Mr. Alpaca *by Alpaca Data Lab*

ONE-PAGE HANDOUT · PRACTICAL FIELD GUIDE

AGENT WORKFLOW SAFETY & GOVERNANCE — FOR SMALL TEAMS

AI Agent Readiness Checklist for Small Firms

Before AI agents touch email, files, CRM, or client data, define these 10 controls.

AI agents are becoming capable enough to do real work. The problem is not capability. The problem is control. Use this checklist to see whether your team has the basic data boundaries, permission rules, and human approval gates in place before agents act inside real business workflows.

WHO THIS IS FOR

Small law firms, CPA firms, bookkeeping firms, consultants, agencies, and professional service teams using ChatGPT, Claude, Copilot, Gemini, Cursor, meeting bots, automation tools, or early AI agents.

The 10 Controls

SKIM IN 3–5 MINUTES

1 AI Tool Visibility

- ☐ We know which AI tools our team uses.
- ☐ We know whether staff use personal or company-managed AI accounts.
- ☐ We have an approved / restricted / prohibited AI tool list.

2 Data Boundaries

- ☐ We have defined what data may enter AI tools.
- ☐ Client, contract, financial, HR, legal/advisory, and confidential documents are classified.
- ☐ Passwords, API keys, credentials, private keys, raw customer records, and production access are prohibited.

3 Prompt Redaction

- ☐ Staff know what to remove before using AI.
- ☐ We have a written prompt redaction checklist.
- ☐ Redaction does not depend only on individual judgment.

4 Human Review

- ☐ AI-assisted client/customer-facing output requires human review.
- ☐ High-stakes output has a named reviewer.
- ☐ AI-generated legal, financial, advisory, HR, or technical output is not sent without qualified review.

5 Meeting Bots & Transcripts

- ☐ We know which meetings may be transcribed by AI.
- ☐ Sensitive client, legal, HR, financial, strategy, or dispute meetings are restricted.
- ☐ Staff know when not to invite an AI meeting bot.

6 Coding & Technical Exposure

- ☐ AI coding tools do not receive secrets, credentials, API keys, or production access details.
- ☐ Repository, log, environment file, and customer data exposure is controlled.
- ☐ Code generated with AI is reviewed before use.

7 Agent Permission Boundaries

- ☐ We have defined what agents can read.
- ☐ We have defined what agents can draft.
- ☐ We have defined what agents can write, send, update, delete, publish, or trigger — and which systems are off-limits.

8 Human Approval Gates

- ☐ Agents cannot send external emails without approval.
- ☐ Agents cannot update CRM or client records without approval.
- ☐ Agents cannot delete files, publish content, send invoices, modify code, or trigger workflows without approval.

9 Ownership & Logging

- ☐ Someone owns AI governance internally.
- ☐ Agent actions are logged or reviewable.
- ☐ There is an escalation path when AI is unsure or a workflow is sensitive.

10 Client Trust

- ☐ If a client asks how we use AI with their data, we have a clear answer.
- ☐ Our answer does not overpromise.
- ☐ Our AI usage rules are written, not just informal.

How to read this checklist

Not sure where your team stands?

If several boxes are unchecked or unclear, your team may not be ready to connect AI agents to real business systems yet. Start by defining tool visibility, data boundaries, prompt redaction, human review rules, and approval gates before agents can send, update, delete, publish, or trigger workflows.

Start with the Free AI Risk Self-Check, then review the fictiona report or request a Mini AI Data Hygiene Snapshot.

- Take the Free AI Risk Self-Check `https://alpacadatalab.com/ai-risk-self-check`
- View the Fictional Sample Report `https://www.alpacadatalab.com/assets/northbridge-ai-data-hygiene-snapshot.pdf`
- Get a Mini AI Data Hygiene Snapshot `https://alpacadatalab.com/#mini-snapshot`

About the Mini AI Data Hygiene Snapshot — an entry diagnostic for Agent Workflow Safety & Governance. It helps small teams identify current AI tool usage, unmanaged account risk, sensitive data exposure, missing redaction rules, human review gaps, and early agent workflow readiness.

Mr. Alpaca
Alpaca Data Lab · v1.0