



AGENT JOB CARD · TEN LINES

# The AI Agent *Job Card*

An agent is a delegated role, not a feature. Do not switch one on until every line below is filled. The gaps you cannot complete are precisely your exposure.

<b>1 PURPOSE</b> One job, in a sentence. If it needs a paragraph, the scope is too wide. _____	
<b>2 TASKS</b> The specific things it does. Anything not listed is out of scope. _____ _____	<b>3 TOOLS AND DATA</b> What it can reach. Read access and write access are different jobs. _____ _____
<b>4 PERMITTED ACTIONS</b> What it may do on its own. Keep this short and honest. _____ _____	<b>5 ACTIONS REQUIRING APPROVAL</b> Anything irreversible, costly, or affecting a person. _____ _____
<b>6 QUALITY STANDARD</b> What "good" output is, and how it is checked. _____ _____	<b>7 ESCALATION CONDITIONS</b> When it must stop and hand to a human. _____ _____
<b>8 LOGGING</b> What it records, so its actions can be reviewed. _____	<b>9 HUMAN OWNER</b> One named person. Not a team. _____
<b>10 SHUTDOWN PROCEDURE</b> How to turn it off quickly, and who can. If you cannot say, you are not ready to start. _____	



COMPLETED EXAMPLE · CV-TRIAGE AGENT

# A recruitment agency's *trriage agent*

The same agent, defined. Same technology as an undefined one, completely different risk.

<p><b>1 PURPOSE</b></p> <p><i>Triage and summarise new inbound job applications.</i></p>	
<p><b>2 TASKS</b></p> <p><i>Read each new CV; produce a structured summary against the role's must-haves.</i></p>	<p><b>3 TOOLS AND DATA</b></p> <p><i>Read-only access to the applications inbox and the role spec. No write access to the ATS.</i></p>
<p><b>4 PERMITTED ACTIONS</b></p> <p><i>Draft a summary and a suggested shortlist tier. Nothing leaves the system.</i></p>	<p><b>5 ACTIONS REQUIRING APPROVAL</b></p> <p><i>Any contact with a candidate. Any rejection. Any write to the ATS.</i></p>
<p><b>6 QUALITY STANDARD</b></p> <p><i>Covers the must-haves, flags gaps, invents no detail. Spot-checked daily by the owner.</i></p>	<p><b>7 ESCALATION CONDITIONS</b></p> <p><i>Ambiguous CVs, anything borderline, anything outside the role spec.</i></p>
<p><b>8 LOGGING</b></p> <p><i>Every summary, with its source CV and timestamp, kept for audit.</i></p>	<p><b>9 HUMAN OWNER</b></p> <p><i>Priya Shah, senior consultant.</i></p>
<p><b>10 SHUTDOWN PROCEDURE</b></p> <p><i>Owner or ops lead disables the workflow in one click. Default to off if anything looks wrong.</i></p>	

**Illustrative example, not a specific agency.** Notice the pattern: reading is permitted, anything that touches a person or is hard to undo needs approval, and one named human owns it with a way to stop it.

[Discuss AI governance](#)

[goodtransformer.ai/services/ai-advisory-for-teams](https://goodtransformer.ai/services/ai-advisory-for-teams) · [hello@goodtransformer.ai](mailto:hello@goodtransformer.ai)