Inc. timely disclosure (TDnet), 20 January 2023; FY3/2023 Q3 earnings call. Counterparties described as "各求人データベースの企業" (each of the job database companies, plural) and resolved through "各社と相談した結果" (consultation with each operating company).

> The contract structure is shared across the category. The question is whether yours holds.

04 CONT. · WHAT CHANGED

This is the only public, on-the-record example of what an agency's contingent exposure to per-placement-fee databases looks like once audited and exercised. Most peer firms have similar exposure profiles. They are simply not yet under securities-disclosure pressure to surface them.

> **A category lesson, not a single-operator story.**
>
> *The structural mechanism — fee-attribution terms that can claim a placement when the candidate has been touched on more than one database — is a property of how these contracts are written across the category. For Startups themselves used plural language describing the counterparties. The question for your desk is not which operator was involved in their case. It is whether your attribution math, across every database your recruiters touch, would hold up to a similar audit.*

What changed: throughput, not access.

Per-placement-fee databases became one of the dominant Japan agency sourcing channels through the 2010s for a defensible reason. Public profile data on LinkedIn was always the largest pool, but working it at the throughput an active desk requires cost more recruiter labor than the database alternative cost in fees, for many roles. The cost-per-qualified-candidate math made these databases the rational primary channel for a lot of agencies.

In the last 24–36 months that calculation has shifted. AI scoring against full profiles at scale is now viable. The episode above is no longer a cautionary tale about a specific counterparty — it is a cautionary tale about **over-concentration in any single sourcing channel during a period when the AI-first alternative did not operate at this throughput.** That period has ended.

> **"The exposure is the cost of channel concentration. The fix is channel diversification.**

05 THE UNIT COMPARISON

¥608K more margin, per placement.

Forget the platform. Walk through one ¥10M placement, two ways. Same client. Same fee from the client. The only thing that changes is which sourcing layer surfaced the candidate. Take BizReach — the dominant per-placement-fee platform — at the typical 25% rate, against a Headhunt.AI-first path.

PATH A · VIA BIZREACH

What you keep, per placement.

Client pays your agency (35% of base)	¥3,500,000
Pay BizReach (25% of recruiting fee)	-¥875,000
Pay Headhunt.AI	¥0
Net margin to your firm	¥2,625,000

PER-PLACEMENT MARGIN *Baseline*

PATH B · VIA HEADHUNT.AI FIRST

What you keep, per placement.

Client pays your agency (35% of base)	¥3,500,000
Pay BizReach	¥0
Pay Headhunt.AI (40 meetings × ~¥6,667)	-¥267,000
Net margin to your firm	¥3,233,000

PER-PLACEMENT UPLIFT +

¥608,000

At 25% — the midpoint of the 20–30% range — the spread is **+¥608K per placement**. A 12-biller desk doing 90 placements converts that into approximately **¥55M of annual margin recovered.**

> Same fee. Different margin.

06 WHAT CHANGED, WHAT DIDN'T

What changed. And what didn't.

If qualified candidates are the question, most principals have one answer: the per-placement-fee database's registered pool. But qualified candidates have always sat in three places. What changed in the last 24–36 months is not the pool — it is the throughput of working it.

Public profile data isn't new.

LinkedIn has been the largest pool of working professionals in Japan for years. Every senior recruiter knows this. What was hard, until recently, was the throughput of actually working that data — Boolean strings, manual longlisting, hand evaluation of tenure and company-tier patterns. Working that pool by hand cost more in recruiter labor than the per-placement-fee database alternative cost in fees, for many roles.

That is the honest version of why per-placement-fee databases became one of the dominant Japan agency sourcing channels through the 2010s. It was not that the candidates weren't elsewhere — they always were. It was that working them elsewhere was more expensive in recruiter labor than the database fee was in cash.

AI scoring is what is new.

Production-quality scoring against full profiles, in Japan, with bilingual signal evaluation, tenure-pattern analysis, and company-tier sequencing has only become viable at scale in the last 24–36 months. The same researcher who used to evaluate a handful of qualified profiles per week of effort can now evaluate thousands in a few minutes, with consistent JD-relative scoring. **The pool was always there. The throughput is what is new.** This is what changes the cost equation against per-placement-fee databases that were once cheaper per qualified candidate.

> **“The pool was always there. The throughput is what changed.**

06 CONT. · THREE LAYERS HEADHUNT.AI SURFACES

Three layers, all now surfaceable.

Headhunt.AI works against three layers of candidate data. All three were surfaceable in theory before. The economics of surfacing them are what shifted.

Layer one. Public profile data.

Public professional profile data — LinkedIn primary, with adjacent sources stitched in — is the largest, most current dataset of working professionals ever assembled. Headhunt.AI sits on 4M+ Japan-focused profiles, refreshed continuously. The competitive question is not whether the data exists; it is whether your search method actually surfaces it. Most senior Japan candidates do not write profiles in keywords Boolean queries match. They are in the data. They never come up in a Boolean search.

Layer two. AI enrichment on top.

Headhunt.AI scores tenure pattern, company-tier sequence, career-trajectory inflection, bilingual signal in surrounding context, and adjacent-industry relevance — fields the underlying profile doesn't expose. The AI evaluates each candidate against the job in full sentences. The Boolean query evaluates against a string.

Layer three. Your own ATS.

Connect Headhunt.AI to your ATS and the same scoring runs across every candidate your firm has ever sourced. The placements you paid per-placement-fee databases to surface over the past decade are still in the system — never re-evaluated. Headhunt.AI re-scores them against the new role. **Candidates you already paid to source become re-placeable on a new search.** No new database fee liability. No new contingent fee.

> "The candidates didn't disappear. They are sitting in your ATS — paid for once, never re-evaluated since.

07 THE EXPOSURE MAP

Flag the exposure. Before the placement.

Most per-placement-fee databases include an ownership window — varying by operator and contract — during which a candidate one of your recruiters has touched on the platform carries a contingent fee, regardless of which channel ultimately produces the placement. The third cost layer in section 03 only becomes a surprise invoice because no one in your firm can see, candidate by candidate, who is still inside which window.

ATS integration, every shortlist.

When Headhunt.AI integrates with your ATS, every candidate already in the system gets re-scored against the new role. The integration also surfaces whether you or your team has logged any prior contact through a per-placement-fee database — which platforms, when, and whether the candidate falls inside that operator's typical ownership window. The output is a column on the shortlist: clean, exposed to one operator, or exposed to multiple — the For Startups scenario in section 04, made visible in advance.

Build shortlists deliberately.

With the exposure column visible, your recruiters build every shortlist with full information. Include a single-operator-exposed candidate when the fit justifies the contingent cost. Filter out multi-exposed candidates entirely on roles where the math doesn't work. Either way, the decision is deliberate. **The third-layer cost stops compounding the moment your team has visibility into it** — not three months later, when an invoice arrives.

Audit defense, built in.

When a database operator's account team eventually sends a fee adjustment notice covering the past 24 months — and at the scale of this category, that is a question of when, not if — your finance function has line-by-line attribution. Every prior contact, every platform, every timestamp. The For Startups outcome in section 04 doesn't happen to firms operating with this visibility. It happens to firms operating in the dark.

> “What For Startups disclosed cost ¥402M to surface. The same exposure now shows up as a column on your shortlist.

08 THE NEW OPERATING MODEL

AI sourcing first. Per-placement databases as fallback.

The argument is not that you should cancel your per-placement-fee database contracts tomorrow. For most desks, that would be premature — these platforms have uses, particularly in active candidate flow and segments where AI scoring genuinely struggles. The argument is sequencing.

01 **Run Headhunt.AI on every search, first.**

Paste the JD, get the ranked list, surface candidates that Boolean queries on per-placement-fee databases would never raise. Reach out cold through your own LinkedIn account, your CRM, or whatever outreach infrastructure you currently use — **nobody on the per-placement platforms has touched them.**

02 **Reach the top-decile candidates before anyone else does.**

The AI-surfaced candidates are, by definition, the ones not currently being approached by every other agency hammering the same Boolean queries on the same databases. Reply rates run structurally higher.

03 **Fall back to your per-placement-fee databases only on the residual.**

If the AI list does not contain enough qualified candidates — which happens on niche searches and narrow technical specialties — your existing platform contracts are still there. **You just use them less.**

Desks running this sequence should expect platform-attributed placements to drop materially within two to three quarters. The subscription becomes a smaller, justifiable line item. The third-layer contingent exposure stops compounding.

> “Tools earn their keep, or they don't. Use them when they do.

09 WHERE THE SAVED SPEND GOES

Where the recovered margin goes.

The spend that comes off the per-placement-fee database line does not disappear — it becomes deployable capital. Three places it tends to go, and one place it should not.

Where it should go.

01 **Senior recruiter compensation.**

The highest-leverage of the three. A senior biller at typical Japan agency comp shares (35–45%) takes a meaningful slice of every additional placement; the firm captures the rest. Reinvesting recovered database spend into senior comp produces near-term retention and output gains.

02 **Researcher and associate hires.**

The AI layer absorbs the candidate-identification work that used to consume roughly 50% of researcher time. The same researcher headcount can now support 2–3× the billing volume. **Adding researchers — instead of replacing them — is the cleanest way to compound the productivity gain.**

03 **The agency's own candidate database.**

Every search you run through Headhunt.AI scores and structures candidates into a workspace your firm controls. Over 6–12 months, that workspace becomes a structured, AI-scored asset that **per-placement-fee databases never give you access to.** The data is yours.

Where it shouldn't go.

It should not become invisible. Recovered spend absorbed into general operating budget without a tracked line item disappears in the quarterly noise. Treat it as a separate budget category for the first 12 months and report against it monthly — otherwise the saving compounds for the firm but not for any individual scorecard.

> "The worst place to send your firm's margin. The best place to be visibly pulling it out of.

10 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.

"We get half our placements through per-placement databases. We can't risk that flow."

You don't have to. The recommendation is sequencing, not replacement. Put Headhunt.AI in front of every search. Keep the platform contracts live for the residual. Most desks find within two quarters that the residual is small enough to renegotiate the contracts down — but you do not need to make that decision until you have proof on your own desk.

"Our recruiters won't change workflow. They like what they have."

Recruiters resist tools that *add* work. Headhunt.AI subtracts work — the longlist building, the Boolean iteration, the noise filtering. The workflow change is a recruiter pasting a JD instead of writing search syntax. After the first week, the resistance we have seen runs in the other direction: recruiters who tested it want it on every search.

"Why not just use LinkedIn directly? It's the same data."

Fair pushback. The data overlap is real — Headhunt.AI's base layer is public profile data, LinkedIn-primary, with adjacent sources. The difference is throughput. LinkedIn-direct sourcing through Boolean queries puts the labor cost on your researcher: a handful of qualified profiles per week of effort, across all the Boolean iteration and manual evaluation. Headhunt.AI scores the same 4M+ profile universe against your JD in minutes and surfaces hundreds-to-thousands of ranked candidates. Same data, very different cost per qualified candidate.

"Per-placement databases have candidates Headhunt.AI doesn't — actively job-seeking ones."

True. Active job seekers who have explicitly registered on a job-board database are accessible there in a way they are not on a public-profile-data base. For roles where active intent is the primary filter — junior-to-mid mass-market roles, certain volume hiring — that flow is real. **For passive senior candidates, which is where the highest-margin agency placements come from, the active-intent filter is a liability, not an advantage.**

10 CONT. · PUSHBACK

"If this is so good, why hasn't a peer firm we know already moved?"

Several have. The ones who moved early are not advertising it — for the same reason no peer agency loudly disclosed its database concentration until forced to by a securities filing. The migration is a margin advantage, and margin advantages are not the kind of thing you publish in a press release. You will mostly hear about it from peer principals at the same maturity stage you are currently at.

"The contract has a termination clause that makes early exit expensive."

Most do. We are not recommending early termination. We are recommending that you **stop being the firm that grows the contract.** Renew at lower seat counts. Negotiate down on per-placement fee tiers. Use the AI-first sourcing layer to make the reduction defensible to the platform's own account team — they do not want you to leave entirely, and they will accept smaller deals before they accept losing you. The leverage is highest when you have proof on your own desk that the alternative works.

"We tried other AI sourcing tools. They were noise."

Most are. The cleanest way to handle this objection is the test in Section 11, not a debate. Run the same JD through Headhunt.AI and through whatever else you have. We benchmark against the leading global tools regularly; on Japan-specific candidate data, the quality lead is significant. If the list is not visibly better on your roles, walk away. The test costs ¥75,000 — less than a single recruiter-day at most agency cost structures.

On the limits.

AI sourcing does not make per-placement-fee databases obsolete tomorrow. There are roles where active-intent filtering on a registered candidate base genuinely produces a better list — particularly volume mid-market roles. Use the platforms there. The argument is structural advantage in specific segments, **not a universal claim.**

> **Lead with AI first. Fall back to per-placement databases on the residual.**

11 THE TEST

A test you can run on a stuck role.

Everything in this brief is theory until your team puts it in front of a real role on your desk. Two minutes. Same-day answer.

THE TEST · IN ONE LINE

One stuck role. 500 ranked candidates. ¥75,000. No contract.

Pick a role your team has worked through your existing per-placement-fee database for two months or more. Run it through Headhunt.AI. Compare the AI-surfaced top 100 against the candidates your team has already approached.

If even one of the top 100 is new to your team, Headhunt.AI is finding profiles your current process is structurally missing — proof of concept on your hardest case, with zero added contractual liability.

01 **Buy a ¥75,000 credit pack.**
500 credits = up to 500 qualified candidate matches scoring 50+ on the ESAI Score. No subscription. No contract. Credits never expire.

02 **Pick the hardest open role on the desk.**
The honest test is the role you have already exhausted on your existing database. Mid-market and contingent roles in segments where AI scoring works well — bilingual finance, IT, sales, HR, marketing, GTM, ops, most engineering — produce the cleanest signal.

03 **Compare top 100 to your existing pipeline.**
The single binary question: *are there candidates on this list my team has not already approached?* If yes — even one of one hundred — you have your answer. ¥75,000 spent. Decision made.

> “¥75K spent. Fresh new candidates returned.

12 SELF-AUDIT

Seven diagnostic questions.

The right starting question is not "should we cancel our database contract." It is a set of **operational diagnostics on database concentration risk.** The point is to find which questions you can answer in numbers and which you can only answer in generalities.

- ☐ Do you know what percentage of your placements last quarter were attributable to candidates first sourced through your per-placement-fee databases, vs. candidates sourced through other channels — measured **per recruiter**, not estimated firm-wide?
- ☐ Do you know your **annualized total per-placement-fee database cost** across subscriptions and per-placement fees, expressed as a percentage of recruiting revenue? Most principals quote 10–15%; the actual figure, once subscription, per-placement fees, and contingent exposure are added together, is often higher.
- ☐ Have you mapped your **contingent fee exposure** on candidates messaged through any of your platforms but placed via another channel? If yes, what is the dollar range? If no, this is the third cost layer your P&L is not currently pricing.
- ☐ If a database operator's account team called tomorrow with a fee adjustment notice covering the last 24 months, would your finance function be able to defend your reported attributions **line by line** — or would you settle?
- ☐ What fraction of your shortlists last quarter were sourced from **passive candidates not actively registered on any job platform**? If below 50%, your sourcing is competing with every other agency hammering the same active pool.
- ☐ Have you run a **structured AI sourcing test** on a real open requisition in the last 12 months — or is your current view based on vendor demos and conversations with peers?
- ☐ If a peer firm announced 15% lower per-placement-fee database spend with the same number of placements tomorrow, what is your **concrete response?**

SCORE INTERPRETATION

6–7: Concentration risk is visible. Run the §11 test next.

4–5: Visibility exists; mechanism does not. The test is the fastest path to a concrete answer.

2–3: Concentration risk is real and likely under-priced.

0–1: Operating without visibility. The exposure is whatever the contracts say it is.

13 WHERE THIS LEADS

The honest take.

Japan's third-party recruiting databases — of which BizReach is the largest — built a real market. They created the infrastructure that allowed the country's mid-to-high-end placement business to scale through the 2010s. Some of that infrastructure remains genuinely valuable. **Some of it is a tax that compounds on agency P&Ls because, until recently, the AI-first alternative did not operate at this throughput.**

The alternatives now exist. The structural cost gap between AI-first sourcing and per-placement-fee platforms is wide enough that even partial migration produces meaningful margin recovery. The contingent exposure layer — the one most agencies have not priced — stops compounding the moment you stop sending tracked outreach through these platforms on candidates you could have reached another way.

None of this is a moral argument about any one operator. The operators are businesses doing what businesses do — they capture margin where the structure lets them. The question is whether your agency is going to keep being a line item on the category's revenue statement, or whether you are going to test the alternative this week and decide for yourself.

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 AI sourcing platform built and run by ExecutiveSearch.AI K.K. — a Tokyo recruiting firm operating an AI-first model since 2018, and a subsidiary of Monstarlab Inc. (TSE: 5255). We are not a vendor that wandered into recruiting. **We are a recruiting firm that built the platform we needed for our own desks, ran it for eight years, and now license it to peer Japan agencies.**

Our 4M+ profile Japan database, the ESAI Score, the bilingual scout-mail engine, and the production numbers throughout this brief are all from the same platform that runs our own desk today. The per-placement-fee database costs we describe in this brief are ones we paid for years ourselves before we built the alternative.

Headhunt.AI

START WITH THE ¥75,000 TEST

One stuck role. 500 candidate matches. One conversation.

Pick the role your team has been working longest. Run it through Headhunt.AI. Compare the top 100 against your existing pipeline. If the candidates are new to your team, you have your answer — and you got it for less than a single recruiter-day.

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