The Next Standard: *Practical AI Steps for CPAs*

By: Marc Hamilton, CPA CGMA CFO & Treasurer Three Notch Group, Inc.

The accounting profession is in a pivotal transition. Artificial intelligence (AI) is no longer theoretical—it is embedded in the daily platforms CPAs already use. From Microsoft 365 to Adobe Acrobat to Enterprise Resource Planning (ERP) tools, the shift is not optional. CPAs in public practice, business and industry, government, and nonprofit sectors must now grapple with how to responsibly harness AI without undermining the very assurance their profession is built on.

The profession has always adapted to new standards: from paper to spreadsheets, from manual reconciliations to ERP systems. AI is the next standard. The risk is not adoption itself but failing to adopt with professional guardrails. Research shows that technology-mature firms produce up to 39% more revenue per employee—an unmistakable signal that modernization is now tied directly to firm economics. The critical question is: how do CPAs embrace these tools while ensuring independence, objectivity, and trust remain intact?

What's Changed Recently: Adoption Guidance

AI is already in the tools we use every day. The question is how to move forward with confidence. Start with the short list below: it's the most current, credible guidance for turning curiosity into governed practice—policies you can defend, workpapers you can audit, and outcomes you can improve.

Responsible AI Implementation Checklist (aligned with ISO/IEC 42001)
Why it matters: Turns "AI governance" into a CPA-ready control set (roles, human oversight, data classification, documentation) mapped to a globally recognized standard.
Use now: Make this your firm's AI control baseline; require an **AI workpaper** (prompt, model/version, inputs, validations, reviewer sign-off) for any AI-assisted deliverable; map your policy to the checklist headings. AICPA & CIMA

AI Risk Management Framework (AI RMF 1.0) + Playbook
Why it matters: The U.S. government's voluntary framework widely adopted across industries. Its **Govern → Map → Measure → Manage** structure is practical and auditable. The Playbook provides copy-ready actions.
Use now: Treat any new agent or automation as an **RMF deployment** with defined risks, success criteria, monitoring, and rollback. Tie your AI workpaper to RMF outcomes (e.g., validation evidence under "Measure"). NIST

AI Management System (AIMS)
Why it matters: The first management-system standard for AI—familiar to CPAs who use ISO-style quality systems. It gives a certifiable structure (policy, objectives, controls, continual improvement) for responsible AI use.
Use now: Borrow the clause structure to formalize your AI policy (scope, leadership, planning, support, operations, performance evaluation, improvement). For larger firms, consider certification to signal maturity to clients and insurers. ISO

"Spotlight: Generative AI in Audits and Financial Reporting" (July 2024)
Why it matters: Not a standard—but a clear signal of regulator expectations: documentation, supervision, and professional skepticism when AI aids audit or financial-reporting work.
Use now: In audit/attest, label **preview** tools as pilot-only, keep **human-in-the-loop**, and retain prompts/outputs/tie-outs in workpapers; ensure partner-level supervision over any AI-assisted procedures. PCAOB

Risk Management Considerations

Insurers are actively recalibrating how AI risk is treated across **Cyber, Tech/Professional E&O, Media, EPL, and D&O coverages**—and the fine print matters. In 2024, **W. R. Berkley** introduced an **"Artificial Intelligence Exclusion — Absolute"** endorsement for certain liability lines; its broad "arising out of" language can sweep in losses linked to the *use, development, or deployment* of AI (even alleged use), so renewals should be reviewed for similar forms. At the same time, major brokers emphasize that most AI events still flow through **existing policies**, but with growing reliance on **exclusions** (or, in some markets, **affirmative AI endorsements**) and tighter underwriting questions about how firms govern, validate, and document AI-assisted work. In other words, the market is moving from "silent AI" (ambiguity) toward explicit terms—good for clarity, risky if you don't read the endorsements. Hunton NLR Marsh WTW

For CPA firms, the play is to **treat AI like any material operational risk**: (1) **Pre-read forms and endorsements** at renewal for AI-related language across all lines, watching for "absolute" AI exclusions or broad "actual or alleged use" phrasing; (2) arrive "underwriter-ready" with an **AI inventory**, your **policy spine** (e.g., AICPA/ISO/NIST mapping), and sample **AI workpapers** that show human-in-the-loop review, validation/tie-outs, logging, and rollback plans; (3) align **engagement letters** and client communications with how you actually use AI (and what you do *not* do), since brokers report questionnaires now probe disclosures and controls; and (4) map likely scenarios (model error, data leakage, IP/media claims, bias/discrimination) to the policies that would respond and the exclusions that might limit recovery. The goal is simple: **no surprises at claim time**—clear coverage intent, documented controls, and evidence you ran AI in a defensible, supervised way. Marsh WTW

AI Tools: A Rapidly Evolving Landscape

If your company or firm runs on **Microsoft 365**, AI now lives inside the tools you already open every day—**Word, Excel, Outlook, and Teams**—as **Microsoft 365 Copilot**. In practice, Copilot drafts and rewrites text in Word and Outlook, summarizes long threads and documents, and answers questions using your Microsoft 365 SharePoint content (subject to your IT permissions). In Excel, Copilot helps you explore a workbook in natural language, propose formulas, and generate explanations or summaries for the ranges you select. For production work, always confirm whether a feature is **General Availability (GA)** or **preview** using Microsoft's official release notes; train staff on GA only and treat preview as pilot-only with added review controls. Microsoft

For **PDF-heavy** work, **Adobe Acrobat AI Assistant** brings summarization and data extraction into Acrobat. Companies can enable admin controls and review the product's security posture (processing/caching behavior, model providers) before allowing client data. This keeps PDF review inside governed tools rather than ad-hoc uploads. Adobe

Many firms also use **research assistants** for external lookups and drafting. If you choose to deploy them, prefer **enterprise** offerings with contractual data protections—e.g., **ChatGPT Enterprise** (admin controls, SSO/SCIM, no training on your business data by default) and **Perplexity Enterprise Pro** (answers with **cited** sources, enterprise controls). Use these for research and first drafts; keep client-identifiable data within governed environments. OpenAI Perplexity

From "AI That Writes" to "AI That Helps You Work"

Most of us first met AI as a **good drafter**—it helps write emails, tidy memos, and summarize long threads. The next step is **AI that helps your work move**, still under your approval. This is often called **agentic AI** (*supervised, task-capable AI that only acts within guardrails you set*): the AI proposes sensible next steps, your workflow tool prepares them, and **you** decide whether anything is actually sent or posted.

In Microsoft 365, that looks like a **Power Automate** flow—"when an invoice arrives, save the PDF, extract key fields, draft a reply, post a summary to Teams"—with a **pause for approval** before any change to a system. When an update is needed, it goes through an **API** (*Application Programming Interface— a secure, rules-based doorway that lets systems exchange data or perform actions with permissions and audit logs*). If nothing is approved, nothing happens; every step is logged.

This approach fits CPA work because it reuses controls you already value—**review before release, clear evidence, and reversibility**. Each run leaves a short **AI workpaper** (inputs, model/version, actions taken, outputs, validations/tie-outs, preparer/reviewer, dates). As features evolve, treat **GA** (generally available) capabilities as production-ready and **preview** as pilot-only with added review.

If you're new to this, the on-ramp is simple: pick one routine task, add a single **supervised** step (e.g., have the AI draft the message you already send), attach the workpaper, and measure minutes saved. When your comfortable, add one **approved API** action (for example, create a "to-review" item—not a journal entry). The goal isn't to automate everything at once; it's to **learn in small, safe steps** that compound over time, keeping assurance intact while curiosity grows.

Curiosity Prompts

Quick wins in daily work

- What **standard memo or email** (client update, variance note, engagement recap) could be safely AI-drafted if a reviewer signs off?
- If Copilot could turn **meeting notes** into next-step checklists, where would that save you the most time?

Close and reconciliations

- What would a **"to-review" queue** look like if an API could add invoice candidates but never post to the ledger?
- Which **data** in your reconciliation workbook always lead to exceptions—and how could AI pre-sort or identify them?

Supervised automations (agentic flows)

- If a Power Automate flow **paused for approval** before any send or write-back, which task would you automate first?
- Where would a **"draft → approve → send"** pattern prevent rework or mis-sends in your team?

Documentation & evidence

- What **two screenshots** do you take every month that an API log could replace with better evidence?
- How could you show, in one place, the **prompt/inputs, model/version, output, and reviewer** for an AI-assisted task?

Client service & advisory

- What **client question** do you answer repeatedly that Copilot could draft—so you focus on tailoring the advice?
- If Acrobat's AI could extract key terms from a **50-page contract**, which fields would you want verified first?

References (APA 7th)

1. AICPA & CIMA. (2025, July 10). Responsible AI implementation checklist (aligned to ISO/IEC 42001). https://www.aicpa-cima.com/news/article/responsible-ai-implementation-checklist-with-iso-iec-42001-alignment
2. ISO/IEC. (2023). ISO/IEC 42001:2023 — Artificial intelligence — Management system. https://www.iso.org/standard/42001
3. NIST. (2023). AI Risk Management Framework (1.0) and Playbook. https://www.nist.gov/itl/ai-risk-management-framework
4. Microsoft. (2025). Release notes for Microsoft 365 Copilot. https://learn.microsoft.com/en-us/copilot/microsoft-365/release-notes
5. Microsoft Learn. (2025). Financial Reconciliation agent (preview). https://learn.microsoft.com/en-us/copilot/finance/reconcile/reconcile-overview
6. Adobe. (2025). Administer Acrobat AI Assistant for enterprises & Security fact sheet. https://helpx.adobe.com/acrobat/using/administer-acrobat-ai-assistant-for-enterprises.html
7. IRS. (2024). Publication 4557: Safeguarding taxpayer data. https://www.irs.gov/pub/irs-pdf/p4557.pdf
8. PCAOB. (2024). Spotlight: Generative AI in audits & financial reporting. https://pcaobus.org/documents/generative-ai-spotlight.pdf
9. Rightworks. (2024–2025). Technology-mature firms earn up to 39% more revenue per employee. https://www.rightworks.com/blog/accounting-technology-firms-earn-more-revenue/

AI Assistance Disclosure:

Portions of this article were drafted and edited with the aid of OpenAI GPT-5; all facts, interpretations, and conclusions were reviewed and approved by the author.