

AI BUYER'S GUIDE

Back-Office AI Buyer's Guide for Henderson CPA Firms

Where AI works in your operations, where it doesn't, and what it costs. Mid-2026.

By Piers Rollinson, DomeWorks

Henderson, Nevada



The 4% Problem

I keep talking to Henderson CPA owners who've been pitched the same line: *Save 18 hours a week with AI*. It's a real number from a major practice-management vendor's own customer survey. It sounds great. It's also five times what the only peer-reviewed study on the same question actually measured. Stanford and MIT looked at 79 small and mid-sized firms last year and found AI users reallocated about 3.5 hours a week, not 18. That gap is the entire problem with how AI is being sold to small CPA firms right now.

Here's what surprised me more. CPA.com's 2025 industry survey found that 96% of small accounting firms have no mature system for measuring AI ROI. Vendors are quoting 30-70% time savings on focused workflows, and almost no buyer has the measurement infrastructure to know which side of that range they actually landed on. That's a lot of capex moving through small firms with no way to tell if the spend is working.

This guide is the buyer's view I wish you had before any of those vendor calls. Every number traces to a citation on the back cover. Vendor claims are labeled vendor claims. Where the evidence is thin, I say so. I haven't taken money from any vendor named in this guide and I don't resell their tools. Where AI does work, I'll show you which workflows. Where it doesn't, I'll show you why and what to do instead.

One scope note. This guide is about operations: client work, document processing, compliance, advisory production. Marketing, sales, and rainmaking are a different conversation. Stay here if your top constraint is the work itself.

TLDR

If you only read one page, here's the shape of the rest.

1. Three buckets. Every back-office workflow falls into one of three categories. *AI absorbs cleanly* means the tool runs end-to-end and you review lightly. *AI assists, human owns* means the tool drafts and you decide. *Stays fully human* covers judgment, attestation, and regulated decisions where AI doesn't help and may hurt. The framework on pages 2-3 is what to memorize. Everything else applies it.

2. The cost math is real but smaller than vendors say. A sensible AI back-office stack for a 15-person Henderson firm runs about \$26,000 a year. Recovered capacity is 30-50 hours per week firm-wide, not 200+. Page 8 shows the numbers.

3. The #1 thing to not get burned by. Two of the most-marketed small-firm bookkeeping platforms collapsed in 13 months. Your contracts need an escrow clause that wasn't required two years ago. Page 9 covers all five common burns.

Vendors are quoting 30-70% time savings on focused workflows, and almost no buyer has the measurement infrastructure to know which side of that range they actually landed on.

CPA.COM 2025 INDUSTRY SURVEY

How to Think About It: The Three Buckets

Most CPA owners trying to figure out AI start with the wrong question. *Which tool should I buy?* The tool catalog is enormous. The vendor pitches are noisy. Every demo looks impressive. Six weeks in you've signed up for two tools, neither one is being used by the team, and you've added a line item to the P&L that nobody can defend.

The question that actually works is upstream of any tool. *Which of my workflows can AI absorb cleanly, which can it assist with, and which should stay fully human regardless of what any vendor says?* Once you've placed each workflow in one of those three buckets, the tool question becomes a 20-minute exercise instead of a six-month wandering. The framework is the whole game.

Bucket 1: AI absorbs cleanly

The tool runs end-to-end on routine cases at roughly 85% or better accuracy. A human reviews exceptions and edge cases. The work shifts from doing to checking, and the time saved is real and measurable.

What lives here today:

- **Invoice extraction on clean books.** Header data, line items, vendor coding, GL routing. Dext, AutoEntry, and Bill.com all do this competently after a 3-6 month learning period. Vendor accuracy claims of 97-99% are believable on routine invoices and not on handwritten fuel receipts.
- **1099 form generation with IRS TIN matching.** Tax1099, Avalara 1099. Binary against IRS records. Either the TIN matches or it doesn't. AI absorbs this completely. The only thing that stays human is the *classification* of who gets a 1099 in the first place.

-
- **Document intake routing for stable tax forms.** W-2s, 1099-NECs, K-1s. Predictable templates, AI files them in the right workpaper folder before your staff opens the email.

Bucket 2: AI assists, human owns

The tool drafts. You review and decide. The accuracy ceiling is 70-85% on independent benchmarks, which is high enough to multiply your throughput and low enough that unreviewed AI output will embarrass you in front of a client.

What lives here:

- **Bank reconciliation on messy books.** Stripe and Shopify net-of-fee payouts, owner-account transfers, duplicate feeds. The DualEntry benchmark in April 2026 found bank reconciliation is one of the two weakest categories for general-purpose LLMs, with the best model still getting roughly one in five tasks wrong.
- **Financial statement preparation for compilations and reviews.** Caseware Validate runs 450 internal-consistency checks per statement and saves real time on subsequent drafts. The CPA still owns the SSARS deliverable. AI does not write the disclosures.
- **Client question triage for substantive replies.** Karbon AI, TaxDome AI. They route well. They draft well enough to start from. They confidently invent details that aren't true if you let them send unedited.

Bucket 3: Stays fully human

AI provides minimal value or the risk profile is wrong. These are judgment calls, attestation, regulated decisions, and anything that attaches CPA liability under Circular 230, AU-C 240, or §6694.

What lives here:

- **Worker classification.** 1099 vs W-2 vs corporate-exempt. This is a legal determination, not an extraction problem. AI can help you draft the contemporaneous documentation. It cannot decide the question.
- **Audit risk assessment, scoping, and the audit opinion.** PCAOB's acting chair flagged AI overreliance as an audit-quality risk in November 2025. AI-driven 100%-population testing creates false confidence when the underlying GL is incomplete.
- **K-1 special allocations, §704(b)/§754 items, multi-state conformity edge cases.** Hallucinated tax-law citations are the most-cited failure mode in independent reporting. The fix isn't a better prompt. The fix is a CPA reading the code.

01

AI absorbs cleanly

85%+

Tool runs end-to-end. Human reviews exceptions.

- Invoice extraction on clean books
- 1099 TIN matching against IRS records
- Document intake routing for stable tax forms

02

AI assists, human owns

70-85%

Tool drafts. You review and decide.

- Bank reconciliation on messy books
- Financial statement preparation

- Client question triage for substantive replies

03

Stays fully human

n/a

Judgment, attestation, regulated decisions.

- Worker classification (1099 vs W-2)
- Audit risk assessment + opinion
- K-1 special allocations, §704(b)/§754

The three-bucket framework. Place each back-office workflow in one of these.

| *The framework is the whole game.*

Workflow-by-workflow read

Twelve workflows ordered roughly easiest-to-hardest. Each entry tags its bucket, names a tool the research actually evidences, gives one accuracy or time-savings number with the source labeled vendor or independent, and flags the failure mode worth knowing. If a workflow doesn't appear here, the research didn't substantiate a useful claim and I left it out.

Bucket 1: AI absorbs cleanly

1. Invoice processing (accounts payable)

OCR plus learned vendor-coding rules. Header and line-item extraction; low-confidence items routed to a human.

- **Tools, \$12-\$240/mo for a 10-20 person firm.** Dext, AutoEntry, Bill.com, Hubdoc, Ramp Bill Pay.
- **Vendor.** Dext: 99% extraction accuracy across 31.4M January 2026 documents.
- **Independent.** Ardent Partners 2025 *AP Metrics That Matter*: best-in-class teams at 49.2% straight-through processing, average 32.6%. Vendor demos imply closer to 100%; field reality is exceptions happen often.
- **Failure mode.** New vendors with no learned rules, multi-page PDFs bundled as one upload, handwritten fuel receipts, vendor name aliases that break learned rules.

2. 1099 form generation with IRS TIN matching

Form preparation, IRS TIN validation, CF/SF state filings.

- **Tools, about \$300-\$5,000/yr depending on volume (300-4,000 forms).** Tax1099 (about \$2.99/form pay-as-you-go), Avalara 1099, TaxBandits (\$0.80-\$2.75).
- **Accuracy.** Binary against IRS records, so percentages don't apply.
- **Failure mode.** TIN mismatches when DBA differs from legal name. W-8 / 1042-S workflows for foreign vendors. State filings outside the CF/SF program. Classification errors upstream that AI won't catch.

3. Document intake routing for stable tax forms

W-2s, 1099-NECs, K-1s with predictable templates.

- **Tools, \$49-\$67/user/mo.** TaxDome, Karbon (Team plan), TaxScout (flat \$49/mo).
- **Vendor.** TaxScout claims 5-layer validation across 180+ form types. No public independent benchmark exists for AI document classification across mixed CPA document sets.
- **Failure mode.** State-specific forms that look federal. Multi-form PDFs uploaded as one packet. Missing-document detection only works if prior year is in the system.

4. Bank reconciliation on clean books

Routine recurring transactions, predictable vendors, no processor payouts.

- **Tools.** QuickBooks Online's 2025 Accounting Agent (bundled, with full agents on Advanced at about \$275/mo per file after the May 2026 price rise). Xero's Just Ask Xero (no extra charge in beta). Booke AI (\$20-60/client/mo).
- **Vendor.** FloQast's Doximity case study: 75-90% of bank transactions auto-matched versus 20-30% with NetSuite native.
- **Failure mode.** The workflow drops to Bucket 2 the moment Stripe, Shopify, or Square net-of-fee payouts appear.

5. Client question triage (routing only)

Categorizing inbound emails and routing to the right team member or workflow.

- **Tools, \$50-\$89/user/mo.** Karbon AI, TaxDome AI.
- **Vendor.** Karbon's 2024 customer survey claims 4 hours per employee per week on triage alone.
- **Independent.** The only peer-reviewed comparison anchor (Stanford/MIT, 3.5 hours per week across all AI tools combined) suggests that 4-hour figure is probably overstated.
- **Failure mode.** Misattribution when contacts share names. The triage layer touches PII, and consumer-tier chats are the wrong place to do it.

Bucket 2: AI assists, human owns

6. Payables matching (PO-to-invoice, exception handling)

Most \$1-5M-revenue clients of Henderson firms don't run formal POs, so this workflow is rare.

- **Tools.** Stampli (no public price).
- **Vendor.** Stampli claims 97-100% PO matching accuracy.
- **Independent.** Ardent Partners 2025: 53% of AP leaders cite high exception rates as a top challenge. The vendor "100% match" framing doesn't survive contact with partial deliveries, freight surcharges, and service POs without goods receipts.
- **Failure mode.** Tolerance configuration is a judgment call, not an extraction problem.

7. Bank reconciliation on messy books

Stripe and Shopify net-of-fee payouts, owner transfers misclassified as revenue, duplicate bank feeds, AI suggestions that persist after a one-time override.

- **Tools.** Booke AI (claims 98% categorization accuracy, vendor).
- **Independent.** The DualEntry / CFO.com April 2026 benchmark identified bank reconciliation as one of the two weakest categories for general-purpose LLMs, with the best model still getting roughly one in five tasks wrong.
- **Failure mode.** Any client with a payment processor or owner-account complexity is exactly where these systems fail.

8. Financial statement preparation (compilations, reviews)

The CPA still owns the SSARS deliverable and the disclosures.

- **Tools.** Caseware Working Papers + AiDA + Validate. Quote-only, several hundred per user per year.
- **Vendor.** Caseware Validate: 450 checks per statement, 40-50% efficiencies on subsequent drafts. One customer cut casting from 25 to 12 minutes.
- **Failure mode.** AI confidently copies forward last year's lease disclosure when the lease was modified. AI does not know what disclosure is missing because it was never trained on what should be there.

9. Client question triage (substantive replies)

- **Tools.** Karbon AI, TaxDome AI.
- **Independent.** Karbon's State of AI survey: 70% of accountants cite data security as a top concern.

- **Failure mode.** AI drafts well enough to start from and confidently invents details if you let staff send unedited. The substantive reply is also where firm voice and warmth live; consumer-tone AI breaks the relationship even when it gets the facts right.

10. Month-end close

- **Tools.** FloQast (\$9,000-\$12,000/year for a small team). Numeric (Starter free, paid quote-only). Sage Intacct's Close Assistant (bundled for Intacct clients).
- **Vendor.** FloQast claims 20% reduction in monthly close time and up to 40% automation of close tasks.
- **Independent.** The DualEntry benchmark named close as one of the two weakest LLM categories. Numeric's founder publicly conceded hallucination risk in flux commentary, mitigated by clickable transaction links so the actual tie-out math is not done by the LLM.
- **Failure mode.** Flux explanations that confidently invent a "seasonality" driver when the real cause is a one-time JE.

11. Tax return preparation (1040, 1065, 1120-S)

Thomson Reuters' "Ready to Review" framing for AI-drafted returns is the right product-strategy posture: AI prepares, the CPA reviews and signs.

- **Tools.** Drake Tax Unlimited (\$1,995-\$2,495/year firm-wide, the small-firm price-leader). Intuit ProConnect (\$1,049 flat for 10+ users plus per-return). Lacerte, UltraTax CS, CCH Axxess Tax.
- **AI overlays.** Blue J (about \$1,500/seat/year), TaxGPT (about \$1,498/year).
- **Independent.** Washington Post's March 2024 evaluation of *consumer* tax chatbots: TurboTax's Intuit Assist wrong on more than 50% of 16 questions, H&R Block's AI Tax Assist wrong on roughly 30%. There is no equivalent independent test of professional tools.

- **Failure mode.** Hallucinated tax-law citations, misread K-1 special allocations, state conformity errors, outdated post-OBDDA SALT guidance.

12. Audit and assurance work

PCAOB Acting Chair Botic warned in November 2025 about overreliance on AI jeopardizing audit quality.

- **Tools.** DataSnipper (Excel add-in, \$300-\$1,000/user/year). MindBridge AI (engagement-based, \$15,000-\$40,000/year for a 10-20 person firm, ICAEW-accredited). Caseware IDEA + AiDA.
- **Independent.** MindBridge has the strongest independent algorithm validation in the category (annual audits by Holistic AI, ICAEW Technology Accreditation). The UK National Audit Office reported £400,000 in efficiency savings using DataSnipper, with routine work completed up to 3x faster.
- **Failure mode.** AI-driven 100%-population testing creates false confidence when the underlying GL is incomplete. AI cannot replace AU-C 240 fraud-risk assessment.

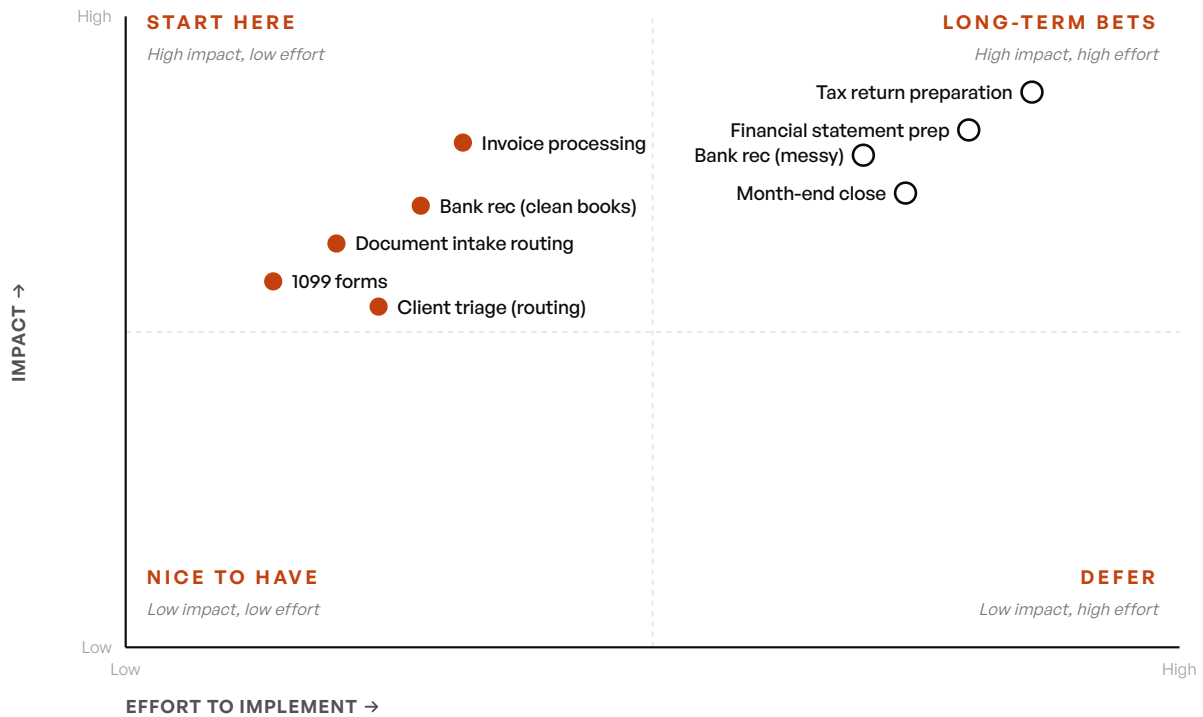
Bucket 3: Stays fully human

13. Worker classification, audit opinion, judgment items

1099 vs W-2 vs corporate-exempt. The audit opinion itself. K-1 special allocations, §704(b)/§754 items, going-concern indicators, fraud-risk assessment, scoping, and any tax position that attaches preparer liability under Circular 230, AU-C 240, or §6694.

AI can help draft contemporaneous documentation. AI cannot decide the question.

Failure mode. The most expensive failure mode in the entire stack is treating any of these as a Bucket 2 task because a vendor demo made it look that way.



● Bucket 1 (AI absorbs cleanly) ○ Bucket 2 (AI assists, human owns)

Workflow impact-vs-effort. Bucket 3 items (worker classification, audit opinion, K-1 special allocations) are deliberately off-chart: AI cannot decide them.

Hallucinated tax-law citations are the most-cited failure mode in independent reporting. The fix isn't a better prompt. The fix is a CPA reading the code.

Which workflows to start with

A first-cut decision tree to place any back-office workflow in one of the three buckets. Walk a workflow through these three questions in order.

Q1. Is the workflow rules-based with a binary correct/incorrect outcome that can be checked against an external source of truth?

- *Yes* → likely **Bucket 1: AI absorbs cleanly**. (Examples: 1099 TIN matching against IRS records, invoice extraction against vendor catalogue, document classification against form templates.) Start here.
- *No* → continue to Q2.

Q2. Does the workflow require professional judgment that attaches preparer or CPA liability under Circular 230, AU-C 240, or §6694?

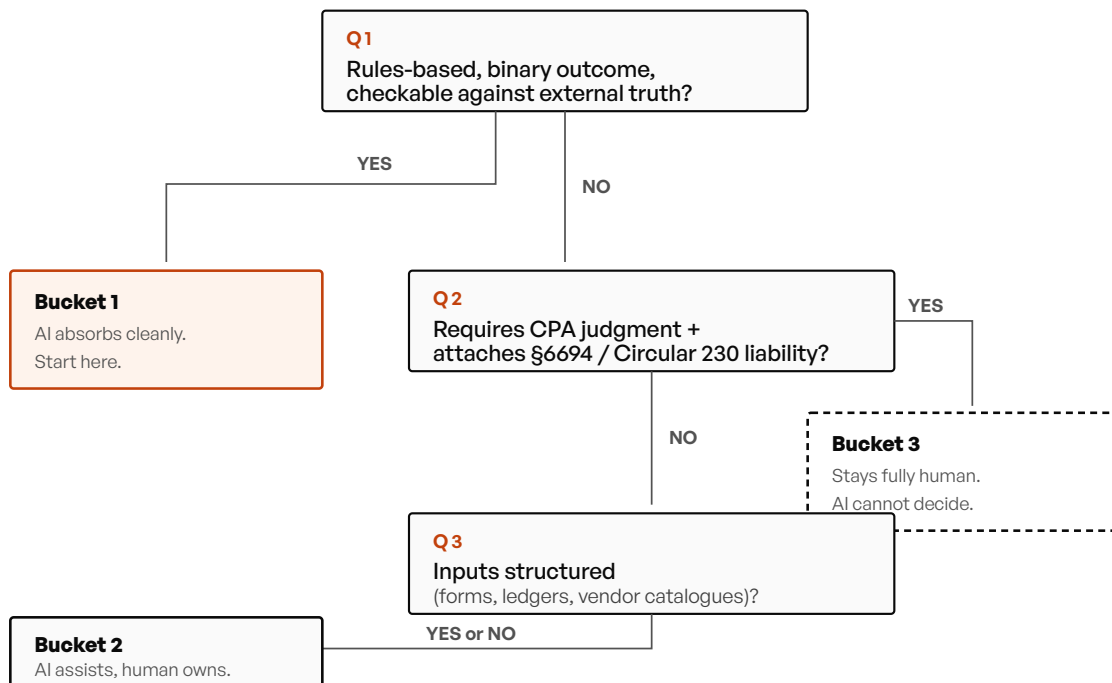
- *Yes* → **Bucket 3: Stays fully human**. (Examples: worker classification, audit opinion, K-1 special allocations, going-concern.) Don't start with AI here. AI can support documentation; it cannot decide the question.
- *No* → continue to Q3.

Q3. Are the inputs structured (forms, ledgers, vendor catalogues) or unstructured (emails, conversations, narrative client notes)?

- *Structured* → **Bucket 2: AI assists, human owns**. High productivity gain available with disciplined review. (Examples: bank rec on messy books, financial statement preparation, payables matching.)
- *Unstructured* → **Bucket 2: AI assists, human owns**, but treat as exploratory. Expect lower accuracy. (Examples: substantive client question replies, draft responses to client follow-up.)

The caveat

This tree is a first cut. Your firm's actual workflow mix, client base, and software stack determine the right starting point. Two firms with identical workflows often end up in different buckets because one has clean books and the other doesn't. The audit produces a sequenced 30-60-90 day rollout for your specific firm based on what's actually in your stack today.



Walk any back-office workflow through these three questions in order. The visual mirrors the prose decision tree on the left.

The Math

What an AI back-office stack actually costs and what it actually saves for a 15-person Henderson CPA firm doing about \$3M in fees.

Sensible stack

Layer	Tool	Monthly cost
Practice management + AI	Karbon Team, 15 seats	\$885
Document capture	Dext accountant plan	\$300
Tax software (amortized)	Drake Unlimited or ProConnect	\$200
1099 filing (amortized over year)	Avalara 1099 / Tax1099	\$150
Ad-hoc LLM	ChatGPT Team or Claude Team, 15 seats	\$375
Audit add-on (if firm does audit)	DataSnipper, ~5 seats	\$250
Total		\$2,160/mo, \$26,000/yr

This excludes firm-side QBO ProAdvisor (free) and client-paid QBO/Xero subscriptions. A firm doing serious audit work adds MindBridge at \$1,500-\$3,500/mo, pushing the stack to roughly \$45,000-\$60,000/yr. A firm with no audit and minimal close work could run as low as \$18,000/yr.

Recovered hours

The peer-reviewed Stanford/MIT figure is 3.5 hours per AI-using accountant per 40-hour week. For 15 employees that's about 52 hours per week of capacity. Discount aggressively for the AICPA 2025 MAP "wait-and-see" cohort and you get 30-50 hours per week of real recovered time. Not the 277 hours implied by Karbon's vendor figure.

Break-even

- Annual stack cost: \$26,000
- Annual recovered capacity at 40 hrs/week × 50 weeks = 2,000 hours
- At \$75/hr internal cost: \$150,000 of recovered capacity
- At \$200/hr realized billable rate (if redirected to chargeable work):
\$400,000

The stack pays for itself if the firm redirects roughly 130 hours per year to billable work. That's about 2.5 hours per week firm-wide, well under the recovered capacity. The arithmetic is easy.


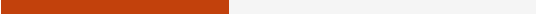
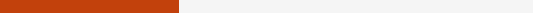
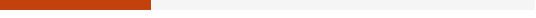
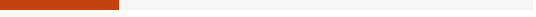
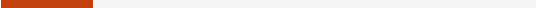
The catch

The risk isn't break-even on paper. The risk is whether recovered hours actually become billable. Firms that absorb hours into earlier evenings rather than client work see no P&L benefit at all. The pre/post measurement system is what makes the difference. Without it, you'll feel like AI helped, your team will agree, and your billable hours per FTE won't move.

The audit produces your firm's stack and your firm's break-even, not a hypothetical. The Tier-1.5 measurement engagement sets up the tracking that turns recovered hours into billable hours instead of into earlier evenings.

15-PERSON HENDERSON FIRM, ~\$3M FEES

\$2,160/mo ~\$25,920/yr

Karbon Team (15 seats) Practice management + AI		\$885/mo
ChatGPT/Claude Team (15 seats) Ad-hoc LLM		\$375/mo
Dext accountant Document capture		\$300/mo
DataSnipper (audit, 5 seats) Audit add-on		\$250/mo
Drake Unlimited or ProConnect Tax software (amortized)		\$200/mo
Avalara 1099 / Tax1099 1099 filing (amortized)		\$150/mo

Excludes firm-side QBO ProAdvisor (free) and client-paid QBO/Xero. Adding MindBridge for serious audit work pushes the stack to ~\$45-60K/yr. Without audit it can run as low as ~\$18K/yr.

The sensible AI back-office stack for a 15-person firm.

The risk isn't break-even on paper. The risk is whether recovered hours actually become billable.

ONE A MONTH, NO DRIP

Want the next monthly note from the field?

I send one email a month with what's actually working in small-firm AI rollouts — and what isn't. Same standard as this guide. Drop your email if you'd like the next one.

To subscribe from the PDF: email piers@domeworks.tech with "Monthly notes" in the subject.

What Not to Get Burned By

Five common burns. The first two are vendor-risk exposure that didn't exist two years ago. The next three are operational. Every one of them has cost a small firm real money in the last 18 months.

1. Vendor concentration risk

Two of the most-marketed small-firm bookkeeping platforms collapsed in 13 months. Bench Accounting shut down in December 2024 and locked roughly 12,000 customers out of their books days before tax season. Botkeeper closed in February 2026 after 11 years and approximately \$90M raised; the Infinite platform was acquired by Xendoo. Both held client books on proprietary infrastructure with no easy export path.

Mitigation: any tech-enabled bookkeeping vendor you sign with should have a contractual escrow clause requiring working papers exportable in QBO or Xero-compatible formats. This is a new line item in vendor evaluation. Two years ago it wasn't required. Today it is.

2. The 5x time-savings overstatement

Vendor surveys quote 18+ hours per employee per week saved. Karbon, FloQast, DataSnipper, Vic.ai all sit in that range. The only peer-reviewed independent study (Stanford and MIT, 79 small/midsize firms) measured 3.5 hours per week. That's roughly a 5x gap.

Mitigation: when evaluating any AI vendor, divide their claimed time savings by 5 as your starting estimate. Then ask the vendor to point you to an independent study that reproduces their number. They almost never can. That answer is itself the data.

3. Hallucinated tax law

The DualEntry / CFO.com April 2026 benchmark tested 19 AI models against 101 real accounting tasks. The best model (Claude Opus 4.7) hit 79.2% accuracy. That means roughly one in five answers was wrong, with bank reconciliation and month-end close as the two weakest categories. Wash-sale rules, multi-state filing requirements, and K-1 special allocations are the most-cited failure points.

Mitigation: every AI-drafted memo, response, or disclosure must trace to a clickable source. If the AI cites a code section, click the link and read the section yourself before relying. Circular 230 §10.22 and §10.34 and IRS §6694 preparer penalties attach regardless of whether AI was involved.

4. Tier-locked pricing

"AI is included" framing is misleading at base pricing. Intuit's most powerful agents require QuickBooks Advanced at roughly \$275 per file per month after the May 2026 price rise. Karbon's strongest features are on Business at \$89/user/mo, not Team at \$59/user/mo. Sage Copilot beyond search is a paid add-on. The actual AI you want is usually one tier above the price the marketing page quotes.

Mitigation: read the pricing page, not the product page. For each AI capability the vendor demos, find the line in the pricing table that includes it. If the demo capability isn't on the tier you're considering, that's the actual price.

5. Confidentiality leak on consumer chats

Karbon's State of AI survey found 70% of accountants cite data security as a top concern. Tankersley and Johnston's accounting-specific safety review ranks model safety as Claude > Copilot > Perplexity > ChatGPT > Gemini.

Consumer-tier chat tools do not provide enterprise privacy controls. PII pasted into a free ChatGPT or Gemini chat is a documented risk pattern.

Mitigation: a one-paragraph AI use policy is the cheapest control your firm can implement. What tools the firm allows. What client data can never be pasted into consumer chats. Who reviews AI-drafted output before it goes to a client. Page 10 has a starting structure.

01

Vendor concentration risk

Bench (Dec 2024) and Botkeeper (Feb 2026) collapsed within 13 months.

MITIGATION Require an escrow clause for QBO/Xero-compatible export.

02

5x time-savings overstatement

Vendor 18+ hrs/wk vs Stanford/MIT peer-reviewed 3.5 hrs/wk. About a 5x gap.

MITIGATION Divide vendor claims by 5 as your starting estimate.

03

Hallucinated tax law

Best LLM hit 79.2% on 101 accounting tasks (DualEntry/CFO.com Apr 2026).
Roughly 1 in 5 wrong.

MITIGATION Every AI memo or citation must trace to a clickable source. Read it.

04

Tier-locked pricing

"AI included" usually means base tier. Real AI is one tier up: QBO Advanced \$275/mo, Karbon Business \$89/user/mo.

MITIGATION Read the pricing page, not the product page.

05

Confidentiality leak on consumer chats

70% of accountants cite data security as their top AI concern (Karbon State of AI).

MITIGATION A one-paragraph AI use policy is the cheapest control your firm can implement.

Five burns. Every one has cost a small firm real money in the last 18 months.

If you're a Henderson firm

I'm based in Henderson. The local context shapes how I think about the rest of this guide.

Where Henderson CPAs cluster

The Nevada Society of CPAs (NSCPA) Southern Nevada Chapter is the easiest place to meet other firm owners working through the same questions. AI is showing up on the agenda more than it did even six months ago. If you're not connected, the chapter membership is worth the \$200ish a year for the conversations alone.

The Henderson Chamber of Commerce has been running AI-focused events through 2026 for SMB owners across sectors. Most attendees aren't accountants, which is part of the value. You hear how dental practices and law firms are scoping the same problem you are, and the patterns generalize more than you'd expect.

The competitive picture locally

The AICPA's 2025 National MAP Survey found that small firms under \$1.5M in fees describe themselves predominantly as "observing developments" on AI. That description matches what I see in Henderson too. The firms that move from observing to doing in the next 12 months will have a real two-year window before this becomes table stakes. The firms that wait until everyone else has it will be paying premium pricing for the same outcomes.

The local advantage

Local time zone is a real advantage I underweight when talking to founders out of the area. Same-day response, in-person walkthroughs at your office, the ability to be in your team's training session rather than over Zoom. The Tier-1.5 measurement engagement in particular is a lot easier to set up when the consultant can sit at the desk of whoever owns the time-tracker for an hour.

If you'd rather have this conversation in person at the Chamber, NSCPA, or your office instead of on a video call, that's an option. The 45-minute walkthrough on the last page can be in-person for any Henderson or Las Vegas firm. Same for the audit itself.

Three Things to Do This Quarter

Three concrete moves any Henderson CPA firm can make in the next 90 days without spending a dollar on new tools. Do them in order. The third one prevents the worst failure modes from page 9.

1. Audit your top 3 highest-volume workflows against the three buckets

Pick the three things your team actually spends the most hours on. Not the things you wish were the priorities. The actual hour leaders. AP processing, bank rec, and client question triage are common picks for a 15-person firm; yours may differ.

For each, place it in one of the three buckets using the framework on pages 2-3 and the decision tree on page 7.5. Don't buy any tool yet. The point of this step is to learn which problems are AI-shaped and which aren't. Most firms discover their #1 hour leak is actually a Bucket 3 task that no tool will help with, and their #2 leak is a Bucket 1 task that's been waiting for someone to notice.

2. Start measuring one process before/after

Pick one Bucket 1 or Bucket 2 workflow from step 1. Time it for two weeks before any AI changes. Record the median hours per week your team spends on it. That baseline is the only honest way to know whether AI helped when you do introduce a tool.

Per the page 1 number: 96% of small firms have no system for measuring AI ROI. Two weeks of tracking puts you in the 4% before you've even bought anything. The measurement isn't about justifying tool spend. It's about

preventing the failure mode where everyone agrees AI helped and the billable hours per FTE haven't moved.

3. Write a one-paragraph AI use policy

Not a 12-page policy document. One paragraph that answers three questions. What AI tools is the firm using or allowing? What client data can and cannot be pasted into those tools? Who reviews AI-drafted output before it goes to a client?

A starting paragraph: *"This firm uses [tool] for [purpose]. Client PII (names, EINs, SSNs, financial details) may not be pasted into consumer-tier AI tools. AI-drafted client-facing communication is reviewed and approved by a CPA before sending. AI-drafted memos and disclosures must include a citation that traces to a primary source."* Edit to taste. The point is having anything written down before staff need to make a call about a client's tax data on a Friday afternoon.

If you'd rather not do these alone, the 45-minute walkthrough on the last page applies all three steps to your firm. No commitment, just a working session.

About & Ways to work together

About Piers

I'm Piers Rollinson. I run DomeWorks, a Henderson-based independent advisory practice helping small firms across professional services figure out where AI actually fits and where it doesn't.

Before DomeWorks, I spent 15 years building products and leading engineering teams. Most recently I was Director of Engineering at Mudflap (Series B fintech serving 500K+ truckers) where I led the AI Platform and Voice AI teams and redesigned the engineering organization around AI and agentic coding. Before that I was an Engineering Manager at DoorDash and Square, and Head of Engineering at Zesty (acquired by Square in 2018). I've been building software since 2009 and managing engineers since 2017.

I'm independent. I don't resell AI tools. I'm not partnered with the vendors named in any of my guides. I get paid by the firms I work with, not by tool referrals. That's the only way the honest-broker positioning works.

Ways to work together

In order from lowest commitment to highest.

01

Free

Free 45-minute walkthrough

I'll walk through your top three highest-volume workflows against the framework you just read, show you which one I'd start with first, and answer questions about specific tools. No slides, no pitch. Henderson or Las Vegas in person, or video.

[Book the walkthrough →](#)

02

\$999

Tier-1 audit

Four-phase deliverable for firm owners who want a written diagnosis. 45-minute discovery call, AI analysis of your specific workflows, custom report identifying 5-7 opportunities ranked by impact and effort, 30-minute walkthrough call. 5+ hours-per-week recovered-capacity guarantee or full refund.

[Book a discovery call →](#)

03

2 weeks

Tier-1.5 measurement engagement

If your bigger problem is that you're not measuring AI ROI at all, you're not alone. 96% of small firms aren't. Two-week engagement to set up the measurement system that turns recovered hours into billable hours instead of earlier evenings. Pre-baseline, per-workflow attribution, conversion-to-billable tracking.

[Ask on the walkthrough →](#)**MID-MARKET****If your firm is bigger than 25 people**

The audit is calibrated to firms with 5-25 back-office staff. If you're larger and looking at AI as an org-wide capability, the fractional CAIO engagement is a different conversation. Strategy work, governance, ROI measurement at the program level.

[Read about the fractional CAIO offer →](#)**Stay in touch**

If you want monthly notes from me on what's actually working in the field, drop your email below. No drip campaign. One email a month, written to the same standard as this guide.

To subscribe from the PDF: email piers@domeworks.tech with "Monthly notes" in the subject.

Sources

Every numerical claim in this guide traces to one of the URLs below. Vendor claims are labeled vendor; independently reported claims are labeled independent.

Workflows and benchmarks

- Stanford/MIT (Choi/Xie) study via Journal of Accountancy, August 2025 (independent, peer-reviewed). 79 small/midsize firms; 3.5 hrs/week reallocated; 21% higher billable hours; close 7.5 days sooner. <https://www.journalofaccountancy.com/news/2025/aug/calculating-ais-impact-on-cpas-new-study-quantifies-time-savings/>
- DualEntry / CFO.com benchmark, April 2026 (independent). 19 AI models, 101 accounting tasks, best model 79.2% accuracy, bank rec and close weakest. <https://www.cfo.com/news/the-best-ai-model-still-fails-1-in-5-accounting-tasks-Claude-Opus-OpenAI-GPT/818100/>
- Washington Post (Geoffrey Fowler), March 2024 (independent). TurboTax Intuit Assist wrong on >50% of 16 questions; H&R Block about 30%. <https://www.washingtonpost.com/technology/2024/03/04/ai-taxes-turbotax-hrblock-chatbot/>
- Ardent Partners "AP Metrics That Matter in 2025" (independent). Best-in-class 49.2% straight-through processing; average 32.6%. <https://ardentpartners.com/ap-metrics-that-matter-in-2025/>
- IOFM AP benchmarks via secondary sources (industry). Manual error rate about 2% vs automated 0.8%; manual cost about \$9.40/invoice vs \$1.45 fully automated. <https://www.ascendsoftware.com/blog/what-good-looks-like-ap-benchmarks-every-modern-team-should-know-in-2025>

- Filed TaxCalcBench (vendor with disclosed methodology). Base ChatGPT/Claude/Gemini about 30-35% accuracy on full tax returns. <https://www.filed.com/measuring-ai-tax-accuracy-filed-vs-chatgpt-claude-gemini>
- DataSnipper case studies + UK National Audit Office citation (independent for the NAO portion). NAO reported £400,000 in efficiency savings, 3x faster routine work. <https://www.datasnipper.com/external-audit>
- MindBridge independent algorithm audit (independent). Annual audits by Holistic AI, ICAEW Technology Accreditation. <https://www.mindbridge.ai/blog/building-trust-in-artificial-intelligence-for-audit/>

Vendor pricing

- QuickBooks 2026 pricing (vendor). <https://stephsbooks.com/news/quickbooks-online-price-increase-2026>
- Karbon (vendor). <https://karbonhq.com/pricing/>
- Dext (vendor via G2). <https://www.g2.com/products/dext/pricing>
- AutoEntry (vendor). <https://www.autoentry.com/pricing>
- Bill.com (vendor). <https://www.bill.com/product/pricing>
- Drake / ProConnect / UltraTax (industry survey). <https://www.thetaxadviser.com/issues/2025/aug/2025-tax-software-survey/>
- Blue J (vendor coverage). <https://venturebeat.com/technology/how-ai-tax-startup-blue-j-torched-its-entire-business-model-for-chatgpt-and>
- Tax1099 / Avalara 1099 (vendor). <https://www.tax1099.com/tax-1099-efile-pricing>
- Sage Intacct via Cargas (third-party). <https://cargas.com/software/sage-intacct/pricing/>
- FloQast via Numeric (third-party estimate). <https://www.numeric.io/blog/floqast-pricing>
- Xero JAX (vendor). <https://www.xero.com/us/ai-in-accounting/jax/>

Industry signals and surveys

- AICPA / CPA.com 2025 AI in Accounting Report (industry). Source for the 4% measurement-maturity statistic. <https://www.cpa.com/news/cpacom-issues-2025-ai-accounting-report>
- AICPA 2025 National MAP Survey small-firm insights (industry). <https://www.aicpa-cima.com/resources/article/key-small-firm-insights-in-the-2025-national-map-survey>
- Wolters Kluwer 2025 Future Ready Accountant (vendor-sponsored). AI adoption 9% → 41% YoY across 2,768 pros in 14 countries. <https://www.wolterskluwer.com/en/news/wolters-kluwer-releases-its-2025-future-ready-accountant-report>
- Intuit QuickBooks 2025 Accountant Tech Survey (vendor-sponsored). 46% daily AI use claim. <https://www.firmofthefuture.com/news/accountant-tech-survey-2025/>
- Karbon State of AI 2025 (vendor). 70% data security concern; 18.5 hrs/week claim; 37% invest in AI training. <https://karbonhq.com/resources/state-of-ai-accounting-report-2025/>
- Journal of Accountancy "Real-Life Ways Accountants Are Using AI", June 2025 (industry). <https://www.journalofaccountancy.com/issues/2025/jun/real-life-ways-accountants-are-using-ai/>

Failures, regulation, vendor-risk

- Bench shutdown (TechCrunch, Dec 2024). <https://techcrunch.com/2024/12/30/bench-to-be-acquired-after-abruptly-shutting-down/>
- Bench customer billing dispute (TechCrunch, March 2025). <https://techcrunch.com/2025/03/14/bench-is-charging-people-for-services-they-already-paid-for-some-customers-say/>
- Botkeeper closure (Accounting Today, Feb 2026). <https://www.accountingtoday.com/news/botkeeper-shuts-down>

-
- Botkeeper post-mortem (CFO Brew, Feb 2026). <https://www.cfobrew.com/stories/2026/02/17/botkeeper-what-went-wrong>
 - Xendoo acquisition of Botkeeper Infinite (Feb 2026). <https://www.cpapracticeadvisor.com/2026/02/27/xendoo-buys-botkeeper-infinite-ai-platform-will-remain-active-and-supported/179027/>
 - PCAOB Acting Chair Botic, November 2025 Baruch College speech. <https://www.kslaw.com/blog-posts/pcaob-acting-chair-urges-shared-definition-of-audit-quality-amid-ai-and-pe-changes>
 - Numeric founder on hallucination risk (TechCrunch, Oct 2024). <https://techcrunch.com/2024/10/10/numeric-grabs-28m-series-a-for-automating-accounting-with-ai/>
 - CPA Practice Advisor "Accounting Technology Lab" Tankersley/Johnston AI safety ranking, May 2025. <https://www.cpapracticeadvisor.com/podcasts/discussing-the-2025-aicpa-ai-symposium-the-accounting-technology-lab-podcast-may-2025/>