



LEVER TALENT

Agentic Workflow Vetting Worksheet

The 4Cs Strategic Harness Framework, applied to one workflow

NOTE: The worksheet below is what we have implemented at Lever Talent, Inc. Feel free to use it as is or as inspiration for crafting a strategically aligned vetting worksheet of your own. Want help? [Speak to a Workforce Orchestration Expert™](#) at Lever Talent, Inc. today.

Instructions: Complete this before you build an agent for anything. It runs one workflow through the 4Cs Harness, the same framework we use on ourselves at Lever Talent and teach in AI Fluency for HR. Your job is to describe the work accurately. Deciding the shape of the build, whether it's one agent or a loop of them, comes after this sheet is done.

Why a harness

A model is a raw engine. The harness is whatever sits between your business and that engine, making sure the right data gets in, the right output gets out, and someone is accountable for everything in between. Prompting is not a harness. This worksheet is how you build one.

The four Cs

C1 Context <i>What it must know</i>	C2 Constraints <i>Where it operates and what it obeys</i>	C3 Connections <i>What it touches</i>	C4 Control <i>Who is accountable and who is affected</i>
---	---	---	--

How to use it

- ☐ One worksheet per workflow. Don't batch three ideas into one sheet.
- ☐ Fill it in order. The four Cs are the substance of this sheet. Everything before them sets up the work, and everything after them summarizes it.
- ☐ Blank fields are the finding. A section you can't complete is telling you the workflow isn't ready, which is useful information, not a failure.
- ☐ Answer from what you actually know. Where you're guessing, say you're guessing.



Define the Work

What you're trying to accomplish and what actually happens today. Everything downstream depends on this being specific.

THE OBJECTIVE

One sentence. What outcome, for whom, and why it matters to the business.

THE TRIGGER

What specific event starts this workflow, in what system, and how often does it fire? Be literal. "A client signs an SOW in HubSpot, roughly six times a month."

WHO OWNS IT TODAY

Name the person or role currently accountable for the output.

The workflow as it runs today

Break it into the steps a person actually performs, even if there's no order. If you can't describe the steps, you can't delegate them.

#	What Happens	Who Does It	Time Per Run



The workflow as an agent would run it

Same steps, reassigned. Mark every step as an agent, human, or agent with a human gate.

#	What the Agent Does	Agent / Human / Gated	Output of the Step

What you want the agent to do

Your intent, not the technical design. How much of this workflow do you want the agent to take?

- ☐ Assist. A person still runs the workflow. The agent removes the tedious part and hands work back.
- ☐ Own with a gate. The agent runs the workflow end-to-end, and a named person approves before the output leaves.
- ☐ Own outright. The agent runs and ships without a review on every run.

DEFINITION OF DONE

Describe a finished, acceptable output precisely enough that a reviewer could grade one without asking you.



C1 Context — what it must know

What internal strategy, material, and capability the agent needs to do this well?

SOURCE DOCUMENTS THE AGENT MUST READ

Name each one, where it lives, and the date it was last updated. Stale sources produce wrong output.

STANDARDS THE OUTPUT MUST MATCH

Brand voice, formatting defaults, client-ready bar, naming conventions, and house terminology.

HISTORY THE AGENT NEEDS

Prior decisions, past work, and positions you've already taken that the agent must not contradict.

COMPETENCIES THE AGENT NEEDS

What does the agent have to be good at to do this well? Describe the capability, not the tool. For example: interpreting a specific kind of data, writing in a special voice, structuring a model, summarizing a call without losing nuance. You'll work out whether you have or need to build the skills.

THE REFERENCE EXAMPLE

Point to one existing artifact that represents the quality bar. If none exists, say so.

- | | |
|---|--|
| ☐ Source documents named and located | ☐ Required competencies described |
| ☐ Standards and voice defined | ☐ A reference example identified |
| ☐ Sources confirmed current, not stale | ☐ Gaps in any of the above noted |



C2 Constraints — where it operates and what it obeys

The process the agent operates inside, the rules it obeys, and the data it depends on.

THE PROCESS IT OPERATES INSIDE

What system, process, or part of a process does this agent sit within? Name the constraints that the process already carries: sequence, dependencies, deadlines, approvals, and handoffs.

IN BOUNDS AND OUT OF BOUNDS

Which parts of the process should the agent interact with, and which parts should it stay out of or not concern itself with? Be explicit about the out-of-bounds. That's the part people forget to say.

HARD RULES AND NEVER-DOS

Write each as a testable statement. "Never name a client in an external draft" can be checked. "Be careful with client data" cannot.

COMPLIANCE OR LEGAL CONCERNS YOU'RE AWARE OF

You're not being asked for a legal opinion. Flag anything that gives you pause, so you know to loop in whoever handles that at your org.

THE DATA

What data will the agent process and use, if any? Where does it live, and what state is it in: clean, structured, complete, or messy?



ADDITIONAL DATA THAT MAY BE NEEDED

Is there data this workflow would need that you don't have today? Does it exist somewhere? Do you know its condition, or would it have to be created?

- ☐ Host process named
- ☐ Data named and its state assessed
- ☐ Out of bounds stated explicitly
- ☐ Additional data needs to be flagged
- ☐ Rules written as testable statements
- ☐ Compliance concerns noted, if any

C3 Connections — *what it touches*

Every system, dataset, and person the workflow touches, and how hygienic each connection is.

SYSTEMS AND DATA

One row per system. Ours include HubSpot, Asana, Google Drive and Gmail, Slack, Fathom, and client-supplied files, etc. Name the connection method and be honest about its hygiene.

System or Dataset	Read / Write	How It Connects	Hygiene Concern

PEOPLE IN THE CHAIN

Who supplies the input the agent depends on, and who receives the output. Name them. A missing human upstream breaks the workflow more often than a missing API.



THE WEAKEST LINK

Which connection is most likely to fail or require a manual step? What happens to the run when it does?

- | | |
|---|---|
| ☐ Every system of record named | ☐ No manual export that silently breaks the loop |
| ☐ Access scoped to least privilege | ☐ Failure behavior known for each connection |
| ☐ Data retention terms confirmed | ☐ Upstream and downstream humans named |

C4 Control — *who is accountable and who is affected*

Who is accountable, how the data is handled, who is affected, and what you tell them.

HUMAN IN THE LOOP

One named person, not a team. What exactly do they review, and what can they override?

WHERE THE GATE SITS

Before the output leaves the workflow, before it reaches a client, or after it ships. If the gate is after, justify it.

CONFIDENTIALITY AND DATA HANDLING

Client confidentiality, assessment data, personal information, contractual or partner terms. What must not leave, and where must it never go.



FAILURE MODES

Describe the two or three most likely ways this produces a bad output, and how each one gets caught.

ESCALATION PATH

When the agent is uncertain or the output fails review, where does it go, and who decides?

EMPATHY FOR THE PEOPLE INVOLVED

How might the people this touches react to an agent being part of it? Cover all three groups: people working alongside the agent, people receiving its output, and people whose data goes in and may have a right to know. Where a reaction is likely, say how you address it.

WHO NEEDS TO KNOW

Every party or team with a stake in this running. Awareness is part of the exposure.

Party or Team	What They Need to Know	Who Tells Them

DOES THIS REQUIRE DISCLOSURE

Decide first whether it does, and to whom. A fully internal workflow may not. Answer yes or no and say why.



DISCLOSURE LANGUAGE

Only if yes. Draft the actual sentence you'd say to the client, candidate, or employee.

KILL SWITCH

Who can stop this workflow, how fast, and how you undo what it already did?

- | | |
|---|--|
| ☐ Reviewer named as an individual | ☐ People impact considered for all three groups |
| ☐ Gate position justified | ☐ Disclosure decision made, yes or no |
| ☐ Confidentiality boundary defined | ☐ Who-needs-to-know list complete |
| ☐ Escalation path defined | ☐ Override and rollback path set |

The Agent Case

Only after the four Cs are complete. This is the evidence, not the verdict.

BASELINE VERSUS TARGET

Today: how long it takes, how often it runs, what it costs in hours. With the agent: the same three numbers.

WHAT THE AGENT DOES BETTER, AND BY WHAT MECHANISM

Name the property of the work that favors an agent. Volume, repetition, consistency across runs, speed, availability outside working hours, and tolerance for tedium. "It would be faster" is a claim. The mechanism is the argument.



WHAT THE HUMAN DOES WITH THE TIME

The freed hours go where? Name the higher-judgment work this unlocks. If the time doesn't get reinvested, the case is weaker than it looks.

WHAT YOU GIVE UP

What does the person currently bring that the agent won't? Empathy, relationship, judgment on edge cases, and institutional memory. Is that loss acceptable, and how is it covered?

COST OF BEING WRONG

If a bad output clears the gate, what's the damage, and is it reversible?

Before You Build

The build decision isn't made on this sheet alone. Weigh it against everything else on your roadmap. Finish these three fields before you move forward.

Disqualifiers

Check any that apply. A checked box doesn't kill the idea; it tells you what has to change first.

- ☐ No individual is accountable for the output.
- ☐ The data needs cleanup before an agent could read it reliably.
- ☐ The workflow changes shape most times it runs.
- ☐ The value of the step is the human relationship or the empathy in it.
- ☐ An error could reach someone before anyone reviews it.
- ☐ You can't describe a good output precisely enough to grade one.

THE CASE, ONE MORE TIME

Having worked the whole sheet, including the disqualifiers, make the case once more. Why should this be an agent build rather than work that stays with a person? If your own answer changed while filling this in, say that. It's the most useful thing on the page.



WHAT YOU COULD FIX WITHOUT AN AGENT

What surfaced along the way that would improve this workflow with no agent at all. A template, a checklist, a form, a deleted step, a clearer owner, a fixed automation. Answer this even if you're confident about the agent build. These are worth doing either way.

WHERE YOU NEED MORE INFORMATION

What became glaringly obvious that you don't know yet. Just name the gaps.

Completed by

Name	Role	Date	Workflow Name

Want help running this on a real workflow?

Speak with a Workforce Orchestration Expert at Lever Talent today! [Learn more.](#)