

FICTIONAL SAMPLE / JULY 2026

AI Opportunity Assessment

Northstar Recruitment

PROMPT -> WORKFLOW -> AGENT -> OPERATING SYSTEM

01	Candidate intake	CONTROLLED WORKFLOW
02	Client proposal	PROMPT + REVIEW
03	Interview notes	AGENT CANDIDATE

THIS COMPANY AND ALL INPUTS ARE FICTIONAL - NO RESULTS ARE CLAIMED

DECISION BRIEF

Start with control, not autonomy.

Northstar does not need an AI transformation. It needs three repeatable operating decisions: one workflow to standardise now, one prompt-led process to improve, and one bounded agent candidate to test after the quality rules are stable.

1

Workflow first

Candidate intake repeats often, has clear inputs and can keep final routing with a recruiter.

2

Prompt + review

Proposal drafting benefits from structure, but commercial judgement stays with the account lead.

3

Agent candidate

Interview-note processing can be tested only after consent, storage and QA rules are agreed.

THE OPERATING DECISION

DO NOW

Define intake criteria and run the workflow manually for two weeks.

TEST NEXT

Use a structured proposal prompt with mandatory account-lead approval.

DO NOT AUTOMATE YET

Candidate rejection, salary decisions and client commitments.

MEASURE

Cycle time, rework rate, missing information and recruiter overrides.

SMALL-TEAM SCAN

Where AI could earn its place

The scan ranks work by repetition, clarity, reviewability and cost of error. A high score is not permission to automate; it is a reason to inspect the workflow more closely.

AREA	PROCESS	REPEAT	REVIEW	DECISION
Sales	Proposal creation			PROMPT + REVIEW
Delivery	Candidate intake			WORKFLOW
Research	Talent-market scan			WORKFLOW
Operations	Interview notes			AGENT CANDIDATE
Client care	Status updates			PROMPT
People	Candidate rejection			KEEP MANUAL

Illustrative assessment logic based on fictional inputs. Scores show process characteristics, not achieved performance.

THREE OPPORTUNITIES ONLY

Focus beats tool sprawl.

Each opportunity is framed as a business decision with an owner, boundary and first measurable version.

01

Candidate intake workflow

NOW / OWNER: RECRUITER

Standardise required inputs, extract a first-pass brief and let the recruiter approve routing.

02

Proposal drafting system

NEXT / OWNER: ACCOUNT LEAD

Turn discovery notes into a structured first draft while keeping scope, pricing and promises manual.

03

Interview-note assistant


LATER / OWNER: OPS LEAD

Test a bounded assistant after consent, storage, redaction and spot-check rules are documented.

PRIORITY 01 / CANDIDATE INTAKE

Build the workflow before the agent.

The first version should run with simple tools and visible human decisions. Automation comes only after the team can see where the workflow fails.



FIRST VERSION / TWO-WEEK MANUAL PILOT

INPUT RULE	No record enters the workflow without a source, role and consent status.
QUALITY RULE	The recruiter must be able to correct extracted fields before routing.
STOP RULE	If confidence is low or context conflicts, route to manual review.
SUCCESS SIGNAL	Less rework without a rise in recruiter overrides or missing context.

CONTROL DESIGN

Autonomy stops where risk becomes real.

These controls are part of the workflow, not an afterthought. They define what the model may suggest and what only a person may decide.

Candidate data	HIGH	Minimise inputs, document consent and prohibit copying special-category data into unapproved tools.
Biased routing	HIGH	AI may structure information, but a recruiter owns every shortlist and rejection decision.
Commercial promises	MEDIUM	Proposal drafts cannot commit pricing, scope or timing without account-lead approval.
Hallucinated context	MEDIUM	Every factual claim must trace back to a source note, CV or approved company record.
KEEP MANUAL		
Final candidate rejection, salary judgement, client commitments and any decision a person cannot explain or appeal.		

SEQUENCED ACTION

The first version should be boring.

The roadmap earns autonomy in stages. Each phase has an owner, an output and a reason to stop before complexity compounds.

DAYS 1-30	DAYS 31-60	DAYS 61-90
Define 1 Choose the intake workflow owner 2 Document inputs and exceptions 3 Set quality and stop rules 4 Capture a manual baseline	Run 1 Pilot the workflow with real work 2 Review every output 3 Log overrides and missing context 4 Improve the prompt and template	Decide 1 Compare against the baseline 2 Keep, change or stop the workflow 3 Select one bounded agent test 4 Set monitoring before autonomy

Decision at day 90: Scale only if quality is stable, exceptions are understood and a named person still owns the result.