

255 MINUTES.

What it actually costs to hire one person, and what happens when you hand the admin to AI.

Adonye Orumbie

Director, Global Talent Acquisition - Adonye AI Agents Hub

THE PROBLEM

Recruiters weren't hired to do this.

Writing intake briefs from Slack messages

45 min / role

Screening 200 CVs by hand

1 full day / role

Writing outreach messages one by one

10 - 15 min / candidate

Chasing interview slots back and forth

2-3 days / hire

Building pipeline reports in spreadsheets

2 hrs / week

~255

minutes of admin

Per single hire - before the candidate
ever gets an offer

What if the job wasn't admin, but it was judgment?

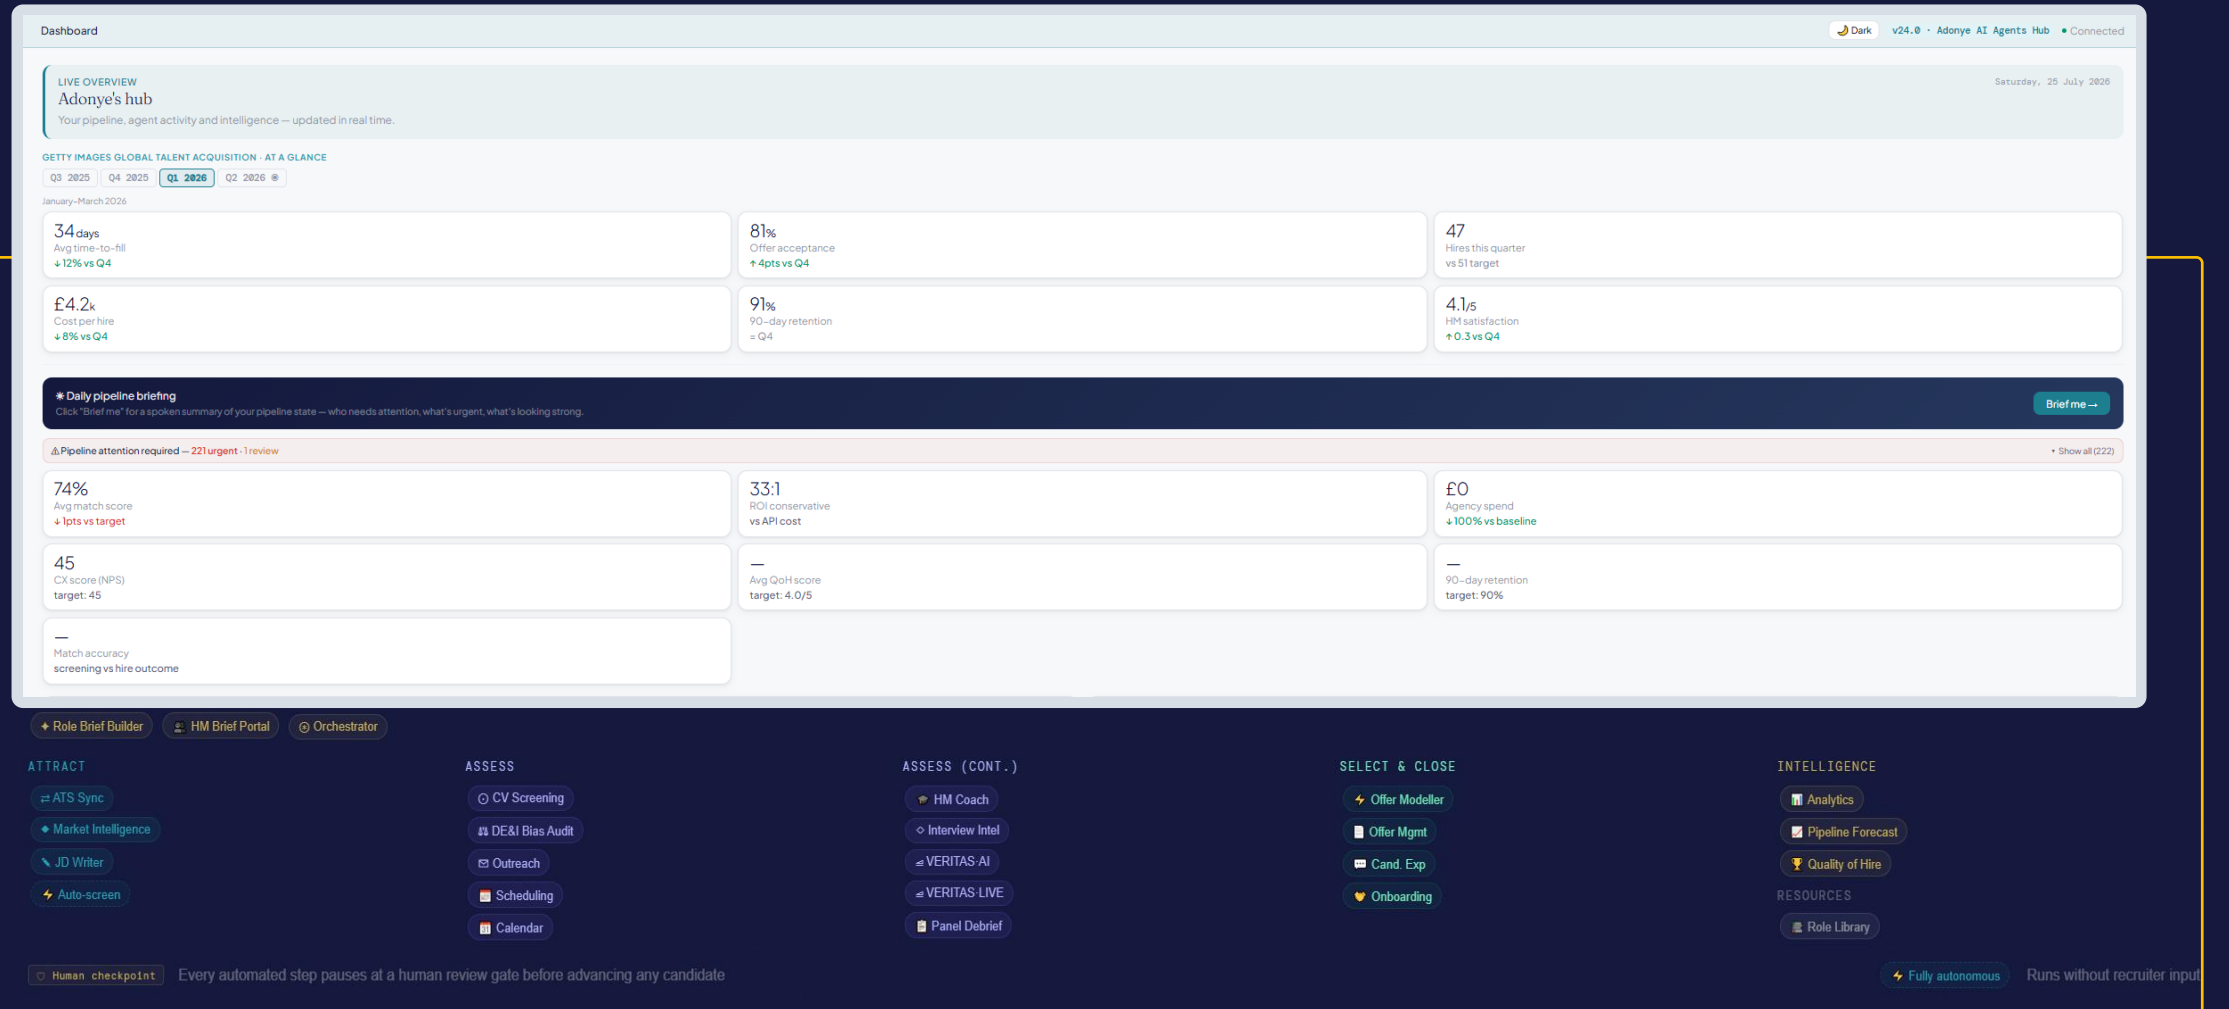
25 years in Talent Acquisition taught me the same lesson every quarter

So, I built something to take the admin, not the judgment.

HOW

25 AI agents. One hub. The full hiring lifecycle.

This is the actual hub, not a mockup. Every agent shares one candidate record, one pipeline, one source of truth.



ROLE BUILDING

One conversation. Five deliverables. One bias caught in real time.

◆ ROLE BRIEF BUILDER

One conversation. *Five deliverables.*

Adonye receives a Slack from the CPO: "We need a Senior Product Designer, Berlin, start Q3." She opens the Role Brief Builder and describes the role conversationally. The agent pushes back on vague or biased criteria — then generates everything simultaneously.

- ✓ Intake brief
- ✓ Job description (inclusive, bias-checked)
- ✓ Sourcing strategy
- ✓ Interview scorecard
- ✓ Outreach sequence

TIME SAVING

Manual process: **~90 minutes** → With hub: **8 minutes**

Role Brief Builder · Agent 19 · Pushing back on criteria

One conversation → 5 deliverables

AO Senior Product Designer for Berlin. Lars wants someone with 10 years minimum and ideally from a FAANG company.

The 10-year minimum correlates with age bias without improving outcomes — can we reframe as 'demonstrated senior-level impact'? On FAANG-only: this creates a diversity risk and pipeline volume problem. What is it about FAANG that matters to Lars specifically?

AI

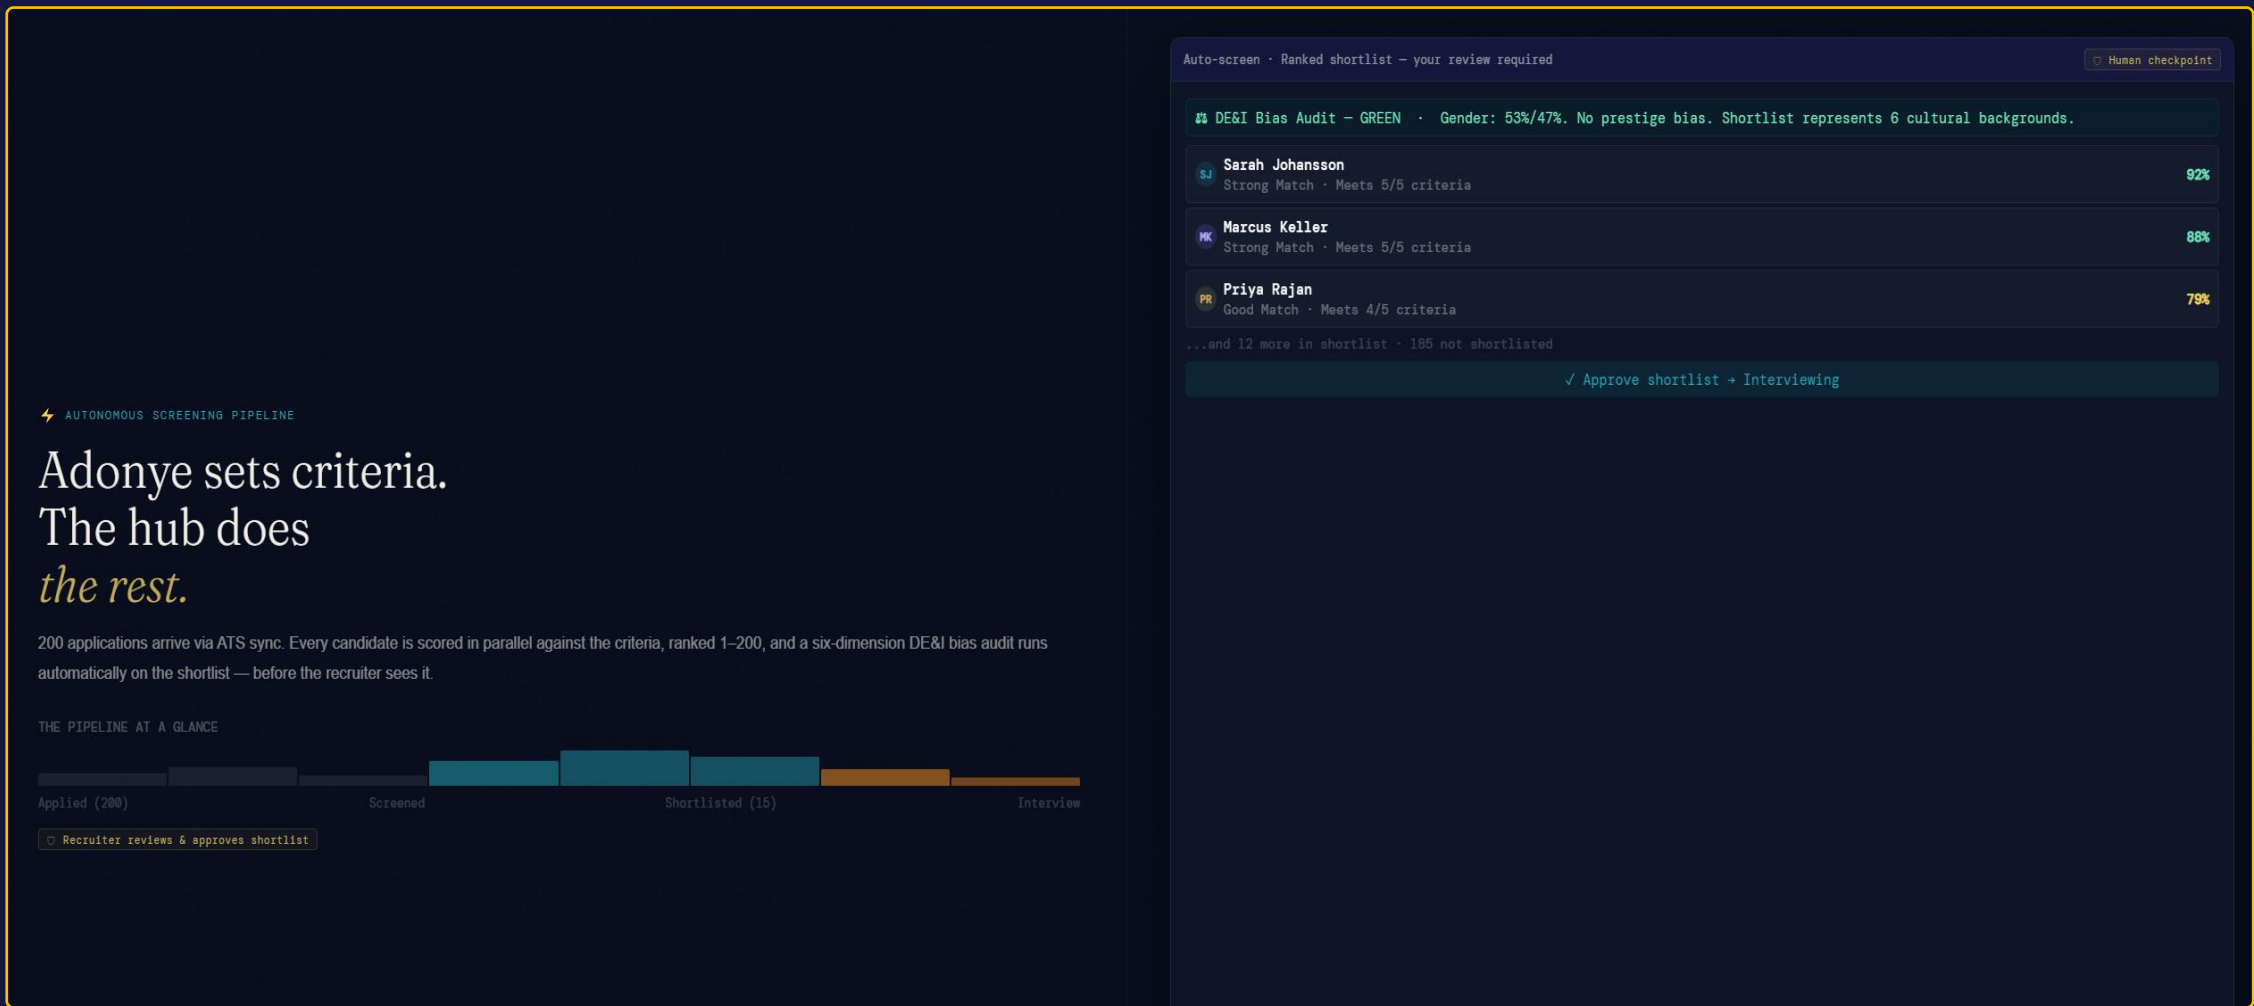
GENERATING 5 DELIVERABLES

✓ Intake brief ✓ Job description ✓ Sourcing strategy ⌕ Scorecard · Outreach

A CPO asks for "10 years minimum, ideally FAANG." The agent pushes back on the age and prestige bias before the brief is even built - 90 minutes of work done in 8.

APPLICATIONS ARRIVE

200 CVs. Ranked, scored, bias-audited - before the Recruiter sees them.



HUMAN GOVERNANCE

16 mandatory checkpoints. Agents assist - humans decide.

GOVERNANCE
Human-in-the-loop checkpoints
Sixteen mandatory decision gates across all agents where humans must review, approve or decide before the process continues. The agents accelerate work — humans remain accountable for every hire.

AI agents assist — humans decide
Every agent in this hub is designed to compress effort, remove bias, and improve consistency. None of them make final decisions about people. The checkpoints below are the moments where human judgement, accountability and legal compliance require a person to review, approve and sign off before the process moves forward.

Brief & JD approval Post-intake · intake Required TA partner reviews and approves the hiring brief and JD before it is published. Role Brief Builder outputs require the same sign-off. Owner: TA Partner High risk	Shortlist review Post-screening · screening Required TA partner reviews all screening results and approves the final shortlist before candidates are contacted or progressed. Owner: TA Partner High risk
Bias audit gate Pre-shortlist · bias Required The DE&I Bias Audit must be run on every shortlist before it is shared with a hiring manager. RED flags require mandatory human review before the list proceeds. Owner: TA Partner + HRBP Critical risk	Outreach message review Pre-send · outreach Required All AI-drafted outreach messages must be read, personalised, and sent by a human TA professional. No message is sent directly from AI output. Owner: TA Partner Medium risk
HM brief review & delivery Pre-interview · hmcoach Required The TA partner reviews the HM coaching brief before delivering it. Bias coaching notes must be shared as preparation support, not criticism. Owner: TA Partner Medium risk	Panel confirmation Pre-interview · scheduling Required Hiring manager confirms interview panel composition and availability before invitations are sent to candidates. Owner: Hiring Manager Low risk

Shortlist approval. Bias-audit sign-off. Hire decisions. Offer sign-off. Every one owned by a named person, not an agent.

HUMAN GOVERNANCE

16 mandatory checkpoints. Agents assist - humans decide.

 VERITAS authenticity review Post-interview · veritas Required All VERITAS risk scores must be reviewed by a qualified TA professional before influencing any hiring decision. Scores are probabilistic indicators only. Owner: TA Partner High risk	 Hire / no-hire decision Post-debrief · interview Required All hire decisions made by the hiring manager and confirmed by the TA partner. The interview intelligence agent structures the conversation — it never makes the call. Owner: Hiring Manager + TA Partner Critical risk
 Offer sign-off Pre-offer · offer Required TA Director and hiring manager approve all offer letters before sending. Offer Modeller scenarios require human judgement on which to proceed with. Owner: TA Director + Hiring Manager Critical risk	 Offer scenario approval Pre-negotiation · offermodel Required The recommended offer scenario must be approved by TA Director before being presented to the candidate. Walk-away positions require hiring manager sign-off. Owner: TA Director High risk
 Candidate communication review Pre-send · candexp Required Every AI-drafted candidate message must be read and approved before sending. Rejection messages require particular care and personalisation. Owner: TA Partner High risk	 Market intelligence validation Pre-brief · market Required Market Intelligence briefs are AI-generated estimates. All salary benchmarks must be validated against current sources before being presented to HMs or the CHRO. Owner: TA Director Medium risk
 Forecast sign-off Pre-presentation · forecast Required Pipeline forecasts are models with assumptions. The TA Director must review all assumptions and confidence levels before sharing with the CHRO or board. Owner: TA Director High risk	 Quality of hire data review Quarterly · qoh Required QoH pattern analysis surfaces probabilities, not decisions. The TA Director and HRBP must review AI-identified patterns before changing any hiring process or criteria. Owner: TA Director + HRBP High risk
 Role Brief Builder output approval Post-generation · rolebrief Required All five Role Brief Builder outputs must be reviewed and approved by the TA partner and hiring manager before use. AI-generated brief does not equal approved brief. Owner: TA Partner + Hiring Manager High risk	 DE&I funnel review Quarterly · analytics Required Quarterly review of pipeline representation data with HRBP and function heads. Agent surfaces the data — humans own the response. Owner: TA Director + HRBP High risk

Interview fraud is a live problem. VERITAS-LIVE is the answer.

VERITAS-LIVE

The interview
is happening.
Live analysis.

While Sarah's interview runs, VERITAS-LIVE transcribes responses in real time and surfaces authenticity indicators, AI-drafted response patterns, and evidence quality flags — giving the interviewer a quiet intelligence feed without disrupting the conversation.

**Live transcript**

Full response captured and analysed in real time across 3 questions

**AI-drafted response detection**

Flags responses with high likelihood of being pre-generated

**Evidence quality scoring**

Rates whether answers contain specific, measurable evidence

ASSESS · VERITAS-LIVE · REAL-TIME

VERITAS-LIVE — Sarah Johansson

Live authenticity analysis. Three questions shown. Navigate between responses with the buttons below.

97%

AUTHENTICITY

Strong

EVIDENCE

High

SPECIFICITY

Q1: WALK ME THROUGH THE LAST PRODUCT YOU SHIPPED WHERE THE DATA TOLD YOU SOMETHING UNEXPECTED.

At Spotify, we had been designing a playlist creation flow for about three months. The qualitative research was positive — users said they loved it. Then we shipped to 5% of users and the engagement data told a completely different story: 70% single-session abandonment. We dug into the sessions and realised the problem was the naming. We called it My Collection and users thought it was read-only. They did not understand they could create something. We changed the entry point to Start a playlist and the abandonment rate dropped to 28% within two weeks. The lesson was that naming is design. I had been so focused on the interaction that I had not stress-tested the language.

✓ Authentic: specific product, specific data, genuine failure narrative. Self-correction (naming is design) is a strong authenticity signal.

✓ Strong evidence: 70% to 28% abandonment with a clear causal change. Quantified and specific.

✓ No AI-generation patterns: includes a genuine mistake narrative, which AI-generated content rarely volunteers.

VERITAS-LIVE surfaces authenticity and evidence signals as the answer happens, flagged for human judgement, never an automated verdict.

Every source documented. Every candidate's data, exportable on request.

SUBJECT ACCESS REQUEST (SAR) — CANDIDATE AUDIT TRAIL

Under UK GDPR / Data Protection Act 2018, a candidate can submit a Subject Access Request (SAR) for all personal data held about them. This tool generates a complete audit trail export for a named candidate covering all hub activity.

🔗 What this covers

- All agent outputs referencing this candidate (screening scores, interview questions used, offer details, communications drafted)
- Pipeline stage history and timestamps
- QoH check-in records if applicable
- Any notes added in the candidate profile
- What data was sent to the Anthropic API (prompt summaries — no raw audio)

CANDIDATE NAME (EXACT)

e.g. Sarah Johansson

SAR REFERENCE NUMBER

e.g. SAR-2026-001

REQUESTING PARTY / REASON

e.g. Former applicant — Senior Product Designer role, applied March 2026

Generate SAR export

Generates a plain-text export of all hub data for this candidate

ⓘ Note: This hub stores data in your browser's localStorage only. It is not transmitted to any third party except the Anthropic API for AI processing. You must also check your ATS, email system, and calendar for additional candidate data not captured here. Anthropic's zero-retention API policy means no candidate data is stored on their servers.

Subject Access Request export: A full audit trail for any named candidate, in line with UK GDPR.

Every source documented. Every candidate’s data, exportable on request.

LIBRARY · AUDIT & COMPLIANCE

Data Sources & Intelligence Provenance

Every piece of market intelligence, compensation data, and external benchmark used in this hub is documented here. This page exists for auditing, compliance, and transparency — you should always know where your data comes from.

 **Auditing statement**

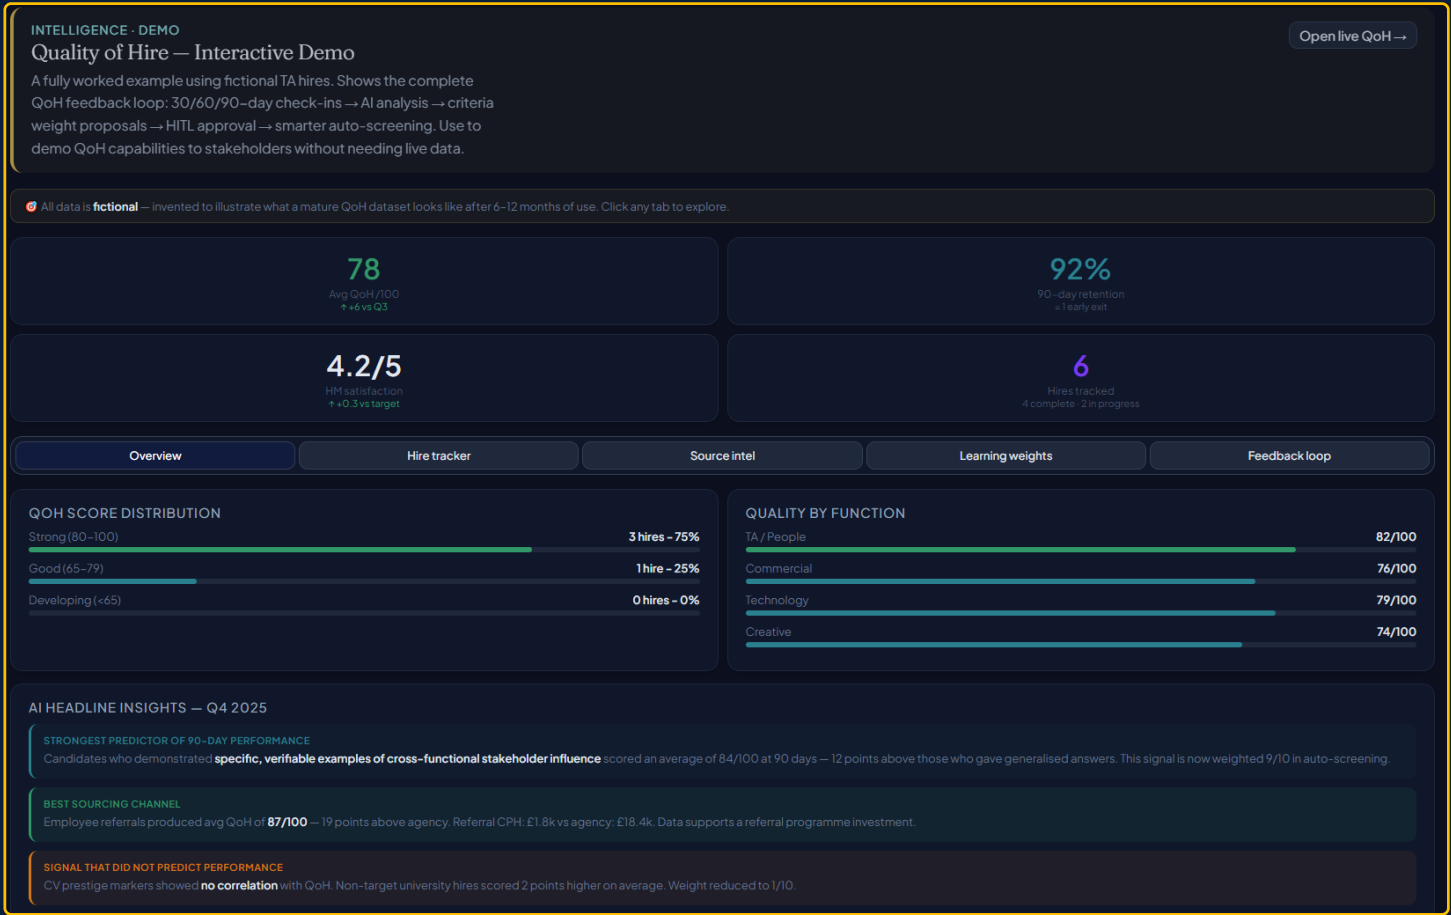
All compensation benchmarks, talent pool estimates, and competitor intelligence are synthesised from publicly available sources. No candidate personal data is transmitted to external services beyond the Anthropic API (for AI processing) and your configured ATS. This document should be reviewed quarterly and updated when data sources change.

COMPENSATION & SALARY BENCHMARKING

SOURCE	DATA TYPE	UPDATE FREQUENCY
Glassdoor glassdoor.com/Salaries	Reported salaries by role, location, and company. User-submitted data.	Continuous
LinkedIn Salary Insights linkedin.com/salary	Anonymised salary ranges from LinkedIn member profiles. Segmented by role, seniority, and geography.	Monthly
Levels.fyi levels.fyi	Verified compensation data for tech roles. Includes base, bonus, and equity breakdowns.	Continuous
Payscale payscale.com	Compensation benchmarks and salary surveys. Used for UK, EMEA, and US role benchmarking.	Quarterly
Robert Half Salary Guide roberthalf.co.uk/salary-guide	Annual salary guide for Finance, Tech, Creative, and HR roles. UK and EMEA focus.	Annual

Every compensation benchmark and market signal traces back to a named, public source.

The hub gets smarter with every hire.



Quality of Hire tracks 30/60/90-day outcomes, and feeds what it learns back into screening criteria.

The hub gets smarter with every hire.

AGENT 14 · COMPETITIVE INTELLIGENCE
Market Intelligence
Generates a sharp, exec-ready talent market briefing for any role or function. Salary benchmarks, competitor hiring moves, talent supply signals, and strategic recommendations.

INTELLIGENCE BRIEF INPUTS

FUNCTION OR ROLE FAMILY
e.g. Talent Acquisition, Engineering, Product Design

MARKET OR GEOGRAPHY
e.g. UK, EMEA, APAC, Global

SENIORITY LEVEL
Director / Senior Leader (£120k+) ▼

BRIEF TYPE
Full weekly briefing ▼

SPECIFIC QUESTIONS (optional)
e.g. Are competitors increasing TA headcount? What is the going rate for a Global TA Director in London?

TA function

Engineering

Product and Design

Executive brief

Generate intelligence brief

Market Intelligence keeps every hiring-manager conversation grounded in current salary and competitor data.

THE RESULTS

What changed?

Q1 2026 data, conservative estimates

145 min

saved per hire

down from 255 minutes of admin

£339k+

conservative annual saving

across agency cost, capacity & time-to-fill

47:1

return on investment

against Anthropic API cost

<4 wks

payback period

faster than average time-to-fill

Full breakdown: 40% agency cost reduction, FTE capacity unlocked, 5–8 day reduction in time-to-fill, and reduced mis-hire cost through structured screening.

WHY THIS IS DIFFERENT

Not a chatbot. Not a plugin. A complete operating system for TA.

 Bias-first architecture Every shortlist is automatically audited across six dimensions before the recruiter sees it — gender balance, name-based diversity, educational prestige bias, age signals, career gaps, and language bias. Not an add-on.	 Human-in-the-loop by design The AI does the work. The recruiter makes the decisions. Every automated step has a defined human checkpoint before any candidate advances, receives an offer, or is declined.	 Organisation-specific intelligence Every agent is configured around the client organisation's data — comp philosophy, EVP language, DE&I policy, competitor landscape, and role taxonomy. Organisation context is injected into every agent prompt automatically. Not a generic AI tool.
 Genuinely autonomous & connected Agents chain to each other automatically — CV screen flows to outreach, interview to debrief, debrief to offer modeller. Autonomous pipeline mode drafts outreach for every shortlisted candidate after approval. No copy-pasting between tools.	 Connected — not isolated Outlook email and calendar integration. ATS sync with Greenhouse, Lever, Workday. Stage changes push back to ATS automatically. Real systems, real data.	 Built to learn Quality of Hire data feeds back into screening criteria. Market Intelligence refreshes weekly. The hub gets more accurate with every placement.

THE NAME

It's literally named after me. That's not an accident.

WHAT ADONYE STANDS FOR

A

Align

Align hiring with business objectives from the first conversation — role brief, headcount rationale, success criteria, and stakeholder expectations set before sourcing begins.

D

Discover

Discover the right talent through intelligence-led sourcing — market benchmarks, competitor mapping, talent pool analysis, and diversity-first search strategies.

O

Optimise

Optimise every stage of the pipeline automatically — autonomous screening, bias-audited shortlists, interview intelligence, and scheduling without manual coordination.

N

Nurture

Nurture every candidate with a personalised, consistent experience — tailored outreach, timely communication, and silver medal pipelines for future opportunities.

Y

Yield

Yield measurable results — offers modelled at the right percentile, debrief-to-hire confidence scores, and quality-of-hire data that feeds back into future screening.

E

Elevate

Elevate the TA function from a transactional service to a strategic capability — freeing recruiters to focus on judgement, relationships, and the decisions that no algorithm should make.

Every agent in the hub is designed around one of these six principles.

The future of TA isn't less human.

It's more human because the 255 minutes of admin finally belong to the machine, and the judgement stays with us.

Thank You!

Adonye Orumbie - Adonye AI Agents Hub

Always happy to discuss this Hub further