



The 6-Question AI Usage Policy

Most companies need one page, not forty. Answer these six questions, sign it, post it — done.

1 · APPROVED TOOLS

Which AI tools and agents are sanctioned — and under which accounts?

Business accounts only. No personal-account AI for company work. List each approved tool by name.

2 · DATA RULES

What may never be entered into an agent?

Credentials · customer PII beyond need · badge data, codes, floor plans, emergency procedures · anything under NDA.

3 · HUMAN ACCOUNTABILITY

Who owns every output that leaves the building?

A named human reviewed it before release. "The AI wrote it" is never an accountability answer.

4 · DISCLOSURE

When do customers know they're talking to an agent?

Recommended default: always, for direct interaction. State your rule here.

5 · TIER BOUNDARIES

What may agents prepare but never execute?

Payments · deletions · contracts · mass sends · anything life-safety or legally binding. Agents draft; humans execute.

6 · INCIDENT PATH

What happens when an agent gets it wrong?

Stop → preserve the log → report to [named person]. No quiet cleanup — it destroys the audit trail and the lesson.

Want the full discipline — delegation, guardrails, evals, audit trails?

AI Agent Command(TM) · the agent management course by Kryos Access Systems · kryossecuritygroup.com