

Agent Accountability Checklist

For Educators - Nonprofits - Independent Professionals

Five plain-language checkpoints for anyone using AI agents in work that serves students, clients, or communities. No technical background required.

AI agents can handle significant routine work. For organisations that serve students, clients, or community members, the stakes of getting it wrong are higher. These five checkpoints cover the moments where a human being needs to be present and deliberate. Three minutes to review. Seconds for something to go wrong without it.

CHECKPOINT 01 - ENTRY — Before you set an agent going

■ Is the task clearly defined?

The agent should be able to complete it without guessing what you want. If you cannot describe the outcome in one sentence, the task is not ready.

■ [pertinent.com/audit.php](#) — *Agent Stack Audit identifies which tasks are genuinely ready.*

■ Is the audience appropriate?

Student data, client records, and donor information carry legal and ethical obligations. Confirm that the agent will not access, process, or transmit protected information.

■ Have you chosen the right tool for this task?

Not every task is suitable for an agent. Tasks involving sensitive relationships, contextual judgment, or irreversible actions need a human.

■ [pertinent.com/scorecard.php](#) — *Readiness Scorecard helps you assess fit before committing.*

CHECKPOINT 02 - INSTRUCTION — What you tell the agent to do

■ Is your instruction specific enough to be executed faithfully?

Vague instructions produce vague outputs. 'Write a follow-up email to the parent' is vague. 'Write a brief, warm follow-up confirming next steps' is specific.

■ [pertinent.com/followup.php](#) — *Follow-up Email Agent helps you draft with precision.*

■ Have you named what the agent should NOT do?

Constraints matter as much as instructions. State what is out of scope, what tone to avoid, and what information should not be included.

■ Are you comfortable if this instruction is fulfilled exactly as written?

Read your own instruction as if you are the agent receiving it. If the literal execution would be wrong or harmful, revise before proceeding.

CHECKPOINT 03 - OVERRIDE — Your right to stop, redirect, or reject

■ Do you know how to stop the agent if something goes wrong?

Before any agent runs, know the off switch. This is not hypothetical.

■ Is there a human being responsible for the outcome?

AI agents do not carry accountability. The person who instructed the agent is accountable for what it produces. Name that person before deployment.

■ [pertinent.com/compliance.php](#) — *Governance tools help map accountability.*

■ Do the people affected know an agent is involved?

Students, clients, and community members have a reasonable expectation of transparency.

CHECKPOINT 04 · AUDIT — Reviewing the output before it becomes an action

■ Have you read the output before it goes out under your name?

AI agents hallucinate, misread context, and miss nuance. A 30-second review before sending, posting, or acting on any AI-generated output is not optional.

■ pertinent.com/auditor.php — *Output Auditor reviews any AI output in under a minute.*

■ Does the output reflect your organisation's voice and values?

Generic outputs carry no relationship. Review for tone, accuracy, and appropriateness to your specific audience before use.

■ Have you documented that this output was reviewed by a human?

For regulated organisations, audit trails matter. A simple record of who reviewed what and when is sufficient.

CHECKPOINT 05 · EXIT — After the agent has completed its task

■ Did the output match the intent?

Compare what you asked for with what you received. If the gap is significant, revise the instruction before using the agent again.

■ pertinent.com/audit.php — *Re-running the Audit after a poor result helps recalibrate.*

■ Was anything produced that should not be shared, stored, or acted upon?

Some outputs should be discarded. Review before assuming the result is usable.

■ What did you learn from this run?

One observation per use improves the next one. Note it. The person who keeps improving their instructions is the one who gets consistent results.

This checklist reflects the Human-to-Agent (H2A) protocol — the interface between human intent and agent execution. Designed to be reviewed in under three minutes. Adapt it to your context.