



HOUSE OF VISIONARIES PRESENTS

Intentional Building with AI

TALENT ACQUISITION · AGENTIC SYSTEMS EVENT

[SHOW & TELL](#)



Presented by:

Mitchell Gilbey

How we find problems first, build vs. buy with discipline, and scale a human-led recruiting engine at VIXIO.

What we'll walk through

01

Decision-Making & Why

Intentional building: finding real problems before adding tools

03

Our AI Layer in Action

Gemini, Google CLI, Greenhouse + MCP, and our TA Brain command centre

05

Impact & ROI

Hours saved, drop-off, time-to-hire, retention, and capacity gains

02

Build vs. Buy

What we build ourselves, what we pay for, and how governance & cost decide it

04

The Workflow, End to End

From JD to offer - where AI helps and where humans decide


06

Rolling It Out

Enablement across the team and the wider business


Start with the problem. Not the tool.

“Intentional building” is our rule at VIXIO: we don’t start with a tool and look for a use case. We start with a friction point in the candidate or recruiter experience, prove it with data, and only then decide whether to build, buy, or leave it alone.



The Forensic problems

Using our AI reporting dashboards built on Vercel linked with our ATS data, we highlight the 3 key bottlenecks linking to an increased Time to hire.




Candidates spending too much time at the sourcing/1st stage

US roles - especially repeatable ones like analysts, where the spec barely changes - were burning disproportionate recruiter hours on sourcing and first-touch calls.



Candidate drop off

9% of candidates disengaged mid-process. Slow handoffs between sourcing, scheduling and feedback were costing us strong candidates.



Hiring Manager feedback

Manual scorecards, ad-hoc interview questions and scheduling back-and-forth stretched an already competitive hiring cycle.

These three signals - not the availability of new AI tools - defined our roadmap.

What we build, what we buy

Every tool decision runs through the same three questions.



Capability

Can we build it ourselves?

Gemini + Google CLI, prompt-engineered inside a governed “gem,” covers JD writing, scoring, scheduling logic and outreach - no new vendor needed.



Governance

Who controls the data & decisions?

Anything touching candidate PII, scoring rubrics or ATS records stays inside our Google/Greenhouse environment, under our access controls.



Cost & speed to value

Is buying faster or cheaper at our scale?

For narrow, high-volume, repeatable needs - like Round 1 screening on US analyst-type roles - a specialised paid tool beats build time.

Our Stack



Gemini + Google CLI

Our AI layer and knowledge base.



“TA Brain” gem

A custom, governed Gemini gem that acts as our recruiting command centre - every step routes through it.



Greenhouse ATS + MCP (beta)

Our system of record. The Greenhouse MCP lets the gem read and act on ATS data directly to automate steps.



Gmail MCP

Connects scheduling and candidate communication straight from the workflow into recruiters’ inboxes.



Round 1 AI (paid)

A specialised, bought tool for first-round screening on high-volume, repeatable US roles.




Human decision layer

Every hiring decision - shortlist, offer, reject - is made by a person, not the AI.


The 3 big AI implementations

The real, live tools our team uses every day - not mockups.


TA Brain command centre (Gemini)



R1 AI voice screening - MitchBot




Candidate chatbot assistant



Our real life AI workflow

From requisition to offer - through the TA Brain.



Inside the Brain Gem: hiring strategy prep

Triggered the moment the JD lands in the gem (step 2) - the gem prepares everything below before a human touches it.

- Enhanced Job Description
- Sourcing Strategy
- Boolean & Talent Searches
- Candidate Outreach Messages
- Follow-up Sequences
- Interview Scorecards (Greenhouse-linked)
- Interview Questions
- Technical Task Briefs

The results, at a glance

Faster, more engaged candidates - and a team that scaled without growing.



CANDIDATE DROP-OFF

9% → <1%


Faster handoffs across sourcing, scoring & scheduling



US HIRING CYCLE TIME

17 days

After introducing Round 1 AI on repeatable US roles



RECRUITER HOURS SAVED

5 hrs / wk

Per recruiter, across a team of 4



TIME-TO-HIRE (OVERALL)

44.8 → 31 days

Executive hires: 61 → 49 days



12-MONTH RETENTION

72% → 89%

Better-matched hires via values & competency scoring



ANNUAL HIRING CAPACITY

~50 → ~75 roles/yr

+50% growth - same team of 4, no added headcount

Enablement across the team and the business

Building the tool was the easy part - this is how we made it stick.



Recruiter enablement

Hands-on training on the TA Brain gem, paired working sessions, and a feedback loop to refine prompts & rubrics.



Hiring manager buy-in

HMs see faster shortlists and consistent scorecards - built trust that kept adoption high across the business.



Usage tracked, not assumed

We monitor gem usage, ATS automation triggers, and R1 AI screen volume to know adoption is real, not just available.



Hiring metrics as the scoreboard

Time-to-hire, drop-off, retention and hours saved are reviewed monthly - the same metrics that justified build vs. buy.



Governance stays live

Access, rubric changes and data handling are reviewed on a cadence, not set once and forgotten.



Continuous ROI review

Capacity and cost-per-hire gains are re-measured each quarter to justify further investment or a course correction.

QUESTIONS & DISCUSSION

Thank you.

VIXIO · Talent Acquisition · Show & Tell

The result: 89% retention, sub-1% candidate drop-off, and 50% more hiring capacity - from the same team.