



Change by Design

Practical AI. Practiced Hand.

A CXD BLUEPRINT · MAY 2026

The AI blueprint for Northern Ireland recruitment agencies.

**12 workflows the top agencies are running today – and
how to build them with the software you already have.**

Northern Ireland · For agency owners, MDs & senior leaders

changebydesign.com

WHY THIS MATTERS TO YOU, RIGHT NOW

The top agencies are pulling away. The gap is widening every quarter.

Two things are happening at the same time. The agencies that have adopted AI early are placing candidates twice as fast and growing revenue at twice the rate of the rest. The agencies running AI-assisted screening are placing more people per recruiter and filling roles faster. Both gaps compound: every quarter you wait, your competitors get further ahead.

4x

Top-performing agencies are 4x more likely to use AI — and twice as likely to have grown revenue in the last 12 months.

Bulhorn GRID 2025

36%

Agencies running AI-assisted screening report 36% more placements per recruiter and a 22% higher fill rate.

Bulhorn AI users, 2025

This blueprint is what to do about it. Inside, you'll find:

- 12 workflows the leading agencies are running today
- **How to build them with your team, using the software you already have, in the next 90 days**
- What you can do yourself
- What to watch out for
- When outside help genuinely pays off

Every workflow in here is built around two questions: **does it save your team meaningful hours, and does it move the placement number?** If a workflow doesn't do both, it's not in this document.

WHY NOW

Six questions every agency should be able to answer.

Most agencies don't choose to overhaul their tooling – they get pushed. The push, right now, is a widening competitive gap and a margin squeeze. Both demand the same response: pick the workflows that matter, build them well, and stop paying for tools that aren't earning their keep.

The agencies pulling ahead can answer all six of the questions below. Most agencies can answer two or three.

-
- 01 How long does it take you to get from brief to shortlist?

 - 02 What percentage of your placements come from your existing database?

 - 03 How many hours per week is each consultant spending on admin?

 - 04 Do you know your candidate response rates by channel and message variant?

 - 05 Can your agency see, in one place, what every AI tool is costing and earning?

 - 06 Are your consultants using AI tools at all – with or without permission?
-

The work to answer those six questions is the same work as building AI workflows that move the numbers. You can't do one without the other. Start in the next quarter and you'll be ahead by the autumn. Wait until next year and you'll be playing catch-up against agencies that already shipped.

Treat this as an operating review, not a tech project. The agencies that win do this once a quarter. The next pages walk through the twelve workflows worth your time.

CONTENTS

How to use this blueprint.

Each workflow uses the same shape – read end to end if you’re planning a transformation, or jump to the three that match your loudest pain and start there.

Challenge	What the workflow fixes.
Tools you’ll need	Named tools that work in Northern Ireland today.
How to set it up	What you actually do, in order.
The upside	Realistic numbers, sourced where possible.
Watch out for	Data protection, bias, and operational risks.
Estimated cost	From-scratch pricing in GBP – most agencies already have some of these.
The tricky bits	Where outside hands pay off.

FINDING PEOPLE

- 01** The 48-hour shortlist
- 02** Reaching candidates who aren’t looking
- 03** Mining the gold in your ATS basement

ASSESSING PEOPLE

- 04** Auto-screening at the point of application
- 05** AI-assisted reference checking

ENGAGING PEOPLE

- 06** Multi-channel sequenced outreach
- 07** The conversational candidate experience
- 08** Interview scheduling, fully automated
- 09** Call notes that write themselves

RUNNING THE BUSINESS

- 10** The intent-signal BD pipeline
- 11** Content at consultant scale
- 12** The AI-ready scorecard

01 The 48-hour shortlist

CHALLENGE

The average UK hire takes around ten weeks. The first ten days are usually spent building a longlist – manual searches, LinkedIn pulls, asking around the desk. By the time you present a shortlist, two of your five candidates have taken counter-offers.

TOOLS YOU'LL NEED

- **GoPerfect, Pin, or Fetcher.** Search candidate databases by meaning, not just keywords. Connect to most ATS systems.
- **Clay or Apollo.** Pull verified contact details and signals (funding, hiring, news).
- **Claude or ChatGPT.** Score candidates against your brief with a structured prompt.

HOW TO SET IT UP

- 01** Write the brief in five fields: role, must-haves (3 bullets max), nice-to-haves, hard exclusions, six-month success.
- 02** Run the search wide on the first pass – you're looking for non-obvious candidates, not just the LinkedIn obvious ones.
- 03** Pipe the top 50–100 results into Clay to add contact data and recent signals.
- 04** Run them through an LLM with your scoring prompt. It returns scores in a format you can paste straight into your ATS.
- 05** A consultant reviews the top ten and picks five for shortlist. Keep a human in the loop on every ranking decision.

THE UPSIDE

Time-to-shortlist drops from ten days to forty-eight hours. Six to nine recruiter hours saved per assignment. Shortlist quality is at least equal – and usually better, because matching by meaning surfaces passive candidates Boolean searches miss.

WATCH OUT FOR

Profile data scraped from public sources is still personal data under UK GDPR – document your lawful basis under legitimate interest with a balancing test. Keep a logged human-review step and run a quarterly bias check on the scoring outputs.

ESTIMATED COST

If you're starting from scratch, you'll need GoPerfect, Clay, and an LLM API. All in: about £750–850 per month per consultant.

THE TRICKY BITS

The sourcing tool is easy. The brief structure, the scoring prompt, and the bias check process are not. A few days with someone who has done this in production saves a quarter of trial-and-error.

PRIORITY

- Start in the next 30 days _____
- Already running – refine _____
- Park for now _____

NOTES

02 Reaching candidates who aren't looking

CHALLENGE

Active candidates are not enough. You're competing for the same shortlist as every other agency. The placements that move you up the league table are the passive ones – senior people who'd move for the right role but aren't on the market today. They're not on Indeed and they're not answering InMails.

TOOLS YOU'LL NEED

- **LinkedIn Sales Navigator.** Watch for people changes inside target companies.
- **Clay or Common Room.** Combine signals from many sources into one feed.
- **SeekOut.** Deep search across 800m+ profiles, including hard-to-find specialisms.
- **SourceWhale or Pin.** Send sequenced messages across email and LinkedIn.

HOW TO SET IT UP

- 01 Define five trigger events that make a passive candidate worth approaching (a missed promotion, a reorg, a competitor's funding round, a manager leaving, a long tenure in role).
- 02 Wire those triggers into a weekly feed using Clay or a no-code tool like Make.
- 03 Have an LLM rank each name against your live briefs. Top ten go to consultants.
- 04 Write outreach that leads with the trigger – "saw your team just closed Series B" lands differently than a generic intro.

THE UPSIDE

20-30% of placements from passive outreach within six months, against an industry baseline of 5-10%.

WATCH OUT FOR

B2B outreach in the UK sits under PECR and UK GDPR – legitimate interest works with a clear opt-out, which you must honour. Inbox fatigue is real: the trigger-led approach works because it's specific. **Generic AI-written outreach makes you look worse than no outreach at all.**

ESTIMATED COST

From scratch: Sales Navigator ~£85/mo per seat, Clay ~£300/mo team, SourceWhale ~£110/mo per user, SeekOut from £3,800/year per seat. If you already have Sales Navigator and a sequencing tool, you're most of the way there.

THE TRICKY BITS

The trigger logic and the outreach writing are where most agencies stall. An experienced pair of hands gets you to a working feed in days, not weeks.

PRIORITY

- Start in the next 30 days _____
- Already running – refine _____
- Park for now _____

NOTES

03 Mining the gold in your ATS basement

CHALLENGE

Your ATS has somewhere between 20,000 and 200,000 candidate records. Maybe 5% have been touched in the last six months. The rest sit there, unsearchable in any useful way, while your consultants go back to LinkedIn to find people who already exist in your database.

TOOLS YOU'LL NEED

- **Recruiterflow or Bullhorn Amplify.** If you're on one of these, turn the AI search feature on.
- **GoPerfect.** Sits on top of your ATS and adds a chat-style candidate search.
- **Custom build.** Vector search using Pinecone or Weaviate. More flexible, more work.

HOW TO SET IT UP

- 01 Audit retention first – UK GDPR rules apply to old records. Prune what you can't justify holding.
- 02 Enrich the top 10,000 records to update employment status, current title, and contact validity.
- 03 Index the records for chat-style search – either through your ATS's AI feature or a bolt-on.
- 04 Build a search interface inside your ATS or Slack. Consultants ask in plain English.
- 05 Train the team. This is more about behaviour change than tech – consultants default to LinkedIn out of habit.

THE UPSIDE

15–25% of placements come from existing contacts in agencies that mine their database well. Real money at typical perm fees. Spend on external sourcing tools drops in parallel.

WATCH OUT FOR

Retention is the sharpest area of ICO enforcement for recruiters – be clear about how long you keep records and how the deletion process works. Tone matters too: a candidate you haven't spoken to in two years doesn't want a sales pitch as the first message back.

ESTIMATED COST

Recruiterflow from £85/mo per user. Bullhorn Amplify is a paid add-on to Bullhorn. Custom vector-search builds run £6,800–17,000 one-off plus £170–425/mo. If you're already on an AI-native ATS, the cost is just turning the feature on.

THE TRICKY BITS

A custom vector-search build is genuinely technical – don't try it without an engineering team. Turning on an ATS-native feature is mostly configuration plus a couple of training sessions.

PRIORITY

- Start in the next 30 days _____
- Already running – refine _____
- Park for now _____

NOTES

04 Auto-screening at the point of application

CHALLENGE

Job-board applications produce 49% of your candidate flow but only 24.6% of your hires. Most are wrong for the role. With fourteen open requisitions per recruiter, the screening backlog is where good candidates get lost – not because they were rejected, but because nobody looked at them in time.

TOOLS YOU'LL NEED

- **Bullhorn Amplify, Recruiterflow, JobAdder AI, or Mokka.** ATS-native screening, plugged straight into your existing workflow.
- **GoPerfect.** Standalone screening that connects to 60+ ATS systems.
- **Manatal.** Lower-cost option for smaller agencies.

HOW TO SET IT UP

- 01 Build a scoring rubric for each role family (group your roles into 8–15 families).
- 02 Calibrate the model: take 50 CVs you placed and 50 you rejected, run them through, adjust until the scores match your decisions about 80% of the time.
- 03 Wire to your ATS so every new application triggers a score.
- 04 Set thresholds – top scores route to consultants, bottom scores to a polite auto-rejection (signed off by a recruiter), middle hold for capacity.
- 05 Audit weekly for the first month. If the model is wrong on more than a couple of CVs, refine the rubric.

THE UPSIDE

Screening time drops by up to 85%, and recruiters reclaim four to six hours a week. Bullhorn agencies running their automation features report 36% more placements per recruiter and a 22% higher fill rate.

WATCH OUT FOR

UK GDPR Article 22 covers automated decisions that have a significant effect on individuals. **A recruiter must sign off the rejection list** (one click is fine) – you cannot fully automate a decision to reject a candidate. Bias audits are essential: log every score, audit quarterly, and be ready to explain how a candidate was scored if asked.

ESTIMATED COST

Bullhorn Amplify and Recruiterflow are bundled with the ATS – usually £85–170/mo per user. Standalone tools £170–500/mo per user. If you're already paying for one of the big agency ATS systems, the AI feature is probably already in your plan.

THE TRICKY BITS

Rubric calibration and the bias-audit process are the two places agencies most often get this wrong. Both benefit from outside experience on the first build.

PRIORITY

- Start in the next 30 days _____
- Already running – refine _____
- Park for now _____

NOTES

05 AI-assisted reference checking

CHALLENGE

References are the unloved part of the placement. Done well, they protect the placement and catch culture-fit risks before they become problems. Done badly – which is most of the time – they're a phone call where the referee says "yes, they were great" and the consultant ticks a box.

TOOLS YOU'LL NEED

- **Refapp, Xref, or HiPeople.** Purpose-built reference tools with question libraries and AI summaries.
- **Typeform plus an LLM step.** Cheaper DIY route using a form and a summary step via Zapier or Make.
- **Talkpush or Hireguide.** Voice-based for older referees who won't fill in a web form.

HOW TO SET IT UP

- 01** Build a reference template library – one per role family, eight to ten questions, a mix of rating scales and short open text.
- 02** Capture candidate consent and referee details at the brief stage, not the offer stage. This is the single biggest operational bottleneck.
- 03** Automate the send: when a candidate hits "offer" stage, references go out automatically.
- 04** When both references come back, an LLM produces a one-page summary against the brief.
- 05** The consultant reads the summary and decides whether a follow-up call is needed.

THE UPSIDE

60–70% of reference-check time recovered. Better data quality because the questions are standard. Fewer post-placement surprises because structured questions surface what "yeah, they were great" hides.

WATCH OUT FOR

The referee is a data subject – you need their consent to process. The tools above handle this; if you DIY, build it in. Defamation risk is real, so use legal-vetted question wording. Keep the form short – ten minutes is the upper limit before referees abandon.

ESTIMATED COST

Refapp from about £3,800/year for a small agency. Xref similar. DIY Typeform-plus-LLM about £50/mo plus your time. If you already have a workflow tool like Zapier or Make, the DIY version costs almost nothing to try.

THE TRICKY BITS

Light touch. The technical setup is the easy bit. The hardest part is operational – building the habit of capturing referee details at the brief stage, not at the offer stage.

PRIORITY

- Start in the next 30 days _____
- Already running – refine _____
- Park for now _____

NOTES

06 Multi-channel sequenced outreach

CHALLENGE

Cold outreach response rates have collapsed. Your consultants send fifty InMails a week and get four replies. The candidates you want receive twenty messages a week from agencies that all sound the same.

TOOLS YOU'LL NEED

- **SourceWhale.** Recruitment-native, multi-channel sequencing across email and LinkedIn.
- **Pin.** End-to-end platform with sourcing and outreach in one.
- **Apollo or Outreach.io.** Sales-built but work well for recruitment.
- **Clay.** Personalisation layer that pulls signals into your messages.

HOW TO SET IT UP

- 01 Map your sequence – four to six touches over two to three weeks across email and LinkedIn.
- 02 Build personalisation tokens. "First name" is the minimum. Add "recent post topic", "recent funding", "role tenure".
- 03 Use an LLM to produce five message variants, pick the best, edit the first line by hand.
- 04 A/B test continuously and kill the bottom-performing variant every fortnight.
- 05 Hand warm replies to a consultant inside an hour. Speed kills conversion if you're slow.

THE UPSIDE

LinkedIn's own data shows AI-assisted messaging achieves 40% higher acceptance rates. **Multi-channel sequencing typically lifts response 2-3x over single-channel.**

WATCH OUT FOR

PECR and UK GDPR apply to B2B outreach. **Voice drift is the brand risk – generic AI outreach is worse than none.** AI lets you send five times the volume; that doesn't mean you should. Pick fewer, better targets.

ESTIMATED COST

SourceWhale ~£110/mo per user. Pin from £255+/mo per user. Apollo £50-100/mo per user. Clay ~£300/mo team. If you already use Sales Navigator and Apollo, you may have most of this – the gap is usually the sequencing layer.

THE TRICKY BITS

Sequence design and voice work are what separate 20% response rates from 2%. The tools are commodities. The thinking around them is not.

PRIORITY

- Start in the next 30 days _____
- Already running – refine _____
- Park for now _____

NOTES

07 The conversational candidate experience

CHALLENGE

Candidates ask the same questions over and over. Salary, location, status, why have I not heard back. Your consultants spend half their inbox on the same five questions, and the candidates who don't hear back inside 24 hours go cold.

TOOLS YOU'LL NEED

- **Paradox (Olivia), Sense, or Mya.** Purpose-built for high-volume hiring.
- **Custom GPT or Claude project.** Lighter option for smaller agencies – train it on your FAQs.
- **Channel choice.** Web chat for inbound, WhatsApp or SMS for active candidates.

HOW TO SET IT UP

- 01 List your top 20 candidate questions. Most agencies already have these in shared docs.
- 02 Write the answers properly. The bot is only as good as the source material.
- 03 Decide what escalates to a human and what doesn't. Salary negotiation always escalates. Dress code never does.
- 04 Pilot on one role family – don't turn the bot loose on everything from day one.
- 05 Read transcripts every week for a month. You'll find five questions you didn't anticipate. Add them in.

THE UPSIDE

60% drop in candidate-side admin queries. Response time on "where am I in the process?" goes from 18 hours to zero. Consultant inbox volume drops by a third.

WATCH OUT FOR

Transparency matters – **tell the candidate they're talking to a bot.** One sentence at the start is enough. Your bot is now part of your brand: robotic and generic equals bad candidate experience. Be generous with escalation in the first six weeks.

ESTIMATED COST

Paradox is enterprise-priced, typically £12,750+/year. Sense similar. A custom GPT/Claude approach is closer to £85/mo plus integration time. If you already pay for ChatGPT Team or Claude Team, you have the model side covered.

THE TRICKY BITS

The knowledge base writing and the escalation logic are where this works or doesn't. Tools matter less than the thinking that goes into them.

PRIORITY

- Start in the next 30 days _____
- Already running – refine _____
- Park for now _____

NOTES

08 Interview scheduling, fully automated

CHALLENGE

Calendar coordination eats 38% of recruiter time. A senior consultant on £55,000 spending two hours a day on calendar admin is £12,750 of pure waste a year. Three-way scheduling between consultant, candidate, and client is the worst of it.

TOOLS YOU'LL NEED

- **Paradox, GoodTime, or Modern Loop.** Purpose-built scheduling, deep ATS integration.
- **Calendly or SavvyCal.** Lightweight and works with a Make/Zapier flow for prep materials.
- **Bullhorn or Recruiterflow.** Both have scheduling built into the ATS.

HOW TO SET IT UP

- 01 Get every hiring manager's calendar permissioned (read-only is fine). Frame it as: "we need this so we stop chasing you."
- 02 Set rules – buffer time, working hours, time zones, round-robin if multiple interviewers.
- 03 Wire to the ATS so the scheduling link goes out when a candidate hits "interview" stage.
- 04 Auto-send prep materials with the calendar invite – job spec, interviewer bios, location or Zoom link.
- 05 Build a human-handoff path for senior candidates who want a person, not a link.

THE UPSIDE

38% of recruiter time reclaimed. Faster scheduling means faster process means fewer candidates lost to counter-offers. The highest-return, lowest-risk workflow in this blueprint.

WATCH OUT FOR

Not much. **The main operational risk is getting hiring managers to use the scheduling link instead of email out of habit.** Frame it as a service to them, not extra work.

ESTIMATED COST

Calendly from £8.50/mo per user. Paradox and GoodTime are enterprise-priced. ATS-native scheduling is usually included. Most agencies already have Calendly or similar – the gap is the prep-material automation layer, which is mostly free.

THE TRICKY BITS

Almost none. Start here if you've done nothing yet. Pays for itself in six weeks.

PRIORITY

- Start in the next 30 days _____
- Already running – refine _____
- Park for now _____

NOTES

09 Call notes that write themselves

CHALLENGE

Your consultants run six to ten calls a day. They're supposed to write notes into the ATS after each one. They don't. Or they write three lines that are useless to anyone else. The result is an ATS full of half-records, and placements that depend on whoever spoke to the candidate first.

TOOLS YOU'LL NEED

- **Fathom, Otter, Granola, Read.ai.** Auto-join calls, transcribe, and summarise. All have free tiers.
- **Zapier or Make.** Write the summary back to the candidate record in your ATS.
- **Bullhorn-Fathom integration.** Native, no glue code needed if you're on Bullhorn.

HOW TO SET IT UP

- 01 Get consent right. Under UK GDPR you must tell people you're recording. One-party consent is fine, but transparency is required.
- 02 Pick one tool and standardise. Fragmented adoption (some on Otter, some on Fathom) destroys the value.
- 03 Build a standard summary template – six headings, every call: context, candidate situation, current package, what they want, fit against current roles, follow-up actions.
- 04 Wire the summary to land in the candidate record automatically.
- 05 Train consultants to trust the summary. The shift is from "I'll write notes later" to "the notes write themselves and I review them."

THE UPSIDE

60% of consultant admin time recovered. ATS data quality jumps overnight. The placement is no longer dependent on the consultant who spoke to the candidate first.

WATCH OUT FOR

Under UK GDPR you must tell people you're recording. Transcripts contain a lot of personal data – set auto-deletion windows. Some candidates will say things they wouldn't want on record; have a clear policy to delete on request. Recording an in-person meeting is harder to handle gracefully than a Zoom – most agencies stick to digital for this workflow.

ESTIMATED COST

Fathom and Granola have free tiers. Paid plans about £12–25/mo per user. Most consultants already use one of these personally – the work is standardising on one and wiring it to the ATS.

THE TRICKY BITS

Light. The tech is straightforward. The harder part is consultant adoption – pair the rollout with a clear "this is how we work now" from agency leadership.

PRIORITY

- Start in the next 30 days _____
- Already running – refine _____
- Park for now _____

NOTES

10 The intent-signal BD pipeline

CHALLENGE

The recruitment market has stalled in pockets – 47% of recruiters say they've cut graduate and entry-level roles. Your BD pipeline is shrinking just when you need it to grow. Cold-calling out of a static list doesn't work the way it used to.

TOOLS YOU'LL NEED

- **LinkedIn Sales Navigator.** People-move signals inside target companies.
- **Crunchbase or PitchBook.** Funding signals.
- **Clay or Common Room.** Pull and combine signals into one weekly feed.
- **SourceWhale, Pin, or Apollo.** Outreach layer.

HOW TO SET IT UP

- 01** Define what "in market" means for your specialism. A CFO joining. A Series B closing. A regulatory filing. A competitor leadership change.
- 02** Build the weekly feed in Clay or Common Room. Two or three sources to start, add more as you go.
- 03** Have an LLM rank the feed against your active brief list and your existing client list (avoid pitching to a client's competitor).
- 04** Lead with the trigger in your outreach. "Congratulations on the Series B" is a different conversation than "I see you're hiring."
- 05** Build a weekly BD ritual where the top five are reviewed and assigned to consultants.

THE UPSIDE

2–3x more BD meetings per consultant in agencies that build this well. Conversion to retainer or exclusive is typically 25–40% higher than cold-list outreach.

WATCH OUT FOR

Conflict of interest is the biggest risk – build a conflict check into routing logic so you're never pitching the same role to two competing clients. Signal quality matters more than volume: a job post plus a leadership change is a much stronger signal than either alone.

ESTIMATED COST

Sales Navigator ~£85/mo per user. Clay ~£300/mo team. Apollo £50–100/mo per user. Common Room from £850/mo for an agency-size plan. If you already have Sales Navigator, you have the foundation.

THE TRICKY BITS

Signal design and routing logic are where this is made or broken. Worth spending a few days with someone who has built one of these before.

PRIORITY

- Start in the next 30 days _____
- Already running – refine _____
- Park for now _____

NOTES

11 Content at consultant scale

CHALLENGE

Your senior consultants should be posting on LinkedIn three times a week. They post twice a month. Job descriptions take forty minutes to write. Marketing has one person; they can't keep up with twelve consultants.

TOOLS YOU'LL NEED

- **Claude (Sonnet or Opus) or ChatGPT (4-class)**. Team or enterprise tier so prompts aren't used for training.
- **Bullhorn or Vincere native AI**. Built-in JD and outreach copy generation.
- **Taplio or Authoredup**. LinkedIn-specific scheduling and drafting.

HOW TO SET IT UP

- 01** Build a voice prompt for each consultant. Take three pieces of their writing, ask the model to extract the voice into a reusable prompt.
- 02** Build template prompts for each content type – JD, LinkedIn post, candidate placement update.
- 03** Set a daily 20-minute cadence. Not "when you have time." Most agencies pair this with a Monday content-planning slot.
- 04** Human always in the loop on anything with a consultant's name on it.
- 05** Measure what works in LinkedIn analytics or Taplio. Feed the patterns back into the prompts.

THE UPSIDE

5x content output per consultant. JD writing time drops from 40 minutes to 8. Note: 70% of UK recruitment firms already use AI for content – this is no longer competitive advantage, it's table stakes.

WATCH OUT FOR

Voice homogeneity is the brand risk – the same prompt across consultants makes your LinkedIn feed feel like wallpaper. The model invents salary ranges and stats, so humans check anything quantitative. Have a brand-safe topic list (politics, religion, conflict zones) and make it explicit.

ESTIMATED COST

Claude Pro or ChatGPT Plus ~£17/mo per user. Team plans ~£21/mo. Taplio ~£33/mo per user. ATS-native AI usually included. Most consultants already pay for one of these personally – moving them onto a Team plan covers the privacy concern.

THE TRICKY BITS

Voice prompt design is where most of the value lives. A half-day workshop with each consultant pays for itself in six weeks.

PRIORITY

- Start in the next 30 days _____
- Already running – refine _____
- Park for now _____

NOTES

12 The AI-ready scorecard

CHALLENGE

You've spent money on AI tools across the agency. Leadership can't answer "is it paying off?" with a number. The board wants ROI. The audit trail of which workflow is producing what value doesn't exist today.

TOOLS YOU'LL NEED

- **Your ATS.** Source of truth for placement, fee, and time data.
- **Make or Zapier.** Pull data from AI tool admin dashboards into one place.
- **Looker Studio (free) or Tableau.** Visualisation. For a small agency, a Google Sheet refreshing nightly is enough.

HOW TO SET IT UP

- 01 Decide the metrics first, tools second. The mistake is buying a dashboard and asking what to put in it.
- 02 Wire the ATS as the source of truth. Fix the data hygiene before building the dashboard (Workflow 09 helps a lot).
- 03 Build the dashboard. Start ugly. Make it work. Prettify later.
- 04 Read it at the same time each month – last week, three actions out, captured.
- 05 Track per-workflow cost and per-workflow output. Kill what isn't earning.

THE UPSIDE

Real-time view of whether AI is paying off. Faster decisions on which tools to keep or kill. **Board confidence in your AI spend.** A scorecard is also the cheapest way to find out which two tools you can cancel today.

WATCH OUT FOR

Vanity metrics are the trap – "CVs screened by AI" tells you the AI is busy; "days to shortlist" tells you the business is healthier. Pick the latter. If consultants are tracking placements in Excel as well as the ATS, the scorecard is wrong – fix the source before you fix the dashboard.

ESTIMATED COST

Looker Studio is free. Make from £8.50/mo. Tableau from £60/mo per user. Mostly consultant time, not software cost. If you already have a BI tool, you're three afternoons of work from this.

THE TRICKY BITS

Metric design is harder than the build. Get the metrics right with someone who has done it before and the rest of the work is mechanical.

PRIORITY

- Start in the next 30 days _____
- Already running – refine _____
- Park for now _____

NOTES

When to bring outside hands in.

Most of what's in this blueprint you or your team can do yourselves. The agencies that succeed treat AI as a discipline, not a project – and most of the work is operational habit-building, not technical complexity. The places where outside help pays off:

The first 30 days.

Picking the right two or three workflows to start with, given the software you already have, your team, and your sector. Sequencing matters more than tool choice.

Process design and bias-audit calibration.

Calibration of scoring rubrics, design of bias audits, the documentation that protects the business. The work where mistakes are most expensive and prior experience saves the most time.

Change management.

Getting twelve senior recruiters to actually change how they work is harder than the technology. Pair every tool rollout with a clear "this is how we work now" from agency leadership.

Pre-rollout pressure test.

A review before you turn a workflow on to your full team catches problems while they're cheap to fix. Cheaper than rolling back six weeks of bad data.

Want a hand?

Change by Design offers fractional Head of AI engagements built for Northern Ireland recruitment agencies. Three tiers – advisory, hands-on, or a 30-day sprint.

30-minute discovery call, no pitch deck. If you like what you hear, we'll have a more in-depth conversation.

hello@changebydesign.com · changebydesign.com



Change by Design

Practical AI. Practiced Hand.

We work with SMB and mid-market businesses across Ireland, the UK, and North America. Our focus is operational AI – the workflows that get installed, used, and measured.