

The Next Standard: *Practical AI Steps for CPAs*

Issue 2: From Copilots to Agents in Your CPA Tech Stack

By: Marc Hamilton, CPA CGMA, CFO & Treasurer

Three Notch Group, Inc.

The accounting profession is rapidly entering yet a new stage with artificial intelligence (AI). In the first article in this series, we treated AI as “the next standard” and focused on safe, governed experimentation: AI-assisted workpapers, alignment with the ISO/IEC 42001 AI Management System standard (ISO/IEC 42001) and the NIST AI Risk Management Framework (NIST AI RMF), and automations that never post without human review. The aim was to learn the new tools without sacrificing the discipline that defines good accounting.

Over 2025 and into 2026, the story moves from “copilots” to “agents.” AI is no longer just helping us draft emails, perform research or summarize Teams threads. It is starting to act inside the systems we use to record, reconcile, and report. Microsoft 365 Copilot is gaining agents with their own identities and permissions. Microsoft Copilot Studio can click through older applications that never exposed an application programming interface (API). ERP and procurement systems are adding embedded copilots. In short, the tools we already use are beginning to act on our behalf.

For CPAs, the question is not whether this will happen, but how it will happen in our environments—and whether we stay in front of the control conversation.

From Copilot to Agent

Most of us have used generative AI in its copilot form: we ask for a draft, it produces one, and we decide what to do next. The initiative stays with the “human in the loop”.

Agentic AI changes that. An agent is designed to notice what is happening, decide what to do next, and carry out steps across multiple systems. In practical terms, an agent might watch an invoice inbox or ERP postings, use rules or goals to decide which items need attention, call tools to move data or start workflows, and adjust over time based on feedback. Instead of you starting every step, the system can begin moving work on its own.

Think of this as a spectrum. At one end, assistive AI simply suggests text or analysis. In the middle, semi-agentic tools can run multi-step flows, but you still

click “send” or “post.” At the far end, fully agentic systems can initiate and complete tasks when set conditions are met. As you move along that spectrum, governance, approvals, and logging must tighten. The governance ideas from Issue 1 still apply here—agentic tools simply demand more of them.

APIs in Plain Accounting Language

Many of the valuable agent behaviors depend on integration.

An API is a structured way for one system to ask another system to get or update data. Instead of reading the screen the way a person would, an API call might say, “Give me all invoices for Vendor X this month,” or “Create a draft journal entry with these fields.” Those calls are permissioned, documented, and loggable. Because the calls are structured, you can apply familiar controls—least-privilege access, change management around how tools are configured, and audit trails that show who or what did what, and when. Screen-clicking “computer use” has its place, but for higher-risk finance workflows APIs will usually be the safer long-term path.

Where Agents Are Appearing

You do not have to buy a separate AI platform to encounter agents. They are arriving inside tools you already have licensed. In Microsoft 365, agents for Copilot can be set up to watch a shared accounts payable inbox, extract key data, draft emails requesting missing information, and update a tracker so the team has a current view of status. ERP and procurement vendors are introducing features that draft journal entries, suggest variance explanations, flag unusual spend patterns, and propose reconciliation actions. Across all of these, the marketing label matters far less than the behavior. What can the agent see? What can it change? Under whose identity does it act? How do we review and override what it does?

Before you enable agent features from any vendor, it is worth sitting down with your vendors and IT teams and asking a few basic questions: Under what identity these agents will run, and how will access be restricted or controlled? Which systems will they touch, and will they use APIs or screen-driven “computer use”? What logs, alerts, and reports will be available to finance and internal audit? And finally, how will usage be priced—by user, by run, or by capacity—so you do not discover later that “experimental agents” have quietly created a new cost center.

Treating Agents as Staff

In the first issue, we introduced the idea of an AI workpaper—a short record of what a model did for you, what data it used, and how you checked the result. As agents show up in finance workflows, a natural next step is to treat them much like staff members.

Each significant agent should have a simple profile that answers four questions. First, purpose: what is it hired to do—draft AP batches, prepare variance narratives, maintain a prepared-by-client (PBC) status tracker? Second, data: which systems and data sets can it see—general ledger and AP detail, but not payroll, HR, or legal files? Third, action rights: is it read-only, draft-only, or allowed to post within defined thresholds? Fourth, ownership: who approves changes, reviews logs, and can disable or roll it back if something goes wrong? One page per agent is enough, and over time these profiles start to look like an org chart for your digital staff.

Linking to Emerging Frameworks

This “agent as staff member” approach fits the direction of AI governance. Under ISO/IEC 42001, each material agent can be treated as its own AI use case within your AI Management System. Under the NIST AI RMF and its Generative AI Profile (NIST AI 600-1), you can map agents to the familiar actions of Map, Measure, and Manage. Under the European Union Artificial Intelligence Act (EU AI Act), AI systems that influence material decisions are expected to operate with documented risk management, human oversight, and event logging. Issue 1 laid the groundwork; here, agent profiles become concrete evidence that specific automations are governed.

Practical Patterns You Can Use

You do not need a grand roadmap to begin. A few focused patterns can demonstrate value while keeping risk contained.

Month-end drafting is a good starting point. Recurring accruals, allocations, and reclassifications are prime candidates for agent assistance. An agent can look at historical patterns, current-period data, and rules to draft proposed entries. Those drafts then flow into your existing approval process. The boundary is clear: the agent drafts; humans post.

Accounts payable intake is another. Much of AP work involves handling documents and getting them into the right workflow. An agent can monitor a shared inbox, classify documents, extract key fields, suggest coding, and flag potential duplicates or unusual amounts. Approvers still own the decisions, and high-risk vendors or categories can be excluded altogether. The same pattern can support prepared-by-client tracking by updating a simple list as documents arrive and drafting concise status updates.

Measuring and Piloting in 90 Days

Agent projects fail when they are driven by novelty rather than measurable benefit. Before introducing an agent, capture a simple baseline: how long the process currently takes (days to close, days sales outstanding, days to reconcile); how much effort it consumes (hours per period, rework rate); and

what the common error signals look like (post-close adjustments, control exceptions, client complaints).

Once the agent is in place, track a small set of metrics. Time saved per cycle is one. How often the agent's drafts are accepted, lightly edited, or heavily rewritten is another. The number and severity of incidents that required rollback or correction, and the level of adoption across teams and entities, are equally important. If you cannot point to a meaningful improvement in speed, quality, or staff capacity after a couple of cycles, it may be time to refine or retire the agent.

You can fit this into a 90-day sequence. In the first month, choose one or two modest use cases and draft agent profiles with input from IT, internal audit, and information security. In the second month, build and test in a limited environment, including at least one "shadow" cycle where the agent proposes actions but people still perform the work. In the third month, pilot with guardrails for a small group of entities, require human approval for all material actions, and review the agreed metrics weekly. At the end of the 90 days, make a clear decision: continue, adjust, or stop.

Talent, Standards, and the Next Standard

Agentic AI will not remove the need for CPAs. It will, however, move some of the routine workload away from manual drafting and tracking toward supervision, exception handling, and explanation. For staff and seniors, that means less time entering the same types of entries month after month and more time understanding why something does not look right, how to explain a pattern to management, or how to improve a workflow.

Firms can lean into that shift by incorporating agent supervision and AI workpapers into staff rotations and training, using agent profiles as teaching tools, and encouraging people closest to the work to identify small, safe opportunities for agent support with clear stop-rules. At the same time, the standards landscape will continue to evolve, and regulators and standard-setters are already exploring how AI affects evidence, documentation, and risk.

In many ways, the arrival of agents brings us back to themes that have always defined good accounting: clarity of purpose, discipline around access, thoughtful division of duties, and careful documentation of what has been done and why. The tools are changing. The principles are not. For now, the opportunity is with the teams that are willing to move beyond experiments and pilots that live only in demonstrations and instead fold AI into their work in quiet, well-governed ways—one process at a time, with clear metrics and clear owners.


References (APA 7th)

- AICPA & CIMA. (2025, July 11). *Responsible AI implementation checklist (with ISO/IEC 42001 alignment)*. Association of International Certified Professional Accountants. <https://www.aicpa-cima.com/news/article/responsible-ai-implementation-checklist-with-iso-iec-42001-alignment>
- Amazon Web Services. (n.d.). *Unlock new value with agentic AI on AWS*. Amazon Web Services. <https://aws.amazon.com/ai/agent-ai/>
- Atlassian. (2025, May 23). *What are AI agentic workflows?* Atlassian Work Life. <https://www.atlassian.com/blog/artificial-intelligence/ai-agent-workflows>
- Google Cloud. (n.d.). *What is agentic AI? Definition and differentiators*. Google Cloud. <https://cloud.google.com/discover/what-is-agent-ai>
- IBM. (2024). *What are agentic workflows?* IBM Think. <https://www.ibm.com/think/topics/agent-workflows>
- Intuit Inc. (2024, November 20). *Intuit launches AI-powered Intuit Assist for QuickBooks, giving millions of businesses a competitive edge*. Intuit Investor Relations. <https://investors.intuit.com/news-events/press-releases/detail/1222/intuit-launches-ai-powered-intuit-assist-for-quickbooks-giving-millions-of-businesses-a-competitive-edge>
- Kachwala, Z. (2025, June 25). Over 40% of agentic AI projects will be scrapped by 2027, Gartner says. *Reuters*. <https://www.reuters.com/business/over-40-agent-ai-projects-will-be-scrapped-by-2027-gartner-says-2025-06-25/>
- Microsoft. (n.d.). *Microsoft Copilot Studio documentation*. Microsoft Learn. <https://learn.microsoft.com/en-us/microsoft-copilot-studio/>
- Raghavan, S. (2025, September 25). *Empower your workforce with agents in Microsoft 365 Copilot*. Microsoft 365 Blog. <https://www.microsoft.com/en-us/microsoft-365/blog/2025/09/25/empower-your-workforce-with-agents-in-microsoft-365-copilot/>
- SAP SE. (2024, October 8). *SAP supercharges copilot Joule with collaborative capabilities to ignite enterprise AI revolution*. SAP News Center. <https://news.sap.com/2024/10/sap-teched-copilot-joule-collaborative-capabilities-enterprise-ai/>
- Tabassi, E. (2023). *Artificial intelligence risk management framework (AI RMF 1.0)* (NIST AI 100-1). National Institute of Standards and Technology. <https://doi.org/10.6028/NIST.AI.100-1>
- TechRadar Pro. (2025, October 27). Experts warn Microsoft Copilot Studio agents are being hijacked to steal OAuth tokens. *TechRadar Pro*. <https://www.techradar.com/pro/security/experts-warn-microsoft-copilot-studio-agents-are-being-hijacked-to-steal-oauth-tokens>
-

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.

© 2026 Marc Hamilton, CPA, CGMA. Published in Alabama Society of CPA's Connections, Q1 2026. Reprinted or distributed by other Professional Accounting Organizations with permission.



The Next Standard: *Practical AI Steps for CPAs*

By: Marc Hamilton, CPA CGMA

CFO & Treasurer

Three Notch Group, Inc.


LEARN
VOYAGE


2026 & 2027
Cruise Schedule
Earn CPE and Enjoy a Caribbean Get Away!

Date	Ship	Ports
May 10-17, 2026	<i>Carnival Celebration</i>	Miami • Amber Cove • St. Thomas • Celebration Key (20 CPE Hours)
Aug 9-16, 2026	<i>Celebrity Beyond</i>	Miami • CocoCay • Cozumel • Grand Cayman • Nassau
May 15-22, 2027	<i>MSC World America</i>	Miami • Puerto Plata • San Juan • Ocean Cay MSC Marine Reserve
Sept 5-12, 2027	<i>Norwegian Luna</i>	Miami • Puerto Plata • St. Thomas • Tortola • Great Stirrup Cay

Join Us!
Register before Cabins are filled!
www.LearnVoyage.com

Questions?
Email: Marc.Hamilton@learnvoyage.com


Florida Institute of Certified Public Accountants
CPE Certification administered through
the Florida Institute of CPA's (FICPA)