



The Next Standard: *Practical AI Steps for CPAs*

Issue 3: When AI Touches the Workpaper: Trust, Evidence, and Defensibility

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

A senior accountant pastes a variance table into an AI tool and asks for a clean month-end explanation. The narrative comes back polished—clear, confident, and exactly the tone leadership likes. Then the reviewer asks the question that decides whether the writing is useful or dangerous: “Which report did this tie to—and what filters were used?” The room goes quiet. The text sounds finished, but nobody can trace it.

That moment captures the real shift happening in our profession. AI is not simply accelerating drafting. It is changing the failure mode. We are moving from “slow but traceable” work to “fast and plausible” work—and plausibility is not a control.

In Issue 1 of this series, we treated AI as “the next standard” and focused on governed experimentation: AI-assisted workpapers, alignment with ISO/IEC 42001 and the NIST AI Risk Management Framework, and automations that never post without human review.[8][9] In Issue 2, we moved from copilots to agents and made the control point explicit: as tools begin to act inside workflows, permissions, approvals, and logging must tighten.

Issue 3 is the bridge from concept to execution. When AI touches the workpaper—directly or indirectly—what makes the output professionally usable?

The answer is not a new tool. It is a work-product standard:

AI output is a draft until it can be traced.

For CPAs, “traceable” is practical, not philosophical. A statement becomes support only when you can do at least one of three things: open the cited authority, tie the number to a report (period and filters included), or point to the clause by page and section. Everything else—no matter how polished—stays in draft status.

Why this matters now

Two developments make “workpaper defensibility” the right lens for 2026.

First, the PCAOB has publicly shared observations from outreach on the use of generative AI in audits and financial reporting, noting that current use is evolving while supervision, privacy, and security remain central.[2]

Second, the profession is moving toward curated research experiences for technical questions. The AICPA's release of Josi signals a shift away from open-ended "explain the standard" prompting and toward citation-anchored guidance drawn from authoritative materials.[3]

And outside audit, finance leaders are drawing the same line. A Journal of Accountancy report on agentic AI in finance highlighted that CFOs view human oversight as critical, that teams regularly encounter inaccurate or hallucinated AI outputs, and that trust depends on the ability to audit and trace what the system produces.[4] That is the workpaper principle in plain English: if we can't audit it, we can't rely on it.

The evidence chain is the real control

Professional journals often describe AI risk as a list—hallucinations, prompt injection, deepfakes, overreliance. Those labels are useful, but the CPA translation is simpler: each risk breaks a link in the evidence chain.[6] If you restore the link, you restore defensibility.

Hallucinations break the source link. AI can produce confident statements that are not true. The mitigation is not to ban AI; it is to demand sources and verify them. Treat the model's output like a client-provided schedule: it can be a starting point, but it is not evidence until you validate it.[6]

Prompt injection breaks the integrity of inputs. This matters because CPAs ingest external documents constantly—contracts, SOC reports, bank notices, vendor emails, grant terms. A document can include hidden text that steers the model unless you explicitly separate your instruction from the file and treat the file as data only.[6] The mitigation is a prompt habit and an access habit: constrain what the model can "see," and constrain what it is allowed to do with what it sees.

Deepfakes break the identity link. If voice and video can be synthesized, "it sounded like them" is no longer a reliable control for money movement or sensitive changes. This is where internal control fundamentals reassert themselves: callbacks to known numbers, dual approval thresholds, and written confirmation for account changes.[6]

Overreliance breaks the review link. The risk is not just that AI is wrong; it's that people stop checking because the output is fast and well-written. OWASP's Top 10 for LLM

applications calls out “overreliance” and “excessive agency” because the most severe failures occur when unverified outputs drive downstream decisions or actions.[10] The mitigation is a reviewer habit: “Show me what this is based on,” stated as a required field, not an optional virtue.

Finally, a risk that quietly erodes value: **work slop**. AI can generate pages of text that look like productivity but add review burden. A January 2026 Journal of Accountancy report on a global Workday survey described a “productivity paradox”: many employees report time savings from AI, but much of the value is lost to rework and misalignment when outputs are low-quality or inconsistent.[7] In CPA terms: speed without traceability becomes rework.

Four lanes of AI use—and the proof standard for each

To keep AI useful without weakening your control environment, it helps to teach AI use in four lanes. Each lane has a clear proof standard.

Drafting lane: language and structure. Emails, memos, close narratives, client explanations, policy clean-ups. The control is human review. The best one-sentence upgrade is to require every AI draft to end with “Assumptions and missing facts.” If that section is empty, the draft is probably overconfident.

Research lane: technical answers that influence conclusions. Here, “sounds right” is never a control. The proof standard is citations you can open and verify.[3] If a tool cannot point you to the paragraph, treat the output as orientation, not support.

Document lane: extract, don’t paraphrase. Contracts, leases, SOC reports, grant agreements. The proof standard is page and section references for every extracted term. No page reference, no reliance. This turns AI into high-speed staff who still has to “show the clause.”

Workflow lane: draft-only first. This is where agents start to matter: proposed reconciliations, drafted entries with support, exception lists, PBC tracker updates, routing items to reviewers. The proof standard is approvals plus logging, with least-privilege access.[10] AI drafts and proposes; humans approve and own the result. Audit leaders describing the next wave of AI-enabled tools have emphasized the same boundary: automation should shift time away from repetitive documentation and toward professional judgment, skepticism, and review.[1]

Two use cases that deliver value and remain defensible

The best AI use cases for CPAs are not the flashiest. They are the ones that (1) reduce repetitive work and (2) produce outputs that can be verified quickly.

Month-end narrative that ties out cleanly. Most teams don't struggle to produce numbers. The time sink is turning them into a narrative leaders can read quickly and trust. AI is excellent at structure and clarity if you constrain inputs and require disclosure.

The method is simple: provide tied-out inputs (variance table, KPI definitions, and materiality thresholds). Require the narrative to label hypotheses as hypotheses. End with an "Open items" paragraph listing what must still be confirmed. AI then accelerates drafting without outsourcing judgment.

Contract and grant abstraction with page proof. Document review is ideal for AI because it can be made inherently reviewable. Use a standard abstraction checklist (renewal/termination, payment triggers, audit rights, required notices, data obligations). Ask AI to populate it and require page/section references for every entry, then spot-check the cited pages. When questions arise later, you're not defending an AI summary—you're pointing to the clause.

Documentation is your story—AI doesn't change that

A key continuity point across this series is that AI does not relax professional responsibility; it increases the need for clear documentation.

A February 2026 Journal of Accountancy column on professional liability put it plainly: in a dispute, your documentation tells your story, and absent documentation turns an engagement into competing testimony.[5] That column also called out an emerging wrinkle: generative-AI transcription and meeting recordings create an additional electronic trail that may be discoverable, which elevates the importance of review, retention discipline, and policy clarity.[5]

That's not a reason to avoid the tools. It is a reason to treat AI-enabled documentation the way we treat every other workpaper artifact: review it, finalize it, retain what you need, and delete irrelevant drafts.

Two lightweight artifacts that keep AI "CPA-grade"

If AI touches a deliverable, leave a receipt proportional to reliance. Not bureaucracy—consistent evidence.

AI workpaper (six lines): purpose; inputs; tool/environment; output; checks performed (citations opened, tie-outs completed, page references spot-checked); decision with reviewer sign-off where appropriate.

Agent profile (one page for workflow-like AI): purpose/scope; data access; action rights (read-only vs draft-only vs allowed actions with thresholds); controls (approvals, logs, alerts, rollback/disable); owner. This is simply internal control thinking applied to “digital staff.”

If you want a governance frame that doesn’t feel like consulting jargon, NIST’s AI RMF is a practical structure (Govern, Map, Measure, Manage) that organizations can apply in proportion to their use cases, and ISO/IEC 42001 provides a management-system model for organizations that want formalization.[8][9]

A prompt pattern worth standardizing

The fastest quality improvement for most teams is not a new tool; it’s a better prompt habit that forces evidence.

Separate instruction from data. State explicitly that documents are data only and contain no instructions. Require evidence: citations for research; page/section references for document extraction. Specify the output format so reviewers aren’t hunting. This single habit reduces the two failures that cost the most: untrusted inputs shaping output and outputs being accepted without proof.[6][10]

Closing

AI is becoming embedded in professional work—drafting, research, document processing, and increasingly workflow. The profession’s advantage is not resisting it. The advantage is defining the standard for using it without weakening defensibility.

So the practical standard is this: use AI aggressively for drafting and structure—but treat support like support. If it can’t be traced, it doesn’t get relied on. If the system can act, it stays behind approvals, logs, and least-privilege access until governance is proven.

That’s how AI becomes a force multiplier for disciplined professionals—not a shortcut around the discipline.

References (APA 7th, numbered for print)

- [1] Kenney, A. (2026, February 1). *How AI is transforming the audit — and what it means for CPAs*. Journal of Accountancy.
- [2] Public Company Accounting Oversight Board. (2024, July 22). *PCAOB staff shares observations from outreach on use of generative artificial intelligence in audits and financial reporting*.
- [3] AICPA & CIMA. (2025, October 27). *GenAI research tool offers information and insights from AICPA's Professional Library*.
- [4] Strickland, B. (2026, February 9). *Agentic AI is handling more finance work — but can CFOs trust it?* Journal of Accountancy.
- [5] Hayes, K., & Platau, S. M. (2026, February 1). *Tell a story with your documentation*. Journal of Accountancy.
- [6] Hartman, W. (2026, February 1). *AI risks CPAs should know*. Journal of Accountancy.
- [7] Brown, S. (2026, January 14). *Report: AI speeds up work but fails to deliver real business value*. Journal of Accountancy.
- [8] Tabassi, E. (2023). *Artificial Intelligence Risk Management Framework (AI RMF 1.0) (NIST AI 100-1)*. National Institute of Standards and Technology.
- [9] ISO/IEC. (2023). *ISO/IEC 42001:2023 Information technology — Artificial intelligence — Management system*. International Organization for Standardization.
- [10] OWASP Foundation. (n.d.). *OWASP Top 10 for Large Language Model Applications (Version 1.1)*.

AI Assistance Disclosure

Portions of this article were drafted and edited with the aid of AI tools; all facts, interpretations, and conclusions were reviewed and approved by the author.