[Company name] AI agent policy

Why this exists. Some of our software now takes actions for us without a person pressing a button each time. This policy says what that software may do and who answers for it.

What counts as an agent. Any software that uses AI to decide on and take an action in our systems, or to send something outside the company, without a person approving each step.

The rules.

1. Every agent has a named owner. No owner, no agent. The owner of each agent is listed on the agent inventory kept by [role].
2. An agent goes on the inventory before it goes into use. That includes agents inside products we already pay for.
3. An agent acts under its own account, never under a person's.
4. An agent gets the least access that lets it do its job. The owner writes down what it may read, what it may change, and what it may send or spend.
5. An agent never changes its own access or instructions. Only its owner does, and each change is recorded.
6. A person approves before an agent does any of these: pays or moves more than [amount]; pays an account we have not paid before; sends a customer anything about [topics]; deletes records; [add your own].
7. What an agent does is recorded, and the record is kept for [period] where the agent cannot alter it.
8. [Types of data] never go to an agent unless the vendor contract allows it and [role] has approved it in writing.
9. A person who uses an agent's output is responsible for it, the same as for their own work.

When something goes wrong. Anyone who sees an agent causing harm may stop it. The stop step for each agent is on the inventory. Then tell [name or role] within [time]. Do not restart the agent until its owner says so.

Review. [Role] reviews this policy and the inventory every [six] months, and whenever an agent's job or access changes.

Approved by [name], [date].