

BRIEFING 08 · RECEIPTS FROM PRODUCTION

# Trusting the AI got us 17.2× ROI.

*We promised agency principals a 15× return on Headhunt.AI credits. The first sixteen weeks of 2026, on our own desk, cleared 17.2× — with zero human review of any candidate or any scout mail. Receipts.*

---

# 17.2×

RETURN ON AI SOURCING CREDITS

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

---

inside THIS BRIEFING

# What follows, in eleven sections.

A second briefing for agency principals. Briefing 01 projected a 15× return on Headhunt.AI credits and a hands-off funnel. This is the production receipts — sixteen weeks of real data on our own desk.

- 
- 01 The receipts.**  
17.2× ROI on credits over the first sixteen weeks of 2026 — above the 15× we promised.
- 
- 02 Hands off the funnel.**  
100% of candidates found by Headhunt.AI. 100% of scout mails written by Headhunt.AI.
- 
- 03 The funnel, by the numbers.**  
123,675 candidates contacted. 3868 replies. 1260 qualified meetings.
- 
- 04 The unit economic atom, recalculated.**  
¥4,266,675 average fee · 39.625 meetings per placement · ¥107,676 per meeting.
- 
- 05 From ¥100,000 to ¥1,720,788.**  
The arithmetic that turns one credit pack into a 17.2× return.
- 
- 06 What hands-off frees up.**  
Recruiter hours move from sourcing to closing — the activity that actually makes money.
- 
- 07 Where the candidates come from.**  
Two pools, one ranked list — and your ATS gets refreshed as a side effect of running.
- 
- 08 The override.**  
The base rate is what you get when nobody touches anything. You can do better.
- 
- 09 Common pushback.**  
Seven honest answers to the questions principals ask about hands-off sourcing.
- 
- 10 A test you can run this week.**  
Two paths: a public credit-pack pilot, or a silent pilot the team does not have to know about.
- 
- 11 The honest take.**  
Where this leads if you do nothing, and where it leads if you act.
- 
-

## 01 THE RECEIPTS

# 17.2× return on sourcing credits.

Briefing 01, first hypothesized in January, projected that ¥100,000 of Headhunt.AI credits would convert to roughly fifteen qualified candidate meetings — and from there to about ¥1.5M of expected revenue. A 15× return. The first sixteen weeks of 2026 are now in the books, on our own desk, and the ratio came in above projection.

¥100,000 of credits produced 16 qualified meetings. At our 2026 average placement fee, those meetings carry ¥1,720,788 in expected revenue. That is a 17.2× return on every yen spent on credits. The math is shown step-by-step on page 09; the receipts are below.

2026 YTD · PRODUCTION RESULTS · ESAI AGENCY

## The unit economics, drawn from real production.

**17.2×**

RETURN ON SOURCING CREDITS

**16+**

QUALIFIED MEETINGS PER ¥100K

**¥1,720,788**

EXPECTED REVENUE PER ¥100K SPEND

**¥107,676**

EXPECTED REVENUE PER MEETING

*Sixteen weeks of production. Same recruiters as prior periods. Same fee structure. The variable is that 100% of sourcing ran autonomously — no human reviewed candidates and no human edited scout mails.*

## 01 CONT. · WHY THE NUMBER LANDED ABOVE PROJECTION

17.2× is two ROI points higher than the 15× we projected. The gap matters less than where it came from. Two things moved.

**Average fee per placement is up.**

*The 2026 average placement fee at our desk landed at ¥4,266,675 — about 6.7% above the ¥4M figure used in Briefing 01. The fee mix shifted into segments where AI scoring is strongest, which raised the per-placement number without a deliberate move upmarket. Each meeting now carries ¥107,676 of expected revenue instead of ¥100,000.*

The second move is the meetings-to-placement ratio. Briefing 01 used 40 as the working number — the industry benchmark. Our 2026 cohort is running at 39.625, almost identical, but slightly tighter. That tighter ratio combines with the higher fee to produce ¥107,676 per meeting versus the ¥100,000 we projected. Multiply by sixteen meetings per ¥100,000 of credits and you land at ¥1,720,788.

Neither shift was the result of human intervention in the sourcing funnel. Both reflect what the autonomous system produced when left alone for sixteen weeks.

“We projected 15×. The receipts say 17.2×. The recruiters did not work harder; they stopped working on sourcing.”

## 02 HANDS OFF THE FUNNEL

# Zero human review of any candidate or message.

Briefing 01 framed Headhunt.AI as a sourcing accelerator. The 2026 production data forces a sharper claim. Across the sixteen-week window, 100% of the candidates this desk contacted were found by Headhunt.AI, and 100% of the scout mails it sent were written by Headhunt.AI — with zero human review of either.

WHAT HUMANS DID AND DID NOT DO · Q1 + APRIL 2026

## The split between AI and human work in the funnel.

### AI · HANDS OFF

- Universe scoring across 4M+ profiles + your existing ATS
- Scout mail composition in business Japanese or English
- Three-message sequencing, auto-stop on reply

### HUMAN · WHERE LEVERAGE LIVES

- Reply triage and meeting confirmation
- The candidate meeting itself
- Client briefing, shortlisting, and closing

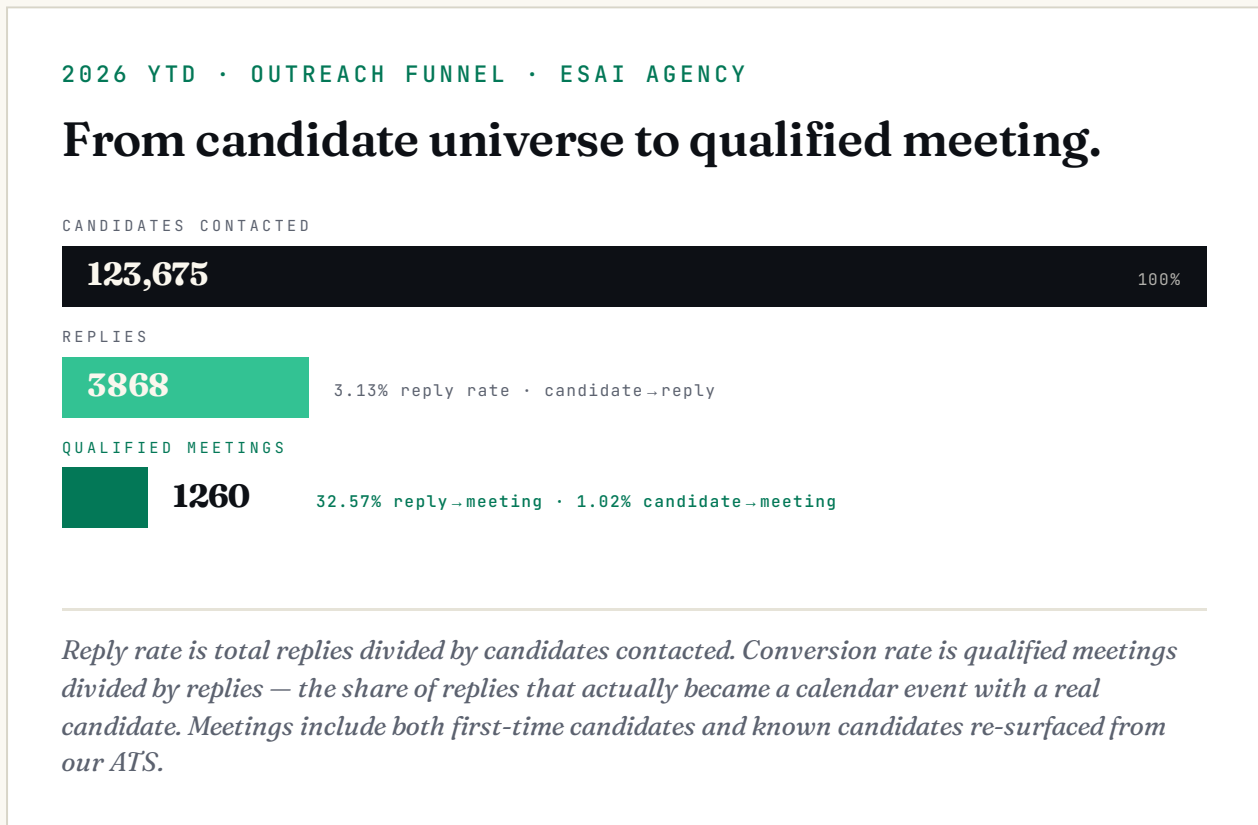
*The autonomous block on the left ran for the full sixteen weeks without human edits. The base rate in this brief is what an agency gets when nobody touches the sourcing layer at all.*

“The point is not that AI helps a recruiter source faster. It is that there are no recruiter hours in the sourcing funnel.”

03 THE FUNNEL

# The funnel, drawn from sixteen weeks of production.

Three numbers describe the full funnel: how many candidates the system contacted, how many replied, and how many of those replies converted into a qualified meeting on the recruiter’s calendar. The shape of the funnel is what produces the 17.2× number.




The 3.13% reply rate is what the autonomous system produces on cold-to-warm Japan-focused outreach with no human review. The 32.57% reply-to-meeting conversion is what the recruiter team produces once a reply lands — that is the human-leverage part of the funnel, and it is healthy. Sixteen meetings per ¥100,000 of credits comes from working backwards through this funnel; the arithmetic is on page 9.

03 CONT. · VOLUME ACROSS THE PERIOD

The 123,675-candidate total was not produced in a single batch. It was produced week-by-week across sixteen weeks of normal operating tempo. The shape of the curve is informative — autonomous outreach scales without adding recruiter hours.

2026 YTD · WEEKLY CANDIDATES CONTACTED · ESAI AGENCY

Sixteen weeks. No additional recruiter hours.



*Weekly candidate volume rises across the period as the system saturates each role's reachable universe and rotates onto new searches. The recruiter team did not add hours during the ramp; the lift came from removing sourcing work from the recruiter week, not from doing more of it.*

SHARE OF 2026 YTD CANDIDATES FOUND BY HEADHUNT.AI

100%

The other figure to remember:  
100% of scout mails were also written by  
Headhunt.AI.

“Volume is the easy part. The hard part is producing volume the recruiters trust without making them work the front of the funnel.”

## 04 THE UNIT

# The unit economic atom, recalculated.

Briefing 01 anchored on ¥100,000 of expected revenue per qualified meeting — derived from a ¥4M placement fee and 40 meetings per placement. Both inputs have moved slightly. The recalculated number is ¥107,676 per meeting.

## 2026 YTD · THE RECALCULATED ATOM

$$\begin{array}{rcccl} \mathbf{¥4,266,675} & \div & \mathbf{39.625} & = & \mathbf{¥107,676} \\ \text{2026 AVG PLACEMENT FEE} & & \text{MEETINGS PER PLACEMENT} & & \text{EXPECTED REVENUE PER MEETING} \end{array}$$

*Both inputs are measured, not estimated. The placement fee is the trailing 16-week average across closed placements. The 39.625 ratio is the trailing meetings-to-placements count over the same window. Industry benchmark for the ratio is 40; our cohort runs marginally tighter.*

The fee shift (¥4M → ¥4,266,675) is a mix effect, not a pricing change. We still charge 30–35% of annual salary on placements; the 2026 cohort skewed slightly higher in salary band, which raised the average. The meetings-to-placement ratio sits inside the normal band of 35 to 50 that we publish in Briefing 01.

“Every part of sourcing rolls up to one question. What does it cost to produce a qualified meeting? The answer, in production, is roughly ¥6250 in credits.

## 05 THE ARITHMETIC

# From ¥100,000 to ¥1,720,788.

Five steps connect a credit purchase to expected revenue. Each step is a measured number from our 2026 production cohort, not a projection. The math compounds, which is why the headline figure surprises people on first read.

**01 Spend ¥100,000 on Headhunt.AI credits.**

The system contacts roughly 1570 candidates with a personalized scout sequence — three messages per candidate, auto-stopping on reply. Total cost: ¥100,000. Recruiter time on this layer: zero.

**02 Receive ~50 replies at the 3.13% production reply rate.**

Roughly fifty replies land in the recruiter team's inbox over the campaign's run. The reply rate reflects autonomous outreach with no human review of either the candidate list or the message text.

**03 Convert ~16 of those replies to qualified meetings.**

The reply-to-meeting conversion rate is 32.57% in our 2026 cohort — the human-leverage part of the funnel. Roughly sixteen meetings reach the calendar.

**04 Each meeting carries ¥107,676 of expected revenue.**

From the recalculated unit economic atom on page 8: ¥4,266,675 average placement fee divided by 39.625 meetings per placement.

**05 Total expected revenue: ¥1,720,788.**

16 meetings × ¥107,676. From a ¥100,000 credit spend with no human hours added on top. A 17.2× return on the credit input.

“The recruiter team did not work the front of this funnel. The numbers are what the autonomous system produced.”

# Recruiter hours move from sourcing to closing.

Briefing 01 measured what consumes a typical recruiter week. 60–70% of the time goes to sourcing and qualification — universe-building, longlist triage, scout drafting, follow-up sequencing, scheduling. Hands-off sourcing removes that block. The hours have to land somewhere.

In our 2026 cohort, those hours land in the only place they can: on the human side of the funnel — the calendar, the conversation, the close.

## Where the freed hours actually went.

*Across the 16-week window, recruiter hours per week stayed flat (about 40 each). Hours allocated to sourcing dropped to roughly 5% — checking the AI's output and flagging issues. Hours allocated to candidate meetings, client briefings, shortlist construction, and closing rose proportionally. Same recruiters. Same week. Different mix.*

This is the answer to the question principals ask first: “If the AI does the sourcing, what do my recruiters do?” They do the work that actually produces revenue. Sourcing has never been the leverage activity. The leverage activity has always been the candidate meeting and the client conversation. Hands-off sourcing is the move that lets the recruiters do more of the high-leverage work without working longer hours.

“The leverage activity in recruiting is not finding the candidate. It is the conversation that follows. Hands-off sourcing buys back the hours where that conversation happens.”

## 07 THE TWO SOURCES

# Where the candidates actually come from.

Two pools feed every search. The agency's existing ATS, which holds candidates the desk has worked before. And 4M+ Japan-focused profiles on the open web. Headhunt.AI scores both pools against the role's specific criteria, deduplicates them, and routes a single ranked list. Candidates the recruiter has met before — and built context with — surface as readily as new candidates.

## THE TWO-SOURCE RANKED LIST

## One scoring model. Two universes. One ranked list.

## SOURCE 1 · YOUR ATS

Candidates the desk has spoken to before. Pulled live via a custom integration we build for each customer's stack — Bullhorn, Salesforce-based, Zoho, Workday, internal databases. No data migration; the ATS stays where it is. Phone numbers, email addresses, and any custom fields the customer wants flow straight onto the result so recruiters can act immediately.

## SOURCE 2 · 4M+ JAPAN PROFILES

Public web profiles focused on the Japan market. Scored against the same role criteria as the ATS pool. The candidates the desk has not seen yet — but should.

---

*Both pools score with the same ESAI Score. Duplicates collapse to a single record. The recruiter sees one ranked list with provenance flagged on each candidate.*

“We score your ATS alongside the open web — and return one list containing both.”

# Your ATS gets refreshed, not just searched.

Most agency ATS records are years stale. The candidate's title is from when you spoke to them. The company is from when you spoke to them. The contact information is whatever they left you with. By the time a relevant role opens for that candidate, half of what you have on file is wrong.

Headhunt.AI fixes this as a side effect of running. Every time the system scores an ATS candidate against a role, it checks the candidate's record against the live 4M-profile database. New title, new employer, new tenure pattern, new visible signals — the ATS record is updated in place. Across sixteen weeks of operation, our own ATS effectively re-cleaned itself, no recruiter hours allocated.

## What the enrichment touches.

*Current title, current employer, current tenure, visible career-trajectory signals, and any public profile changes — written back to the ATS field-by-field, with provenance flagged so the recruiter can see what changed and when. Phone numbers, email addresses, and other private contact fields live in the ATS already; those flow through onto the result intact, alongside the public-profile refresh, so recruiters have everything they need to reach out the moment they open the candidate.*

This is the part of the value that compounds. The longer Headhunt.AI runs against your ATS, the cleaner your ATS becomes. The cleaner your ATS, the better the next ranked list. The better the ranked list, the more meetings.

“The graveyard ATS becomes a living one. Not as a separate project — as a side effect of running searches.

# The base rate is what you get without trying.

17.2× is the production base rate. It is what the system produces when no human reviews any candidate before contact and no human edits any scout mail before it sends. The base rate exists so that any agency adopting Headhunt.AI knows the floor.

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. We see this on our own desk in narrower role pockets where domain knowledge concentrates.

## What this brief does not claim.

*The 17.2× number is not the maximum. It is what the autonomous system produces with the humans pulled out of the funnel. Recruiters who choose to spend 10 minutes per role tightening targeting or rewriting one or two scout mails for senior candidates can push the number higher. We publish the base rate because that is what an agency can rely on; the ceiling depends on how the agency uses the override.*

The right way to think about Headhunt.AI is not as an autonomous system that replaces recruiter judgment. It is as a sourcing layer that produces a defensible base rate without recruiter time, and that recruiters can selectively improve where their judgment is strongest. The base rate gives you back the hours; the override is where domain expertise concentrates.

“Hands-off is the floor, not the ceiling. The 17.2× is what an agency gets without trying. What it gets when the recruiters apply judgment is higher.”

# Seven honest answers about hands-off.

Seven questions principals ask once they hear “100% AI sourcing, zero human review.” Each gets a direct answer, not a deflection.

***“100% AI-written scout mails sound like spam waiting to happen.”***

Fair pushback. The reply rate is the proof. 3.13% on cold-to-warm autonomous outreach in Japan is well above the floor for spammy templated outreach, which lives at 0.3–0.8%. Recipients reply because the messages reference each candidate’s actual profile, current role, and visible career signals — not because the messages are short or aggressive. The composition model was trained on what worked across our own desk for years before it was made hands-off.

***“If the AI is wrong about a candidate, nobody catches it before send.”***

True at the candidate level — and that is by design at the base rate. But the composition is per-candidate, not per-template. If the system has bad evidence, the message is generic, the reply does not come, no harm is done. The cost of a wrong candidate that nobody opens is one credit. The cost of a senior recruiter spending an hour qualifying that same candidate is one recruiter-hour. The economics favor letting the autonomous layer run.

***“Our brand voice cannot tolerate AI-written messages.”***

Sometimes true at the senior end of the market. The override exists for this. Most agencies enable hands-off for their core mid-market role volume — where the math works decisively — and hold human review for senior or named-account roles. The base rate does not require uniform adoption.

***“We already run our own outbound infrastructure. Why pay for yours?”***

Of course you do — every recruiting firm has some outbound capability. The question is whether it produces the volume and conversion the hands-off rate produces. If you want state-of-the-art outbound at scale, two paths exist: a managed-service tier where we run the autonomous infrastructure on your behalf, or implementation consulting where we help your team build the same capability on your own stack. Both are heavily discounted for Pro and Enterprise Headhunt.AI customers.

***“Reply rate is one thing. Are these meetings actually qualified?”***

32.57% reply-to-meeting conversion is the answer. Reply rates are easy to manipulate; meeting conversion is not. A reply that converts to a calendar event with a real candidate at the right level is a meeting. The 1260 meetings on our desk produced placements at the standard 39.625-to-1 ratio — which means the meetings are qualified at the same standard our recruiters apply to manually-sourced candidates.

***“We already have an ATS. We do not need another database.”***

You do not get another database. Headhunt.AI scores your existing ATS in place via a custom integration — Bullhorn, Salesforce-based, Zoho, Workday, internal — alongside the 4M public profiles. Both pools feed one ranked list. The ATS does not move. Page 11 has the structure.

***“Why is the briefing on a base-rate number? Surely the real ROI depends on how we use the system.”***

The real ROI does depend on how an agency uses the system. We publish the base rate because base rates are what a business can plan against. The agency gets at least 17.2x by doing nothing in the funnel; whatever an agency adds via override is on top of that. We did not want to publish a number that depended on a recruiter being unusually skilled or attentive.

“The base rate is what we can defend in a board meeting. The override is where the agency’s own judgment compounds on top of it.

## 10 THE TEST

# A test you can run this week.

Everything in this briefing is theory until it is on your own desk against your own roles. Two ways to put it there. The first is the public credit-pack pilot — small, fast, transparent. The second, on the next page, is a silent pilot the team does not have to know about.

## PATH 1 · THE CREDIT-PACK PILOT, IN ONE LINE

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

## How it works.

### 01 Buy a ¥75,000 credit pack.

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

### 02 Pick one open role and paste the JD.

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

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

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

# Or run it as a silent pilot for one team.

Some agency principals want to validate the system without putting it through committee, change-management, or recruiter-buy-in. For those situations, the cleanest test is the one the team never sees coming.

Pick one billing recruiter who is consistently short on qualified meetings. We deploy Headhunt.AI against that recruiter's active roles, with their existing scout sequence rules and meeting-booking conventions. The team experiences it as "extra qualified candidates suddenly appearing in their inbox and on their calendar." They do not need to learn a new tool. They do not need to rewrite a workflow. They do not need to change anything. The pilot runs in the background; the recruiter just gets more meetings.

## What the recruiter sees.

*Replies start arriving from candidates they did not know they were chasing. Meetings book themselves onto their calendar at the normal pace. The recruiter does the meeting, runs the conversation, and writes the shortlist. Everything downstream of the reply is unchanged. The thing they never knew was happening was the thing that produced the upstream volume.*

This path is for principals who want to see whether hands-off sourcing actually moves the meeting count for one of their recruiters before they justify it as a firm-wide investment. If the count moves, you have your answer. If it does not, you also have your answer — and nobody on the team had to be convinced of anything in advance.

“The cleanest test is the one the team never knew was running. Either the meeting count moves or it does not.

## 11 WHERE THIS LEADS

# The honest take.

The receipts are now on the table. Sixteen weeks of production at ESAI Agency, every candidate found by Headhunt.AI, every scout mail written by Headhunt.AI, no human review. 17.2× return on credits. Recruiter hours flat. Meeting volume up.

Two questions follow. The first: is this base rate stable enough to plan against? On our desk, sixteen weeks is a meaningful sample — 123,675 candidates contacted, 1260 meetings on the calendar — and the variance across weeks is small enough that we treat the number as a working baseline. The second question: what does an agency do with the answer? That is the question the briefing leaves with the reader.

The agencies that move sourcing to a hands-off model in 2026 buy back recruiter hours that were previously locked up in the front of the funnel. The agencies that do not will operate against the same shrinking contingent fee pool with the same cost structure they have today.

---

## REMINDER

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

---

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

about HEADHUNT.AI · FOR AGENCIES

# Built by an agency. For agencies.

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

Our 4M+ profile Japan database, the ESAI Score, the bilingual scout-mail engine, and every production number in this briefing all come from the same platform that runs our own desk today. The 17.2× number is what the autonomous system produced for our recruiters. It is the same system, on the same terms, available to peer agencies.

## Headhunt.AI

START WITH THE ¥75,000 TEST

# 500 candidate matches. One conversation. No contract.

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

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