FICTIONAL EXAMPLE | NOT A CLIENT REPORT

MARLOW KINCAID RECRUITMENT GROUP

AI Opportunities Report

PREPARED FOR	PREPARED BY	DATE
Marlow Kincaid Recruitment Group	Rubi	September 2026

CONTENTS

What's in this report

SECTION 01 — **Executive summary**
Headline findings, your three priorities and the decisions leadership needs to make

SECTION 02 — **Your AI vision and what value means**
The stated ambition, how value is measured, and where leadership and staff differ

SECTION 03 — **Where you are today: AI maturity**
Placement on the MIT CISR maturity model, the evidence, and what it takes to reach Stage 2

SECTION 04 — **Your people**
AI usage, people profile, training focus, pain points and fears

SECTION 05 — **Opportunities**
Overview | Prioritisation | Priority Workflow Cards | Next Steps

SECTION 06 — **90-day roadmap**
Actions, owners and what "done" looks like, plus who owns AI

SECTION 07 — **Conclusion**

APPENDIX A — **About this report**
Who was interviewed, anonymisation, assumptions, method and limitations

APPENDIX B — **Benchmarks**
Maturity, adoption and recruitment-sector data; EU AI Act timetable

APPENDIX C — **OPERA scoring detail**
Rubric and evidence for all 21 scored workflows

Short on time? Section 01 summarises the key findings and recommendations.

SECTION 01

Executive summary

All findings in this report are aggregated and anonymised. No statement is attributed to an identifiable individual. See Appendix A.

HEADLINE

AI is already part of daily work at Marlow Kincaid, but mostly on personal accounts, often with candidate data, and with no signed-off policy, while clients are told in writing that AI isn't used in screening. "AI-first" has also been widely heard as a message about job cuts. Fix the guardrails and the message first. Then the best returns are in work that assesses no one: the weekly contractor cycle and consultants' admin.

WHERE YOU ARE NOW

AI MATURITY

Mid Stage 1, "Experiment and Prepare"

On the MIT CISR Enterprise AI Maturity Model. A Copilot pilot exists (30 licences, London only, about a third used in a typical week), but there is no signed policy, no measured pilot and no clean data foundation yet (Section 3).

PEOPLE

40 people interviewed across all three offices

Three in four use AI at work, and most are open to it, but cautiously: the largest group (19 of 40) are "Cautious Observers", who will move once they see proof. The top concerns are the EU AI Act, job security and data (Section 4).

YOUR THREE PRIORITIES

1 Put guardrails in place, swap personal AI for approved tools, and say clearly what "AI-first" means for people.

WHY IT MATTERS clients now ask about AI in tenders, questionnaires and contracts, and some require notice or consent before AI is used in screening. The firm can't give a verified answer, because nobody knows which tools are in use: IT data shows personal ChatGPT traffic from around 60% of London consultant machines every week. And job security came up in 23 of 40 interviews, mostly unprompted, which puts retention, and with it conversion, at risk.

NEXT STEP within a month, sign off a short policy, a register of AI in use and one approved client answer. Roll out one approved assistant by role, with training, then block the rest. Give one honest message in all three offices.

2 Automate the weekly contractor cycle, from timesheet to cash.

WHY IT MATTERS contract work is over half of NFI. Of about 800 UK timesheets a week, around 130 are missing or wrong every Monday, and fixing them takes 70–80 person-hours, although nine in ten fixes follow rules. Errors flow on into about 290 credit notes a year, each restarting the client's payment clock. Nothing here assesses a candidate. We estimate automated checks and chasing could release roughly 30–45 hours a week, with people keeping every pay decision.

NEXT STEP a deep-dive with contractor services to write down the rules and set a baseline, then a pilot run alongside the current process for at least a month (never going live on a Monday or in December).

3 Give consultants and resourcers back selling time, starting with CV formatting.

WHY IT MATTERS consultants put admin at 40–60% of their week, desk NFI averages about £205k against a £250k ambition, and deals are lost on speed (about 4.5 days from brief to first CV). A CV takes about 20 minutes to format, and London resourcers alone spend 100+ hours a week on it. An approved tool with a quick human check could release roughly 70–100 of those hours (an estimate).

NEXT STEP a deep-dive on templates and checks, then a pilot in one or two London divisions, measured on minutes per CV and time to first CV.

WHY THESE PRIORITIES

We ranked the opportunities on business value, cost and effort, risk, readiness and fit, using OPERA (Rubi's framework for judging whether a workflow is a good candidate for AI, scoring operational pain, process repeatability, exposure to risk, readiness and appetite) and the prioritisation in Section 5. Of the 30 workflows we identified, we scored 21. The strongest are high-volume admin that assesses no one. AI screening and ranking of candidates would save the most time but carries the most risk: it is high-risk under the EU AI Act (obligations apply from December 2027), some client contracts restrict it, and candidate data isn't clean yet. Keep it off for now, and use the time to prepare.

DECISIONS LEADERSHIP NEEDS TO MAKE

- Agree what "AI-first" means for people, including how released time will be used (moved into selling, or absorbed as roles fall vacant), and say it in all three offices.
- Sign off the AI use policy now, as a first version to revise, and fund one approved assistant, allocated by role, not seniority.
- Set one position on candidate screening and the ATS match scores, with a review date.
- Name an executive sponsor, an AI lead and an owner for each pilot, and agree one or two measures before each starts.
- Bring the Amsterdam works council in early on anything that could monitor staff.

THE RECOMMENDED NEXT STEP

> Foundations in the first month, then two small, measured pilots, not a large programme: the contractor timesheet cycle and CV formatting, each preceded by a short workflow deep-dive (Section 6).

All findings are aggregated and anonymised. Figures are estimates based on interview evidence; see Appendix A for method and limitations.

SECTION 02

Your AI vision and what value means

HEADLINE

The stated ambition is “AI-first”. The plan leadership actually describes is “better with AI”: back office first, and people keep every decision about candidates, pay and clients. The label, not the plan, is what staff heard.

AI VISION

The leadership group has committed publicly to Marlow Kincaid becoming an AI-first business. It was meant as ambition: the firm that uses AI properly, rather than one dragged into it. The plan leaders describe in practice is more measured. Start in the back office, where nobody is being assessed by a machine. Prove real savings there. Keep candidate-facing AI switched off until the regulatory position is clear. People keep every decision about candidates, pay and clients.

That is a “better with AI” approach, in which AI does the checking, chasing, drafting and formatting. We think it is the right approach for a firm in a regulated, people-centred market. But the label and the plan don’t yet match, and the label is the one people heard. Leaders emphasise different outcomes. These fit together, but they haven’t yet been written into one plan.

WHAT VALUE MEANS TO YOU

Value is defined more clearly here than in most organisations we see. Leadership manages the business by:

- **Conversion** (operating profit over net fee income): about 14% today, with a board ambition of 22% within three years, on NFI of about £34m that isn't expected to grow much. Each point of conversion is worth about £340k of operating profit.
- **Productivity:** NFI per head of about £136k, against about £170k at the best-run firms of this size. Average desk NFI is about £205k, against an ambition of £250k.
- **Cash and close:** debtor days around 50, and a month-end close of eight working days against a five-day target.
- **Speed:** about 4.5 days from brief to first CV, where some clients now expect a shortlist within 48 hours.
- **Keeping good people,** particularly fee earners.

Staff costs are about 70% of NFI, so leadership sees capacity as the main lever, starting with consultants spending more of their week selling. How released capacity elsewhere is used is an open decision (Section 1). Any investment is expected to pay back within twelve months.

> **Gap:** the measures exist, but the baselines don't. Time-to-fill can't be measured reliably while around 40% of placements have no interview or offer date recorded. Productivity figures drawn from current ATS data should be treated with caution. We recommend agreeing one or two measures for each pilot, and taking a baseline before it starts (Section 5d).

ALIGNMENT CHECK

Leaders broadly agree on where to start. There are four gaps between what leadership intends and what the rest of the business experiences:

1. **The message and the plan.** "AI-first" at the January kick-off, followed by the line "doing more with the people we have", has been heard in all three offices as a signal of job cuts. No follow-up message has been given. Job security came up in 23 of 40 interviews, mostly without being asked, and staff have noticed vacancies not being refilled. This is a risk to the conversion goal itself. Several participants linked recent departures to the uncertainty, and the people who leave first are the ones with options.

2. **Candidate-facing AI.** Views differ sharply on switching on the ATS match scores and auto-shortlisting, which the firm already pays for. Sales-facing teams see them as a competitive necessity. Governance and risk roles, and others, want them kept off until the position is clear. They have been off since spring with no decision date, which several people described as a default rather than a decision.

3. **London and the other offices.** The Copilot pilot is London only. Participants in Dublin and Amsterdam described waiting for London and using free tools in the meantime. In the Netherlands, the works council must consent in advance to any system that could be used to monitor staff. That typically takes around four months, so a “pilot next month” timeline doesn’t fit.

4. **What clients are told and what is known.** Tenders, questionnaires and some contracts ask whether AI is used in screening. The answer given (no) is true of the firm’s approved tools. Nobody can currently verify it for the personal ones.

SECTION 03

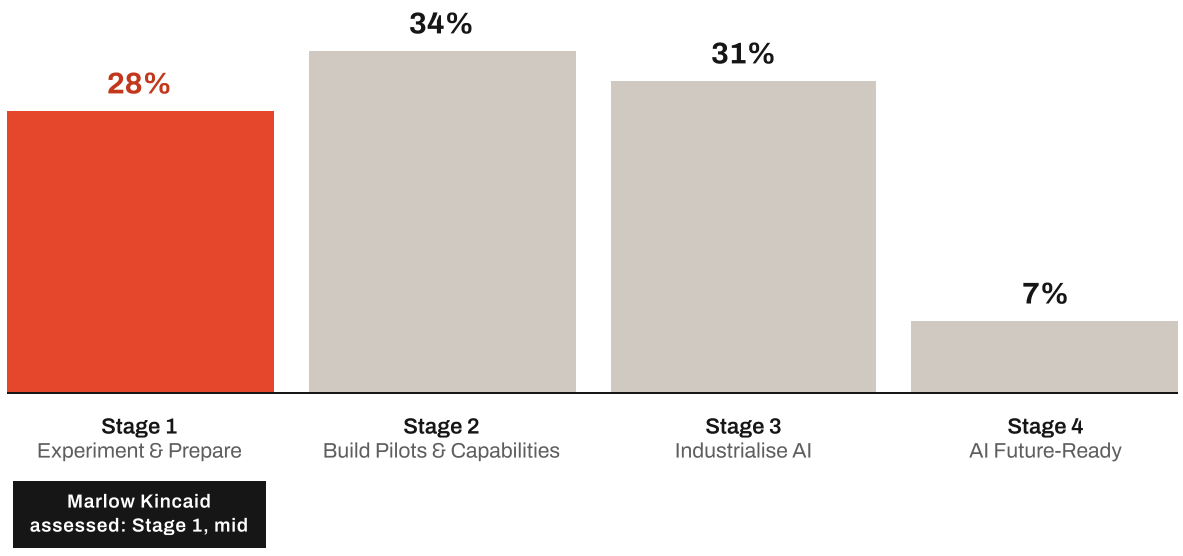
Where you are today: AI maturity

HEADLINE

Mid Stage 1, “Experiment and Prepare” (confidence: high). AI use is widespread, but it is individual and ungoverned.

Where enterprises sit on the MIT CISR Enterprise AI Maturity Model

% of responding enterprises



Source: MIT CISR, Building Enterprise AI Maturity (2024); 2022 Future Ready Survey, n = 721. Client placement: Rubi assessment, n = 40 interviews.

STAGE	ALL-INDUSTRY SHARE (MIT CISR, 2022, N = 721)	MARLOW KINCAID
1. Experiment and Prepare	28%	Mid Stage 1
2. Build Pilots and Capabilities	34%	
3. Industrialise AI	31%	
4. AI Future-Ready	7%	

This is Rubi's assessment against the MIT CISR criteria, not an official MIT CISR score. It is based on 40 interviews (n = 40).

EVIDENCE

- **Policy and governance:** an AI use policy has been in draft since spring and isn't signed off, partly because the regulatory picture kept moving. There is no register of AI tools in use and no approved answer for clients. A vendor switched on an AI feature by default (an engagement score in the video-interview platform), and it was switched off in 2025. Nobody is responsible for checking vendor releases for new AI features.
- **Tools and use:** a Copilot pilot of 30 licences, London only, allocated largely by seniority, with no training and no owner. About a third are used in a typical week, while 40–70 people are waiting for one (estimates varied). IT data shows personal ChatGPT traffic from around 60% of London consultant machines at least weekly (device traffic, not a headcount), and unapproved AI browser extensions on about 25 London machines, some able to read ATS pages.
- **Pilots and measurement:** no pilot has a named owner, a baseline or a measure. The ATS vendor's match scores and auto-shortlisting are paid for but switched off.
- **Data and architecture:** an ATS of around 400,000 candidate records, an estimated 30% of them duplicates. Two finance ledgers. A group timesheet platform that doesn't handle Dutch rules. Critical knowledge held in spreadsheets and in people's heads.
- **Culture and skills:** strong curiosity, especially in sales-facing roles, but no structured AI training. An optional lunch-and-learn drew few people.

Widespread daily use and a sanctioned pilot put you past the very start of Stage 1. What keeps you from Stage 2 is the absence of a policy, training, measured pilots and trusted data. That is why we place you mid-stage.

WHAT IT TAKES TO REACH STAGE 2

1. Sign off a first-version AI use policy, with a simple green, amber and red list, a register of AI in use and one approved answer for clients.
2. Replace personal tools with one approved assistant, allocated by role, with hands-on training. Then block the rest.
3. Run two small pilots, each with a named owner, a baseline and one or two measures (Section 5).
4. Give the ATS data an owner, and start fixing duplicates and retention at source.

WHY IT MATTERS

MIT CISR found that enterprises in Stages 1 and 2 performed below their industry average financially, and those in Stages 3 and 4 performed above it. The biggest financial gain came from moving from Stage 2 to Stage 3.

Further benchmarks, including recruitment-sector data and the EU AI Act timetable, are in Appendix B.

SECTION 04

Your people

HEADLINE

AI is already in daily use, and most people are open to it, but cautiously. The main concerns are regulation, job security and data. The biggest gap is permission: people don't know what they're allowed to do.

AI USAGE (SELF-REPORTED)

FREQUENCY OF AI USE AT WORK	PARTICIPANTS
Daily	7 of 40 (17.5%)
Weekly	5 of 40 (12.5%)
Regularly, frequency not stated	4 of 40 (10%)
Occasionally	14 of 40 (35%)
Never at work	10 of 40 (25%)

17 of 40 participants said they use personal or free AI tools for work, mostly ChatGPT and free translation tools. 13 hold a Copilot pilot licence. Use is heaviest in sales-facing and commercial roles, and lightest in finance and contractor services. Outside London there is almost no access to an approved assistant.

PEOPLE PROFILE

People profile: six types of people facing AI (n = 40 participants)



Share of participants. Overconfident Non-Adopters (2.5%) shown unlabelled; counts in the table below.

TYPE	COUNT	WHAT THIS GROUP NEEDS
1. Adaptable and Curious Open to AI and already exploring it on their own initiative.	13 of 40	Approved tools, room to experiment, advanced training, and help to focus
2. Willing Pivoters Ready to change how they work once shown how, and what's allowed.	3 of 40	Clear rules on what's allowed, structured hands-on learning and reassurance
3. Resistant and Static Prefer current ways of working and push back on change.	0 of 40	—
4. Overconfident Non-Adopters Feel they don't need AI, so rarely use it.	1 of 40	Real examples from similar roles, and a part in pilots
5. Cautious Observers Interested, but waiting for proof that it works and is safe.	19 of 40	Proof: parallel runs, low-risk first uses and measured results
6. Enthusiastic Early Adopters Eager, heavy users who move ahead of guidance.	4 of 40	Champion roles, with clear guardrails on data and checking
Unclassified Not enough evidence to assign a type.	0 of 40	—

Source: "The 6+1 Types of People Facing AI" (In AI We Trust, with Senaj Lelic).

WHAT THIS MEANS

There is almost no resistance to overcome: nobody we spoke to was resistant, and 17 of 40 are curious or enthusiastic. But nearly half are Cautious Observers, found at every level and in every office, who will move once they see proof. That shapes the approach:

- **Prove it before rolling it out.** Run changes in parallel with the current process, particularly anywhere near pay. People still remember a previous platform change that went badly.
- **Enthusiasm is running ahead of the guardrails.** The heaviest users are also where candidate data most often goes into personal tools, and where checking the output is least habitual. Champions need guardrails, not just freedom.
- **Fear is the blocker, not skill.** Some people avoid showing interest in AI in case it looks like training their own replacement. Tools launched without an explanation will be quietly ignored.
- **Offices outside London feel changes are done to them.** "With us, not to us" came up repeatedly.

> Nobody's told us not to. But **nobody's told us we can.**

PARTICIPANT – DIAGNOSTIC INTERVIEW

TRAINING FOCUS

Start with a short, mandatory AI basics module for everyone: what can and can't go into a tool, and how to check what comes back. This also supports the EU AI Act's AI-literacy expectations. Follow it with hands-on, role-based sessions built around the priority workflows. For consultants and resourcers, that means CV formatting and ads. For contractor services, it means the timesheet exception process, designed with the team. In Amsterdam, it means approved translation. Treat checking AI output as a skill in its own right: invented certifications and invented job benefits were among the errors people described. Brief managers first, so they can answer their teams' questions about AI and jobs, and add AI to the induction for new consultants.

MAIN PAIN POINTS

1. No approved tools or clear guidance on what's allowed (19 of 40).

MAIN FEARS

1. The EU AI Act: not knowing what's in scope, for whom, or when (24 of 40).

2. Consultant and resourcer admin (CV formatting, scheduling, notes, compliance packs) crowding out selling time (13 of 40).
3. Poor ATS data: duplicates, missing fields, and search that doesn't find people (11 of 40).
4. Client and tender questions about AI with no approved answer, and slow turnaround (11 of 40).
5. Chasing and re-keying in the weekly contractor cycle (10 of 40).

2. Job security, with "AI-first" heard as cuts (23 of 40).
3. Candidate and client data leaking through personal tools (18 of 40).
4. Confidently wrong output: invented facts, wrong numbers, outdated rules (14 of 40).
5. Unfair or impersonal treatment of candidates: bias, automated rejection, robotic messages (12 of 40).

Being monitored or scored by AI at work was raised in 8 interviews; in the Netherlands it is a works council matter.

Counts are the number of interviews in which a theme came up. 40 of around 250 people were interviewed, invited through leadership referrals. Colleagues who weren't interviewed may see things differently.

SECTION 05

Opportunities

HEADLINE

30 workflows identified and 21 scored. The best candidates are high-volume admin that assesses no one. Candidate screening would save the most time but carries the most risk, and should wait.

What is OPERA? OPERA is Rubi's framework for assessing whether a workflow is a good candidate for AI. It scores five things out of 20 each: Operational pain, Process repeatability, Exposure to risk (a higher score means safer to automate), Readiness and Appetite. Totals place each workflow in one of four bands. Full scoring detail is in Appendix C.

5A. OVERVIEW

Workflows by OPERA band (30 identified; 21 scored; n = 40 interviews)

0 80–100 Ready for full automation	**21** 50–79 Partial automation, human in the loop	**0** 30–49 Leave as-is for now
0 0–29 Needs redesign first	**4** Insufficient data to score	**5** Better fixed without AI

All 21 scored workflows fall in the "partial automation" band. That is typical of recruitment, where nearly every workflow involves people's data or a judgement somewhere, and it means a person stays in the loop throughout. Three sit on the boundary with "Leave as-is for now" (candidate screening, database search and Dutch temp-worker checks), because high value is offset by high risk. The band alone doesn't separate the opportunities, so the ranking in 5b also weighs value, complexity and time to result.

WORKFLOW, GROUPED BY FUNCTION	OPERA (CONFIDENCE)
CONTRACTOR OPERATIONS AND FINANCE	
Contractor timesheet exceptions and approval matching	70 (High)
Contractor onboarding: document and SDS chasing	70 (High)
Invoice PO checks and remittance matching	70 (Medium)
Month-end reconciliations (intercompany and bank)	65 (High)
Dutch temp-worker compliance and payroll checks	55 (High) · boundary with “Leave as-is”
CONSULTANTS AND RESOURCING	
CV formatting onto client templates	75 (High)
Interview scheduling and reminders	70 (High)
Dutch-language CV and job ad translation	70 (High)
Call notes and interview feedback write-ups	60 (High)
Life Sciences compliance packs and references	60 (High)
Prospect and market research	60 (Medium)
First-pass screening of applications	50 (High) · boundary with “Leave as-is”
Candidate database search and re-engagement	50 (High) · boundary with “Leave as-is”
MARKETING, BIDS AND CLIENT QUESTIONS	
Job ad drafting with a compliance check	75 (High) · close to “Ready”
Tender and bid responses	70 (Medium)
Client and tender questions about AI	65 (High)
LEGAL, DATA, IT AND REPORTING	
Management and client reporting packs	70 (High)

WORKFLOW, GROUPED BY FUNCTION	OPERA (CONFIDENCE)
Vendor contract and DPA first-pass review	70 (Medium)
Candidate database clean-up (duplicates and retention)	65 (High)
Subject access request (DSAR) search	60 (Medium)
Internal knowledge assistant (HR policies and how-to guides)	65 (Medium)

Not scored (insufficient data)

- **Employment contract paperwork across three jurisdictions.** Raised, with occasional template errors, but time, systems and appetite weren't discussed.
- **IR35 file notes and dispute chronologies.** Some volumes were given, but readiness, and appetite beyond pre-filling notes, weren't established.
- **Practical questions from workers and candidates** (payslips, relocation, registration), often multilingual and out of hours. No volumes, and no owner.
- **Checking the sales pipeline for forecasts.** Several hours a week spent verifying by phone. This is mainly a data quality problem; repeatability and readiness weren't discussed.

Better fixed without AI (quick wins using standard automation or a process change)

- Self-service password, MFA and ATS access resets, plus a joiner and leaver workflow driven from the HR system. Access requests are around 45% of about 180 IT tickets a week.
- A right-to-work question on application forms, routed to a review pile rather than rejecting automatically. One estimate puts it at 4–5 hours a week back per resourcer on UK roles. Legal should confirm the approach first.
- A structured job brief form for consultants: about a third of briefs arrive without a rate or IR35 status.
- An owner, and a mandatory, validated field, for POs and rates, so changes travel from consultant to platform to invoice.
- ATS interview reminders set up for Amsterdam, with Dutch templates.

5B. PRIORITISATION

WORKFLOW	OPERA (CONFIDENCE)	MAIN VALUE TYPE(S)	APPROX. VALUE (ESTIMATE)	COMPLEXITY	ADOPTION LIKELIHOOD	TIME TO FIRST RESULT (ESTIMATE)
1. Timesheet exceptions and approval matching	70 (High)	Capacity released; cost avoided; cash (fewer credit notes)	~30–45 hrs/week, from 70–80 hrs today, plus ~3 hrs/week of approval matching	Medium: platform integration, undocumented rules, limited IT capacity	Medium–High: the team wants the chasing gone, but is wary after a past system change; needs a parallel run	8–12 weeks
2. CV formatting onto client templates	75 (High)	Capacity redeployed into selling; fewer costly slips	~70–100 hrs/week in London resourcing; ~2–3 hrs/week per consultant who formats their own	Low–Medium: approved tool and templates now; ATS integration later	High: already done on personal accounts; must be as quick as ChatGPT	3–6 weeks
3. Client and tender questions about AI	65 (High)	Non-financial (protecting accounts); capacity released	2–6 hrs saved per questionnaire; turnaround in days rather than weeks	Low, once the policy and register exist	High: asked for across leadership, legal, bids and sales	4–6 weeks
4. Job ad drafting with a compliance check	75 (High)	Capacity released; non-financial (compliance, more applications)	~5–8 hrs/week of post-publication fixing, plus fewer non-compliant ads	Low: rules exist; ATS posting route needs closing	High in marketing; Medium for consultants ("three clicks more" won't be used)	3–6 weeks
5. Contractor onboarding: document and SDS chasing	70 (High)	Capacity released; cost avoided (late starts are lost billing days)	~10–15 hrs/week, from ~25 hrs of chasing, plus fewer late starts	Medium: status lives in three places	High: strongly wanted by those who chase	8–12 weeks (after item 1)
6. Invoice PO checks and remittance matching	70 (Medium)	Cash (debtor days); capacity released	Up to half of PO queries avoided (participant estimate), plus faster cash	Medium: needs PO ownership fixed first	High: keeps a person on anything client-facing	6–10 weeks

WORKFLOW	OPERA (CONFIDENCE)	MAIN VALUE TYPE(S)	APPROX. VALUE (ESTIMATE)	COMPLEXITY	ADOPTION LIKELIHOOD	TIME TO FIRST RESULT (ESTIMATE)
7. Interview scheduling and reminders	70 (High)	Capacity redeployed; non-financial (fewer no-shows)	~5–8 hrs/week in Amsterdam coordination; ~2–3 hrs/week per consultant who schedules	Medium: client calendars aren't visible	High, if candidates still deal with a person	6–10 weeks
8. Dutch-language CV and ad translation	70 (High)	Non-financial (removes a data risk); capacity released	~5–8 hrs/week from template work; the main gain is taking CVs out of free tools	Low: a paid tool is already licensed	High: the team would switch in a week	2–4 weeks

These are initial estimates from interview evidence, intended to decide what to investigate first. They are not a basis for investment on their own.

On money: we haven't converted hours into money. Released hours only create value if they go into fee-earning work (capacity redeployed) or avoid a cost, such as a role that isn't refilled (cost avoided). Which of these leadership intends is one of the decisions in Section 1, and it needs communicating. For scale, on current NFI every £34,000 of extra annual operating profit adds about 0.1 points of conversion. Salary costs and the value of an extra selling hour weren't provided, so both should be agreed in the deep-dives. The tools' own running costs also need counting.

NOT YET: AI SCREENING, RANKING AND MATCH SCORES

First-pass screening and database matching are the largest time sinks in resourcing, and competitors are marketing AI matching. But these workflows sit at the floor of the risk scale. Recruitment is a high-risk use under the EU AI Act, and those obligations now apply from 2 December 2027. Some client contracts already require notice or written consent. GDPR rules on automated decisions apply today. And the tools would run on candidate data that is heavily duplicated. We recommend keeping them switched off, and using the time before December 2027 to prepare: clean the candidate data, prepare an assessment template, review client clauses line by line, and map which roles have EU exposure. Then revisit with a date. In the meantime, the right-to-work form question above captures some of the time safely.

5C. PRIORITY WORKFLOW CARDS

Contractor timesheet exceptions and approval matching

OPERA 70 · Partial automation, with a human in the loop · Confidence: High

What it is: every week around 1,100 contractors submit timesheets (about 800 UK, 200 Netherlands and 100 Ireland). By the Monday 10am deadline, around 130 UK timesheets are missing or wrong. The team chases and fixes them before the Tuesday 2pm payroll cut-off, which participants put at 70–80 person-hours across the two days. Between a quarter and a third of client approvers approve by email or signed PDF, and each approval is matched to its timesheet by hand, which takes a morning every week. About 50 corrections a week rely on rules held in one or two people's heads.

The opportunity: check every timesheet as it arrives against the assignment, rate card, PO and approver. Flag the wrong ones before 10am, chase the missing ones through the channel each contractor actually responds to, and match email approvals automatically. People handle the one in ten that needs judgement and own every pay decision. This suits partial automation. Writing the rules down is part of the value in itself: it removes a dependency on individuals that has already led to overpayments.

Likely value: capacity released of roughly 30–45 hours a week (estimate). That assumes 40–60% of today's 70–80 hours, plus most of the approval matching. There should be fewer errors downstream. Participants reported 20–25 overpayments a quarter, and many of the ~290 credit notes a year trace back to this step. Each credit note takes about 45 minutes of finance time and restarts the client's payment clock. The released time could go into debt chasing and data quality, or be absorbed as roles fall vacant. That is a leadership decision.

Key risks: medium. The workflow involves contractor pay and personal data, and a wrong pay outcome hurts people directly, so nothing changes pay without a person. No candidate is assessed. The team is wary after a previous platform change went badly, so any change should run in parallel first and never go live on a Monday or in December.

What we'd need to confirm: the full set of rules and who holds them; how the timesheet platform, ATS and email can be connected, and by whom, given IT capacity; the real share of email approvals; baseline hours and error counts; and whether the Dutch and Irish timesheets follow later.

CV formatting onto client templates

OPERA 75 · Partial automation, with a human in the loop · Confidence: High

What it is: before a CV goes to a client, it is moved onto the Marlow Kincaid template. Contact details are removed, the layout is fixed and sometimes a summary is added, across around 15 variants by division and client. It takes about 20 minutes a CV. London resourcers each do about 25 a week ("a day a week, each"), and consultants 10–15 a week. The ATS parser doesn't handle it well, and many consultants already do it in personal ChatGPT accounts, often with candidate details included.

The opportunity: an approved tool that puts the CV on the right template for the division and client, strips the contact details and drafts the summary. A person then checks it in a couple of minutes before it goes. This suits partial automation, and it removes the single biggest reason for pasting CVs into personal accounts.

Likely value: roughly 70–100 hours a week released across London resourcing alone (estimate), from about 115 hours today down to a 2–8 minute check per CV. Add 2–3 hours a week for each consultant who formats their own. The value is mainly capacity redeployed: faster first CVs (currently about 4.5 days from brief) and more time with candidates. Fewer slips too: contact details left on a CV have cost at least one fee this year.

Key risks: medium-low. It involves candidate personal data, so the tool needs a proper contract and data processing terms, with EU data residency for Dublin and Amsterdam work. Formatting doesn't assess the candidate, but AI can invent details (an invented certification was one example described), so a person reads every CV before it's sent. Some pharma contracts restrict automated tools in selection, so confirm how formatting is treated before extending it to those accounts.

What we'd need to confirm: the template variants; whether the tool can sit inside the ATS (people said anything slower than ChatGPT on their phone won't be used); baseline minutes per CV and time to first CV; and how removal of contact details is checked.

Client and tender questions about AI

OPERA 65 · Partial automation, with a human in the loop · Confidence: High

What it is: clients now ask about AI in tenders, supplier questionnaires and contract clauses. They ask whether AI is used in screening or ranking, whether candidates are told, what human oversight exists, and whether the firm is ready for the EU AI Act. Nine of the last 14 tenders had an AI question. There have been about 30 Dutch client AI clauses or questionnaires this year, against 4–5 last year, and 7 pharma questionnaires with AI sections in one quarter in Ireland. Prospects raise it in first meetings too. Each answer is routed for legal review, and turnaround has sometimes slipped to weeks.

The opportunity: first, the non-AI foundation: a register of AI in use and one approved, true answer, signed off by Legal. Then an answer library that drafts a first response to each questionnaire from approved wording and past submissions, with sources, for a person to check and sign. This suits partial automation.

Likely value: mainly non-financial. It protects key accounts and renewals (at least one framework renewal is waiting on an answer) and shows the firm to be a safe pair of hands. It also releases 2–6 hours per questionnaire (estimate) across legal, data protection, bids and office leaders, with turnaround in days rather than weeks.

Key risks: medium. These are written, and sometimes contractual, statements to clients. The answer is only as true as the register behind it: while personal tools are in use, the firm can't honestly say "no AI anywhere". Legal signs off the approved wording and any departure from it.

What we'd need to confirm: which clients have contractual AI clauses (they haven't yet been reviewed line by line against the tools in use); the full register; who owns the answer library; and how often the wording needs refreshing as EU AI Act guidance develops.

5D. NEXT STEP FOR EACH OPPORTUNITY: A WORKFLOW DEEP-DIVE

Before building any of these, the person who owns or does each workflow should complete a second, more in-depth Rubi call. That call:

- establishes a proper baseline (current time, volume and error rates);
- validates the value, including how released time would be used;

- maps the tools and systems involved, integrations, security and data risks, and the training needed.

For the contractor cycle, it should include the people who run it every week. For CV formatting, it should include a resourcer and a consultant from each pilot division. Every workflow then goes through the same repeatable process: value, risk level, ownership and governance. The approach used on the first workflow can then be reused on the next.

SECTION 06

90-day roadmap

HEADLINE

Foundations in the first month, two small pilots in months two and three, and a decision at day 90.

TIMING	ACTION	OWNER (ROLE)	WHAT "DONE" LOOKS LIKE
Weeks 1–2	Agree what "AI-first" means for people, including how released time will be used, and say it in all three offices. Brief managers first.	Chief executive, with the people lead and office leaders	One message delivered everywhere; managers can answer their teams' questions
Weeks 1–4	Sign off a first-version AI use policy (green, amber and red list), start the AI register, and agree one approved answer for clients and tenders	Legal and data protection, with an executive sponsor	Policy published and read; register v1 complete; approved answer in use
Weeks 2–6	Reallocate and extend one approved assistant by role, with training and data-loss controls (Amsterdam after the works council process; paid EU translation extended there if Legal confirms). Then block unapproved tools	IT lead, with the learning and development lead	Licences held by the roles that need them; usage measured at 30 days; personal-account traffic falling
Weeks 1–4	Open the conversation with the Amsterdam works council on the assistant and anything that could monitor staff	Netherlands office lead, with legal	Request submitted, setting out purpose, data and safeguards
Weeks 2–5	Workflow deep-dives for the timesheet cycle and CV formatting: rules, baselines and measures	Operations sponsor and revenue sponsor, with the people who do the work	Rules written down; baseline hours, errors and time to first CV recorded
Weeks 5–12	Pilot 1: timesheet exception checks and chasing, run in parallel with the current process for at least a month	Pilot owner in contractor services	Exceptions flagged before 10am; hours and credit notes compared with the baseline
Weeks 5–10	Pilot 2: CV formatting with the approved tool in one or two London divisions	Pilot owner in resourcing	Minutes per CV and time to first CV against the baseline; every CV checked
Week 12	Review both pilots (keep, change or stop), choose the next workflows, and set the position and date for candidate screening	Leadership group	Decisions recorded; next deep-dives booked

Owners are roles, to be confirmed by leadership. Avoid contractor-services changes on Mondays, Tuesdays and in December: from an October start, the timesheet parallel run should end by mid-December or start in the new year. Limited IT capacity means outside implementation help is likely.

WHO OWNS AI

Today, nobody does. We recommend an executive sponsor who answers to the board for AI, and a named AI lead who owns adoption and training day to day. Each approved tool needs an owner for its settings, usage and vendor updates, including checking whether a release has switched on a new AI feature. New tools follow one route: a request, an IT security and data protection check, then the register and the approved list. In the Netherlands, anything that could monitor staff goes to the works council first.

SECTION 07

Conclusion

Marlow Kincaid starts with real advantages. People already use AI, leaders agree on where to start, and value is measured in terms the business already manages by. What's missing is permission and proof. That means a policy people can follow, an approved tool as easy to use as the one on their phone, an honest message about jobs, and two pilots with numbers attached. Moving from Stage 1 to Stage 2 within three months is realistic. The single most important thing to do first is to sign off the policy and swap personal AI for an approved tool, because every other opportunity, and every answer you give a client, depends on it.

Rubi's proposal for supporting this work is available separately on request, so the recommendations in this report can be assessed on their own merits.

DISCLAIMER

This report is based on a limited number of self-reported interviews and is provided for information and discussion only. It does not constitute legal, financial, regulatory or technical advice, including on the EU AI Act, data protection or employment law. Estimates of time, value and cost are indicative and should be validated before any decision is made. The organisation remains responsible for its own decisions, including the selection, implementation and use of any AI tools. To the extent permitted by law, Rubi accepts no liability for any loss arising from reliance on this report.

APPENDIX A

About this report

WHO WAS INTERVIEWED

- **40 participants (n = 40):** 6 in the leadership group, 13 managers, 8 specialists and 13 frontline staff (consultants, resourcers and support roles). 24 are based in London, 8 in Dublin and 8 in Amsterdam.
- **Dates:** 7–18 September 2026.
- **Method:** structured diagnostic calls of about 20 minutes each, run through Rubi's Conversation Engine using role-adapted question sets (leadership, managers, specialists and frontline). Participants were invited through leadership referrals, and a few calls were shorter than planned.

ANONYMISATION

All findings are aggregated and anonymised, and no statement is attributed to an identifiable individual. We don't report findings by office and role together. Where a function has fewer than three participants, it is described as part of a wider group. Several incidents described to us are reported only in general terms, and quotes are few and stripped of identifying detail. Names, clients and specific people mentioned in the interviews have been removed.

ASSUMPTIONS ABOUT THE ORGANISATION

AREA	ASSUMPTION
Sector	Specialist recruitment and staffing: permanent, contract, executive search and outsourced recruitment (RPO and MSP), plus Dutch temporary agency work. Offices in London (head office), Dublin and Amsterdam.
Size	About 250 staff, including about 140 fee earners; around 1,100 contractors and temporary workers on assignment each week; NFI of about £34m. The “Who owns AI” section is included.
Commercial drivers	(1) Conversion (operating profit over NFI), from about 14% to 22% in three years; (2) NFI per head and per desk; (3) cash (debtor days) and a faster close; (4) keeping fee earners, and speed to market.
Productive time	Selling time: conversations with clients and candidates, and business development.
Sensitive data	Candidate CVs (sometimes including health or other special-category data), identity and right-to-work documents, contractor pay and bank details, client-confidential briefs and organisation charts, and employee records.
Regulation and constraints	EU AI Act (recruitment is high-risk; obligations from December 2027); UK GDPR and GDPR, with three data protection regulators; the Dutch works council’s consent rights; the Dutch temporary-work collective agreement; UK right-to-work penalties; IR35 off-payroll rules; client contracts requiring notice or consent before AI is used in screening; equality law in job ads.
Risk profile used for OPERA	Anything that assesses, ranks or filters candidates, or could monitor staff, is the highest risk (0–5). Work that processes candidate or pay data without assessing anyone is medium (6–15, depending on controls). Internal and company data is lower risk.

METHOD AND LIMITATIONS

- Findings come from self-reported interviews, not system data, measurement or audit.
- 40 of around 250 people were interviewed (about 16%), invited through leadership referrals, which may skew the sample.
- OPERA scores are a structured synthesis of that evidence. They are indicative, not precise, and are meant to decide what to investigate first, not to justify investment on their own.

- Time, volume and financial figures are participants' estimates and carry the confidence rating shown. Some didn't agree (for example, the share of email approvals ranged from a quarter to a third, and the Copilot waiting list from 40 to 70), so we use ranges. All need validating before investment decisions.
- Theme counts are the number of interviews in which a theme came up, based on our reading of each transcript.
- People chose what to tell an external consultant. Some views, particularly about risk and their own AI use, may be under-reported.
- Where the report links signals to each other (for example, uncertainty and departures), these are associations, not established causes.
- Statements about the EU AI Act and other regulation describe the position as we understand it at the time of writing. They are not legal advice.

APPENDIX B

Benchmarks

Most figures below measure adoption (whether AI is used), not maturity stage. We found no credible published data on AI maturity stage specifically for recruitment firms. The recruitment-sector surveys below are published by software vendors, so treat them as indicative.

MIT CISR ENTERPRISE AI MATURITY MODEL

STAGE	2022 SHARE (N = 721)	2025 SHARE (N = 152)	2025 GROWTH VS. INDUSTRY	2025 PROFIT VS. INDUSTRY
1. Experiment and Prepare	28%	13%	−26.5 pp	−15.1 pp
2. Build Pilots and Capabilities	34%	23%	−6.8 pp	−1.4 pp
3. Develop AI Ways of Working / Industrialise AI	31%	46%	+4.7 pp	+0.8 pp
4. AI Future-Ready	7%	18%	+13.9 pp	+9.9 pp

Sources: MIT CISR, Building Enterprise AI Maturity, 2024 (2022 Future Ready Survey); MIT CISR, Grow Enterprise AI Maturity for Bottom-Line Impact, 2025 (2025 survey). The 2025 sample is much smaller. Growth and profit are percentage points above or below the industry average. The 2025 data suggests many organisations have moved on since 2022, so a Stage 1 position is now further behind the field than the 2022 chart suggests.

BROAD ADOPTION

- **McKinsey, The State of AI, 2025:** around 88% of organisations use AI regularly in at least one function; about a third have begun scaling it; around 6% qualify as “AI high performers”.

- **McKinsey, The state of AI in 2026: On the road to ROI, 2026** (1,719 respondents, surveyed May–June 2026): nearly nine in ten organisations use AI regularly in at least one function, and 44% report scaling it across the enterprise. 37% attribute some profit (EBIT) impact to AI, and about 6% remain high performers. 80% of respondents say AI has improved their own productivity. The gap between individual productivity and business value is the one this report aims to close.

RECRUITMENT AND STAFFING SECTOR

- **Bullhorn, GRID 2026 Industry Trends Report, 2026** (about 2,300 recruitment professionals worldwide, including the UK, Ireland and Benelux): only 10% of firms have AI built in across their whole workflow, and fewer than half use AI for any single recruitment function. Among firms using AI screening, 46% say it at least halved screening time.
- **The Access Group, AI in Recruitment 2026 Outlook, 2026** (159 agency leaders in the UK, Australia and New Zealand): 78% of agencies use AI in some form, but only 6% have made it part of their core operating model. 47% are running pilots without a clear strategy or measurement. About 70% use AI to write job adverts.

Reading across: Marlow Kincaid's everyday individual use is in line with the sector. Like most agencies, what's missing is structure: a policy, approved tools, training and measured pilots.

UNAPPROVED AI AT WORK

- **Microsoft UK, survey by Censuswide, 2025** (2,000+ UK employees): 71% have used AI tools their employer has not approved.
- **KPMG and the University of Melbourne, Trust, attitudes and use of AI: A global study 2025, 2025** (48,000+ people in 47 countries): almost half of employees who use AI at work admit using it against policy or uploading sensitive company information to public tools, and 66% rely on AI output without checking it. Only 40% say their workplace has a generative AI policy.

EU AI ACT TIMETABLE (AS AT SEPTEMBER 2026)

OBLIGATION	APPLIES FROM
Prohibited practices, including emotion recognition in the workplace	2 February 2025
AI literacy for staff who use AI (Article 4)	2 February 2025
Transparency obligations (Article 50)	2 August 2026
High-risk obligations for AI in employment and recruitment (Annex III, point 4), such as screening, ranking and filtering candidates	**2 December 2027** (previously 2 August 2026)

Source: Regulation (EU) 2026/1744 (the "Digital Omnibus" on AI), published July 2026, as summarised by the European Commission and legal commentators. The rules apply to the Irish and Dutch businesses. They may also reach UK work where a system's output is used in the EU, for example shortlists for EU clients. This needs confirming with legal advisers. The delay gives time to prepare; it doesn't remove the obligations, the client contract clauses or today's GDPR rules on automated decisions.

NATIONAL BUSINESS ADOPTION

- **ONS, Artificial intelligence in UK businesses: 2023 to 2026, 2026:** about 35% of UK businesses with 10 or more employees use at least one AI technology, up from about 12% in late 2023, and 49% of businesses with 250 or more employees.
- **CSO, Information Society Statistics – Enterprises 2025, 2025:** more than 20% of Irish enterprises with 10 or more employees used AI, and 58% of large enterprises.
- **CBS (Statistics Netherlands), 2025:** 22.7% of Dutch enterprises with 10 or more employees used AI in 2024, and 59.2% of those with 500 or more.
- **Eurostat, 2025:** 20.0% of EU enterprises with 10 or more employees used AI in 2025, up from 13.5% in 2024.

APPENDIX C

OPERA scoring detail

RUBRIC

Each dimension is scored out of 20. Scores are given in bands, and all five dimensions are weighted equally. Totals are rounded to the nearest 5, and scores within 10 points of each other should be treated as equivalent. Where a total sits within 5 points of a band boundary, both bands are named.

SCORE	OPERATIONAL PAIN	PROCESS REPEATABILITY	EXPOSURE TO RISK (HIGHER = SAFER TO AUTOMATE)	READINESS	APPETITE
16–20	Named as a top frustration by several participants, with time or cost quantified	Same steps, inputs and output every time	No personal, confidential or regulated data; errors cheap and visible	Data accessible and clean; system has an API or export; owner identified	Several participants actively asked for it or have tried to solve it
11–15	Raised unprompted by more than one person as a real drag	Mostly consistent, with some judgement or exceptions	Some sensitive data, manageable with human review	Data exists but needs cleaning or consolidation; integration possible with effort	Clear willingness, with some conditions
6–10	Mentioned once, or an annoyance rather than a problem	Varies by person, client or case; significant judgement	Sensitive data or consequential decisions; needs sign-off	Data fragmented or held in personal tools; no clear owner	Mixed views, or passive willingness
0–5	Not described as painful	Different every time	Regulated data, decisions affecting outsiders, contractual restriction, or works council approval required	Data unavailable or locked in	Resistance, or the people who do the work weren't asked

TOTAL	BAND
80–100	Ready for full automation
50–79	Partial automation, with a human in the loop
30–49	Leave as-is for now
0–29	Needs redesign first

Confidence ratings: High means all five dimensions were evidenced by two or more participants; Medium means three or four were well evidenced; Low means two or fewer.

WORKFLOW SCORES

1. Contractor timesheet exceptions and approval matching — 70 / 100

DIMENSION	SCORE	EVIDENCE
Operational pain	18	Around 130 of ~800 UK timesheets missing or wrong at 10am every Monday; 70–80 person-hours across Monday and Tuesday. Raised by leadership, contractor services and finance.
Process repeatability	15	About nine in ten fixes follow rules, but the rules aren't written down; one in ten needs a judgement call.
Exposure to risk	12	Contractor pay and personal data; pay errors hurt people directly. Nobody is assessed, and a person keeps every pay decision.
Readiness	10	Timesheet platform, ATS and email don't connect cleanly; rate and PO changes don't always sync; IT capacity is limited.
Appetite	16	Top of the list for operations and finance leaders; the team welcomes losing the chasing, on condition of a parallel run.
Total	**70** (71 unrounded)	Partial automation, with a human in the loop.
Confidence	**High**	All five dimensions evidenced by several participants.

2. Contractor onboarding: document and SDS chasing

70 / 100

DIMENSION	SCORE	EVIDENCE
Operational pain	**17**	About 150 new UK engagements a month; ~60 files a week, 40% incomplete, 12–15 hours a week chasing; ~25 a month start without an SDS (~12 hours a week chasing); four late starts in one month.
Process repeatability	**15**	Reminders and status tracking are mechanical; the document checks themselves are judgement.
Exposure to risk	**12**	Identity, right-to-work and IR35 documents. Automation would only chase and track; every check stays with a person.
Readiness	**9**	Documents in the ATS, status in a shared spreadsheet, SDSs in email; three SDSs logged against the wrong contract last quarter.
Appetite	**16**	Strongly wanted by the people who do the chasing.
Total	**70** (69 unrounded)	Partial automation, with a human in the loop.
Confidence	**High**	All five dimensions evidenced by two or more participants.

3. Invoice PO checks and remittance matching

70 / 100

DIMENSION	SCORE	EVIDENCE
Operational pain	**16**	About 40% of euro invoices are queried, nine in ten over the PO, taking a substantial share of credit-control time; about 290 credit notes a year from timesheet errors.
Process repeatability	**16**	Checking the invoice PO against the client's live PO is a rule; remittance matching is mechanical apart from short payments.
Exposure to risk	**15**	Client commercial data, no candidate data; errors are visible; a person sees anything that goes to a client.
Readiness	**8**	PO field is free text with no owner; the correct PO sits in emails and client portals; separate ledgers.
Appetite	**14**	Finance roles would welcome it and prioritise cash, provided no email goes to a client without a person seeing it.
Total	**70** (69 unrounded)	Partial automation, with a human in the loop.
Confidence	**Medium**	Pain and appetite well evidenced; risk and readiness mainly from one or two participants.

4. Month-end reconciliations (intercompany and bank) 65 / 100

DIMENSION	SCORE	EVIDENCE
Operational pain	14	Close takes eight working days against a five-day target; intercompany about 2.5 days a month; bank reconciliations 1.5–2 days, a third auto-matched.
Process repeatability	16	Participants put intercompany at ~80% and bank reconciliations at ~90% mechanical; judgement sits in the exceptions.
Exposure to risk	13	Confidential financial data and audit evidence; nothing posts to the ledger without approval; no personal data.
Readiness	7	Two ledgers with no common chart-of-accounts mapping; a bank export format changed; heavy spreadsheet use.
Appetite	13	Wanted, but only after a parallel run over two or three closes with matches that can be explained.
Total	**65** (63 unrounded)	Partial automation, with a human in the loop.
Confidence	**High**	All five dimensions evidenced by two or more participants.

5. Dutch temp-worker compliance and payroll checks

55 / 100

DIMENSION	SCORE	EVIDENCE
Operational pain	**16**	About 25 hours a week of compliance work across roles; ~25 hours entries wrong each week at 15–20 minutes; phase-tracking errors have led to back-pay.
Process repeatability	**12**	Rules-based (CAO phases, equal pay, minimum wage), but exceptions are frequent.
Exposure to risk	**8**	Pay, pension and CAO compliance with inspection risk; works council may need informing. Flag only, never decide.
Readiness	**7**	The group timesheet platform doesn't handle Dutch rules; client pay scales in a spreadsheet; the rules are known by very few people.
Appetite	**12**	A top wish in the Netherlands business, strictly as "flag, not decide"; little room for change before the January rule changes.
Total	**55** (55 unrounded)	Partial automation, with a human in the loop. On the boundary with "Leave as-is for now".
Confidence	**High**	All five dimensions evidenced by two or more participants.

6. CV formatting onto client templates

75 / 100

DIMENSION	SCORE	EVIDENCE
Operational pain	**18**	About 20 minutes a CV; ~25 a week per London resourcer (“a day a week, each”); 10–15 a week per consultant. Raised by leaders, managers and frontline staff.
Process repeatability	**14**	Same steps every time, but around 15 template variants by division and client.
Exposure to risk	**11**	Candidate personal data going to clients. Not an assessment, but contact-detail slips and invented details need a human check; some client contract terms.
Readiness	**12**	Templates exist; the ATS parser handles it poorly; an approved assistant could do it now, with ATS integration later.
Appetite	**18**	The most-requested fix, already done on personal accounts.
Total	**75** (73 unrounded)	Partial automation, with a human in the loop.
Confidence	**High**	All five dimensions evidenced by several participants.

7. Interview scheduling and reminders

70 / 100

DIMENSION	SCORE	EVIDENCE
Operational pain	**16**	About 120 interviews a week at three hours a day in one office; ~6 hours a week per consultant elsewhere; 30–50% rescheduled; ~8% no-shows.
Process repeatability	**12**	The mechanics repeat, but exceptions (site managers, PAs, shift workers, time zones) are constant.
Exposure to risk	**14**	Contact and calendar data; decides nothing about anyone; messages to candidates need a human tone.
Readiness	**10**	Outlook and WhatsApp; client calendars aren't visible; ATS reminders exist but aren't set up in Dutch.
Appetite	**16**	Strongly wanted by coordinators and consultants, provided candidates don't feel they're dealing with a robot.
Total	**70** (68 unrounded)	Partial automation, with a human in the loop.
Confidence	**High**	All five dimensions evidenced by two or more participants.

8. Dutch-language CV and job ad translation

70 / 100

DIMENSION	SCORE	EVIDENCE
Operational pain	**14**	About 25 CVs a week at 30–40 minutes in one team (~15 hours); 5+ hours a week for other consultants; 3–4 hours a week on bilingual ads.
Process repeatability	**13**	Language is repeatable; matching job-title levels is judgement and has gone wrong before.
Exposure to risk	**12**	Full CVs currently go into free tools. An EU-hosted, contracted tool with a human check makes this manageable.
Readiness	**16**	A paid translation tool is already licensed in part of the business and could be extended.
Appetite	**17**	Participants said the whole team would switch “the same week” to an approved EU tool.
Total	**70** (72 unrounded)	Partial automation, with a human in the loop.
Confidence	**High**	All five dimensions evidenced by two or more participants.

9. Call notes and interview feedback write-ups

60 / 100

DIMENSION	SCORE	EVIDENCE
Operational pain	**16**	About 25 calls a week and 6+ hours of write-ups; notes typed at night from memory; weak handovers when consultants leave.
Process repeatability	**15**	Mostly transcribing and structuring what was said.
Exposure to risk	**5**	Recording candidates and referees needs consent; Dutch works council consent required; possible client clauses. Unapproved note-taking bots have been found joining calls.
Readiness	**9**	Calls happen on mobiles; logging a call takes many clicks; ATS integration unknown.
Appetite	**15**	A top request for several consultants; a question about approved note-takers has gone unanswered for months.
Total	**60** (60 unrounded)	Partial automation, with a human in the loop.
Confidence	**High**	All five dimensions evidenced by two or more participants.

10. Life Sciences compliance packs and references

60 / 100

DIMENSION	SCORE	EVIDENCE
Operational pain	**17**	About 90 packs a month at ~3 hours each, two-thirds formatting and re-keying; references take ~45 minutes each; a placement was lost to a slow reference.
Process repeatability	**14**	The same data goes onto different client forms; the reference call itself is judgement.
Exposure to risk	**5**	Candidate and referee data; some pharma contracts need written consent for automated tools in screening or selection; EU data residency expected.
Readiness	**10**	Client templates vary; the data sits in CVs, the ATS and email.
Appetite	**16**	Strongly wanted, provided it never assesses the candidate and data stays in the EU.
Total	**60** (62 unrounded)	Partial automation, with a human in the loop.
Confidence	**High**	All five dimensions evidenced by two or more participants.

11. Prospect and market research

60 / 100

DIMENSION	SCORE	EVIDENCE
Operational pain	12	About 80 outreach emails a week at 20–30 minutes each if done properly; three to four days to map a market for a search.
Process repeatability	13	Research steps repeat; choosing who matters and writing are judgement.
Exposure to risk	10	Public data about named people; some search clients restrict undisclosed AI in the process.
Readiness	12	Public sources and existing tools; company records in the ATS are messy.
Appetite	13	Keen in some sales roles; conditional on accuracy in others.
Total	**60** (60 unrounded)	Partial automation, with a human in the loop.
Confidence	**Medium**	Two participants; readiness and risk mainly from one.

12. First-pass screening of applications

50 / 100

DIMENSION	SCORE	EVIDENCE
Operational pain	18	150–310 applications an ad; 12–15 hours a week per resourcer; 60–70% fail basic criteria; around 40% of one outsourced team's time.
Process repeatability	14	Knock-out criteria are mechanical; the second pass is judgement.
Exposure to risk	3	High-risk under the EU AI Act; client contracts require notice or consent; GDPR rules on automated decisions; bias risk; works council.
Readiness	9	Vendor features exist but are switched off, and would run on duplicated data.
Appetite	8	Views are sharply split across the business.
Total	**50** (52 unrounded)	Partial automation, with a human in the loop. On the boundary with "Leave as-is for now".
Confidence	**High**	All five dimensions evidenced by several participants.

13. Candidate database search and re-engagement

50 / 100

DIMENSION	SCORE	EVIDENCE
Operational pain	**15**	About 400,000 records rarely searched; paying for InMails to reach known candidates; a fee lost when a known candidate couldn't be found.
Process repeatability	**12**	Search is repeatable; deciding fit is judgement.
Exposure to risk	**5**	Where "finding" becomes "ranking" under the EU AI Act is unclear; retention rules apply to old records.
Readiness	**5**	Duplicates, stale job titles, no skills tagging and no search across Dutch and English.
Appetite	**15**	Wanted by resourcers and consultants.
Total	**50** (52 unrounded)	Partial automation, with a human in the loop. On the boundary with "Leave as-is for now".
Confidence	**High**	All five dimensions evidenced by two or more participants.

14. Job ad drafting with a compliance check

75 / 100

DIMENSION	SCORE	EVIDENCE
Operational pain	**15**	About 900 ads a month; ~8 hours a week fixing live ads. A June audit of 100 live ads found 38 without the agreed salary format and 11 with wording that carried a discrimination risk.
Process repeatability	**16**	Standard structure, tone of voice and rules.
Exposure to risk	**15**	No candidate data; ads are public, so a person approves them; an inclusive-language check lowers the risk.
Readiness	**13**	Tone and salary rules exist; ATS multi-posting lets ads bypass checks.
Appetite	**17**	Wanted by marketing and consultants, and seen by leadership as a safe place to start.
Total	**75** (76 unrounded)	Partial automation, with a human in the loop. On the boundary with “Ready for full automation”.
Confidence	**High**	All five dimensions evidenced by two or more participants.

15. Tender and bid responses

70 / 100

DIMENSION	SCORE	EVIDENCE
Operational pain	**13**	About 40 tenders a year at 40–60 hours; a third of each spent finding and adapting past answers.
Process repeatability	**15**	60–70% of content is reused.
Exposure to risk	**15**	Own content, no candidate data; signed declarations mean every word is read.
Readiness	**11**	Around 350 past answers, badly named; no current fact sheet.
Appetite	**16**	Strongly wanted, on condition that drafts show their sources.
Total	**70** (70 unrounded)	Partial automation, with a human in the loop.
Confidence	**Medium**	Pain and repeatability mainly from one participant; risk and appetite from several.

16. Client and tender questions about AI

65 / 100

DIMENSION	SCORE	EVIDENCE
Operational pain	**17**	9 of the last 14 tenders had an AI question; ~30 Dutch client AI clauses or questionnaires this year (4–5 last year); 7 pharma questionnaires in a quarter; turnaround sometimes weeks.
Process repeatability	**13**	Three underlying questions, worded differently each time.
Exposure to risk	**10**	Written, sometimes contractual, representations; Legal must sign off.
Readiness	**8**	No register of AI in use, no approved wording; past answers scattered in personal folders.
Appetite	**17**	Asked for by leaders, legal, data protection, bids and business development.
Total	**65** (65 unrounded)	Partial automation, with a human in the loop.
Confidence	**High**	All five dimensions evidenced by several participants.

17. Management and client reporting packs — 70 / 100

DIMENSION	SCORE	EVIDENCE
Operational pain	16	About two analyst days a week on the weekly internal pack, little of it analysis; ~3 analyst days a week on client packs; the monthly board pack also takes significant leadership time.
Process repeatability	15	Same metrics every week; the cleaning needs judgement.
Exposure to risk	14	Internal and client performance data; errors reach the board or clients, so checks stay.
Readiness	8	Four sources that don't join; division codes differ by office; dependence on individuals.
Appetite	15	Strongly wanted, on condition that the data is fixed first.
Total	70 (68 unrounded)	Partial automation, with a human in the loop.
Confidence	**High**	All five dimensions evidenced by two or more participants.

18. Vendor contract and DPA first-pass review — 70 / 100

DIMENSION	SCORE	EVIDENCE
Operational pain	14	A standing queue of vendor contracts, many now with AI terms, at 2–3 hours each for the AI ones; vendor due diligence is behind.
Process repeatability	15	The same handful of clauses are checked every time.
Exposure to risk	13	Confidential contracts; a lawyer signs off.
Readiness	12	Standard positions exist; an approved assistant is available.
Appetite	14	Wanted "cautiously", with the tool itself reviewed first.
Total	70 (68 unrounded)	Partial automation, with a human in the loop.
Confidence	**Medium**	Two participants; risk and readiness from one.

19. Candidate database clean-up (duplicates and retention) — 65 / 100

DIMENSION	SCORE	EVIDENCE
Operational pain	16	About 400,000 records, an estimated 30% duplicates; ~38,000 Dutch records past retention, ~1,500 a month cleared by hand; London not yet counted.
Process repeatability	14	Retention is mostly rules; unclear duplicate pairs need a person to decide.
Exposure to risk	10	Deletions can't be undone and wrong merges lose history; propose-then-confirm keeps it manageable.
Readiness	9	Consent dates missing on about half of Dutch records after a 2022 migration; no data owner.
Appetite	16	Strongly wanted by data, IT and operations roles.
Total	**65** (65 unrounded)	Partial automation, with a human in the loop.
Confidence	**High**	All five dimensions evidenced by two or more participants.

20. Subject access request (DSAR) search — 60 / 100

DIMENSION	SCORE	EVIDENCE
Operational pain	16	About 35 requests a month group-wide at ~6 hours on average, and rising.
Process repeatability	13	Search and collation about 60% mechanical; redaction is judgement.
Exposure to risk	8	Full personal data including health; one-month statutory deadline. The tool would only find, never decide.
Readiness	8	Data across the ATS, timesheets, video platform, mailboxes, Teams and phones; no single search.
Appetite	14	A top wish in governance and operations roles.
Total	**60** (59 unrounded)	Partial automation, with a human in the loop.
Confidence	**Medium**	Two participants; risk evidenced by one.

21. Internal knowledge assistant (HR policies and how-to guides) 65 / 100

DIMENSION	SCORE	EVIDENCE
Operational pain	13	Dozens of repeat HR policy questions a week; the same ~20 questions from every new-starter cohort.
Process repeatability	16	The answers already exist in policies and guides.
Exposure to risk	14	Internal content; case-specific HR matters must go to a person.
Readiness	9	Three sets of country policies; ~140 guides, about half out of date, with no owner.
Appetite	13	Wanted by people, learning and IT roles.
Total	65 (65 unrounded)	Partial automation, with a human in the loop.
Confidence	**Medium**	Three participants, but pain not precisely quantified.

INSUFFICIENT DATA TO SCORE

WORKFLOW	WHAT'S MISSING
Employment contract paperwork	Time per document, the systems used and appetite.
IR35 file notes and dispute chronologies	Readiness (where evidence is held and whether it can be searched) and appetite beyond pre-filling file notes.
Practical questions from workers and candidates	Volumes, time spent and an owner. Described as multilingual and out of hours.
Checking the sales pipeline for forecasts	Repeatability and readiness. Mainly a data quality problem in the ATS.