

● SAMPLE DELIVERABLE

AI EXECUTIVE ASSESSMENT · FINAL REPORT & READOUT

The AI Executive Assessment

Two weeks of interviews, a 96-person staff survey, and a six-system review — turned into evidence, a scored roadmap, and a decision your partner group can defend.

PREPARED FOR

Meridian & Frost, PLLC
Dana Whitfield · Managing Partner

ASSESSMENT

March 2–13, 2026
Readout March 17, 2026

PREPARED BY

Brad Taylor
Fractional Chief AI Officer

Executive summary

Over two weeks I interviewed ten people, surveyed 96 staff, and reviewed the six systems where the firm's work actually lives. The question was never "should M&F use AI" — 41 of your people already do, without rules. The question is where AI earns its keep, what it must never touch, and in what order to move. This report answers all three with numbers.

41/96

STAFF ON PUBLIC AI TOOLS

▼ ungoverned, last 90 days

9

PASTED CLIENT DATA INTO FREE TOOLS

▼ self-reported; likely undercount

~6,850 h

ANNUAL HOURS IN TOP-5 WORKFLOWS

≈ \$582K/yr at \$85/h blended

40%

TAX TEMPLATES WITH 3+ VERSIONS

▼ blocks a knowledge assistant

The recommendation, up front

Adopt a **hybrid, commercial-first architecture** — enterprise Claude plus Microsoft 365 with a governed knowledge layer — and a phased 90-day implementation starting with governance and one visible quick win. Four paths were on the table: (1) advisory only, (2) implementation only, (3) **advisory + phased implementation — recommended**, (4) wait. Waiting is the one option I'd argue against on the evidence: shadow use is already growing without you.

In this report

- 3 **Method & timeline** — who I talked to, what I looked for

- 4 **Findings: AI usage & risk** — what's already happening without governance

- 5 **Findings: workflow & data map** — six systems, three data tiers

- 6 **Findings: where the hours go** — 18 candidates captured, measured

- 7 **The scored top five** — and why engagement letters go first

- 8 **Architecture recommendation** — commercial, hybrid, or private

- 9–10 **Governance starter kit & the 90-day roadmap**

- 11–12 **Investment & return · recommendation & next steps**

Method & timeline

Ten days on site and remote, three instruments: structured interviews, an anonymous all-staff survey, and hands-on review of the six systems where client work lives. Everything in this report traces back to one of the three.

How the two weeks ran

- Mar 2 — Kickoff with Dana & Marcus**
Scope agreed: whole firm, all four service lines, no sacred cows.
- Mar 3–6 — Stakeholder interviews**
Ten sessions, 60–90 minutes each, notes shared back for correction.
- Mar 5–10 — Staff survey + systems review**
96 of 118 staff responded (81%). Admin access to all six systems with Tom Osei.
- Mar 9–13 — Workflow capture & scoring**
18 candidate workflows captured, hours validated against Karbon time data.
- Mar 13 — Fieldwork ends**
Draft findings circulated to Marcus for factual review.
- Mar 17 — Readout session (this report)**
90 minutes with the partner group; decision requested by Mar 31.

Who was interviewed

PERSON	ROLE
Dana Whitfield	Managing Partner
Marcus Bell	COO
Priya Raman	Tax Partner
Jordan Pike	Audit Partner
Elena Ruiz	CAAS Director
Tom Osei	IT Director
Four staff (anonymized)	Tax senior · CAAS associate · audit senior · admin coordinator


What I looked for

Where the hours actually go — not where people think AI belongs. Opinions about AI are cheap; timesheets, inboxes, and the K: drive don't editorialize. Every candidate workflow on page 6 was sized from Karbon time entries and interview-validated estimates, not enthusiasm.

Findings — what's already happening

The firm has no AI policy, no approved-tool list, and no data-loss-prevention rule. It also has 41 people using public AI tools for client work. That combination is the single most urgent finding in this report: the AI decision has already been made by your staff — the only open question is whether the firm governs it.

Staff using public AI tools in the prior 90 days, by department



Anonymous survey, 96 respondents. 41 total — 43% of respondents. Denominators are respondents per department.

Three facts that set the risk picture

9 staff admitted pasting client information into free AI tools. Anonymous self-reporting on a sensitive question — treat 9 as the floor, not the count.

3 personal ChatGPT accounts held client data in chat history at review time — outside firm control, outside any retention policy, and within §7216's definition of tax-return information.

Zero controls existed: no policy, no DLP rule, no contractual no-training commitments. Nobody broke a rule, because there was no rule to break.

What staff said (anonymized, quoted with permission)

"I know I probably shouldn't paste it, but the deadline wins."

— Tax senior

"Everyone on my team uses it. We just don't talk about it in front of the partners."

— CAAS associate

"Give me a version I'm allowed to use and I'll stop using mine tomorrow."

— Admin coordinator

Read those three together and the fix is obvious: staff aren't reckless, they're unequipped. A sanctioned tool plus a clear policy converts most of this risk into adoption.

Findings — workflow & data-source map

Any AI system is only as safe as the data you connect to it. Six systems hold the firm's work; each was reviewed for what lives there, its sensitivity tier (defined on page 9), and how ready it is to sit behind an AI tool.

SYSTEM	WHAT LIVES THERE	TIER	AI-READINESS
CCH Access	Tax returns, workpapers, client tax-return information.	RED	Well-permissioned but §7216-bound. Integrate last, via governed connectors only, after consent language is refreshed.
QuickBooks Online	Client books across ~140 CAAS engagements.	RED	Clean per-client separation; usable for close automation once tiering is signed. API access is solid.
Microsoft 365 / SharePoint	Email, Teams, templates, procedures, deliverable drafts.	YELLOW	Highest-value corpus, weakest hygiene — see the sprawl note below. Permission model is sound; content needs cleanup.
Karbon	Practice management: jobs, time, client correspondence.	YELLOW	Ready. The best structured source in the firm — it sized every workflow on page 6 and can feed intake automation directly.
SafeSend	Return assembly and delivery to clients.	RED	Vendor's built-in AI features only, and only after their data-handling addendum is reviewed. No third-party bolt-ons.
"The K: drive"	Legacy shared drive: 19 years of everything, unpermissioned.	RED	Not AI-ready and should never be connected as-is. Migrate what's current to SharePoint; archive the rest read-only.


Version sprawl is the quiet blocker

About **40% of tax templates on SharePoint have three or more live versions** — same filename pattern, different content, no marked current copy. Harmless for a human who knows which one Priya's team trusts; poisonous for a knowledge assistant, which will answer confidently from the wrong version. Corpus cleanup is on the roadmap (days 46–90) precisely because it gates workflow #3.

Findings — where the hours go

Eighteen candidate workflows were captured during interviews and sized against Karbon time data. The top five alone consume roughly **6,850 hours a year** — the equivalent of three and a half full-time people spent on document chasing, searching, and drafting from scratch.

Top five candidate workflows, by annual hours



Hours are annualized from Karbon time entries where coded, and interview-validated estimates where not. Advisory letters also carry revenue upside: letters that don't get written are planning engagements that don't get sold.

How the hours were counted

Each workflow got a per-instance time estimate from the people who do it, multiplied by annual volume from Karbon, then sanity-checked against coded time where it exists. Where the two disagreed, I used the lower number. Nothing here assumes AI captures 100% of a workflow — capture assumptions appear on page 11, not here.

The other thirteen

All 18 candidates were scored; the full list ships as an appendix workbook. Five were killed outright at scoring (page 7 method, page 12 list) and eight scored below 60 — real but not worth doing before the top five. The backlog becomes a living document if the engagement continues: re-scored monthly, additions welcome.

The scored top five

Every candidate was scored on four factors, composited to 0–100. The data tier drags on *speed*, not value — Red-tier work is worth doing, it just can't go first because governance and §7216 consent have to land before AI touches client tax data.

<p>Impact</p> <p>Annual hours or dollars at stake, from page 6.</p>	<p>Feasibility</p> <p>1–5: today's tools, your data, your team.</p>	<p>Data tier</p> <p>Green / Yellow / Red — governance drag on start date.</p>	<p>Time-to-value</p> <p>Weeks until someone's week visibly improves.</p>
--	--	--	---

#	WORKFLOW	ANNUAL HOURS	FEAS.	TIER	SCORE	PROPOSED START
1	Engagement letters & proposals from precedent	1,400 h	5	YELLOW	<div style="width: 87%;"><div style="background-color: #e67e22;"></div></div> 87	● April
2	Client document intake & organizer chase	2,100 h	4	RED	<div style="width: 84%;"><div style="background-color: #e67e22;"></div></div> 84	● May, post-consent
3	Internal knowledge assistant	1,800 h	4	YELLOW	<div style="width: 81%;"><div style="background-color: #e67e22;"></div></div> 81	● Post-cleanup
4	CAAS close checklist + variance narratives	950 h	4	RED	<div style="width: 74%;"><div style="background-color: #e67e22;"></div></div> 74	● Q3
5	Post-filing client advisory letters	600 h + revenue	3	RED	<div style="width: 71%;"><div style="background-color: #e67e22;"></div></div> 71	● Q4

Why engagement letters go first, not the biggest number

The doc chase has more hours (2,100 vs 1,400), but it's Red-tier: it touches client tax data and can't start until §7216 consent language is refreshed. Engagement letters are Yellow-tier, feasibility 5 of 5, and visible to every partner who reviews one. The first project's real job is to recruit believers — a quick win in week six converts skeptics better than any slide. Then the doc chase follows with the governance already in place.

Architecture — the options, priced

Three architectures were evaluated against \$7216 compliance, three-year total cost, and what a two-person IT team can realistically operate. Costs include licenses, implementation, and the governed knowledge layer where applicable.

OPTION	3-YR COST	\$7216 / COMPLIANCE FIT	BURDEN ON A 2-PERSON IT TEAM
All-commercial SaaS Enterprise AI subscriptions only, no knowledge layer	\$95–140K	Adequate for Tier 1–2 with no-training contracts. Tier 3 work stays manual.	Low. But it caps value: no firm-knowledge grounding, so drafting quality plateaus.
Hybrid, commercial-first RECOMMENDED Enterprise Claude + M365 + governed, permission-aware knowledge layer	\$120–180K	Strong. Contractual no-training commitments; Tier 3 touched only through the governed layer, after consent refresh.	Moderate and honest for Tom's team: vendor-managed models, one integration layer to own.
Private AI Self-hosted models on firm-controlled infrastructure	\$600–900K	No <i>additional</i> \$7216 benefit at M&F's profile — the obligations attach to disclosure and consent, not to where the model runs.	Heavy: model ops, GPU capacity, security patching. Not a two-person job.

Why private AI is deferred, plainly

It costs 4–6× more and, for this firm today, buys nothing the hybrid doesn't already deliver. Enterprise contracts with no-training commitments plus the data-tier policy meet the compliance bar. Private AI is the right answer for some firms — it is the expensive wrong answer for a \$24M firm with two people in IT.

Three triggers that reopen this decision

1. The wealth practice approaches RIA registration at ~\$1B in assets under advisement.
 2. A client contract imposes data-residency terms the commercial stack can't meet.
 3. A material change in the enterprise tools' data commitments.
- Until one fires, this decision stays closed — revisiting it quarterly by mood is how firms burn a year.

Governance starter kit

Three tiers, one page of policy, a short approved list. Governance that fits on a page gets followed; a 40-page policy gets ignored and then audited.

TIER 1 · GREEN

Public or publishable: marketing copy, published guidance, blank templates. **Any approved tool.**

TIER 2 · YELLOW

Internal operations: SOPs, training material, anonymized examples. **Approved enterprise tools only.**

TIER 3 · RED

Client tax-return information (§7216), PII, client financials, audit workpapers. **Governed systems on the approved list only — never public or personal tools.**

 **Approved from day one**

Claude Enterprise — primary tool, Tier 1–2 (Tier 3 later, only inside the governed knowledge layer) · **Microsoft 365 Copilot** — Tier 1–2, CAAS pilot of 14 seats · **SafeSend built-in AI features** — after vendor data-handling review. All three carry contractual no-training commitments.

 **Prohibited from day one**

Personal and free AI accounts for any firm work · **meeting bots on client calls** (consent) · **any tool without a contractual no-training commitment.** Enforced by an M365 DLP rule, handled as coaching on first contact — the 41 people on page 4 need a door in, not a lecture.

Four decisions the partner group must make

DECISION	OWNER	BY
1 · Adopt the data tiers and the one-page AI use policy as written, or with edits.	Dana Whitfield	Mar 31 partner mtg
2 · Choose the engagement path — advisory, implementation, both, or wait (page 12).	Dana Whitfield	Mar 31 partner mtg
3 · Sign the Claude Enterprise contract and confirm the day-one approved list.	Marcus Bell	Apr 3
4 · Turn on the M365 DLP rule with the coaching-first response standard.	Tom Osei	Apr 10

The 90-day roadmap

Ninety days, three phases, every line dated and owned. The clock started Mar 2 — this assessment was the first two weeks, so the plan below begins already in motion.

Days 0–14 · Mar 2–13 — This assessment **DONE**

Interviews, survey, systems review, workflow scoring, this report. Readout Mar 17; partner decisions due Mar 31.

Days 15–45 · Mar 17–Apr 16 — Governance signed, rollout begun

Mar 31: partner group signs the tiers, policy, and engagement path (page 9). **Apr 1:** advisory engagement starts, if chosen. **Apr 3–10:** Claude Enterprise contract signed, DLP rule live, department training runs in four waves led by champions Marcus nominates. **Apr 13:** engagement-letter quick win kicks off with Priya Raman's team — precedent library assembled, build scoped, target ship mid-May.

Days 46–90 · Apr 17–May 31 — Build the second layer

Corpus cleanup starts on the SharePoint tax templates (~290 duplicates to retire; a named owner per template family) — this gates the knowledge assistant. **Intake-automation SOW** is scoped and signed once §7216 consent language is refreshed with counsel. **Champions program and measurement go live:** six named champions, weekly-active-user reporting from the admin consoles, first monthly executive report delivered June 30.

Day 90 checkpoint · early June — Prove it or stop

Three tests: engagement letters shipping in production, unsanctioned-tool use falling on the spot survey, and weekly active use above 40%. If two of three miss, we stop and rescope rather than push on. You should never fund month four on momentum.

What deliberately isn't in the first 90 days

No Tier 3 client-data integrations (consent first), no knowledge assistant build (corpus first), no firm-wide mandates (adoption is recruited, not ordered). The killed list on page 12 stays killed. Sequencing is the strategy: every phase clears the ground the next one stands on.

\$ Investment & return

Costs first, in full. The assessment fee is credited against the first advisory month if you continue, so nothing below double-counts it.

LINE ITEM	COST	NOTES
Advisory retainer	\$15,000/mo	Quarterly commitment. Assessment fee credited to month one.
Enterprise licenses (92 seats)	≈ \$61K/yr	Claude Enterprise ≈ \$60/seat/mo (78) + Copilot pilot \$30/seat/mo (14).
Engagement-letter workflow build	\$28K	Fixed-fee SOW. The April quick win.
Document-intake automation	≈ \$52K	SOW after \$7216 consent refresh. Targets the 2,100 admin hours.
Knowledge assistant	\$45–65K	Budgeted, not committed — gated on corpus cleanup finishing.

The worked math, without heroics

At stake: **6,850 h/yr × \$85/h blended ≈ \$582K**. I do not assume full capture. At a steady-state 55–65% across the top five, that's **3,750–4,450 h ≈ \$320–380K/yr, recurring**. Year-one all-in cost lands around \$310–345K, then drops as one-time builds roll off. The break-even test is stricter: the top two workflows alone, at 60% capture, return ≈ 2,100 h ≈ **\$178K/yr against \$80K of one-time build cost**. The plan clears the bar even if everything after workflow #2 disappoints.

The honest cost of waiting

"Wait" is not the zero-cost option it looks like. Shadow use grew to 41 of 96 with no policy at all — it does not shrink on its own. Nine people have already pasted client information into free tools; the exposure compounds monthly, and \$7216 penalties attach per disclosure. Meanwhile the 6,850 hours keep getting spent by hand. Waiting buys certainty about only one thing: that both numbers are worse at the next measurement.

Hour savings convert to money only when partners reassign the time to client work or reduce overtime and contractor spend — the monthly report (the advisory deliverable) tracks realized hours against this page so the math stays audited, not aspirational.

→ Recommendation & next steps

Four ways forward. Any of the first three is defensible; my recommendation and the reasoning are below each. The decision belongs to the partner group on Mar 31.

1 · Advisory only

\$15K/mo. Governance, vendor selection, and the roadmap get executive ownership; your team or other vendors build. Works if you have build capacity — today, with two people in IT, you don't.

2 · Implementation only

Fixed-fee builds (\$28K letters, ~\$52K intake) without the advisory seat. Cheapest path to shipped software, but nobody owns adoption, governance, or the vendor decisions — the usual way tools end up unused.

3 · Advisory + phased implementation

RECOMMENDED

The advisory seat owns governance, adoption, and sequencing; builds ship in the order the scores dictate. Everything in this report — the roadmap, the checkpoint, the monthly report — assumes this shape.

4 · Wait

Defensible only if you believe the page 4 numbers will hold still. They won't (page 11). If you wait, at minimum adopt the policy and DLP rule anyway — those cost days, not months.

What happened next (noted for this sample)

At the Mar 31 partner meeting, M&F chose option 3. The advisory retainer started Apr 1, 2026; the engagement-letter workflow shipped May 12 and the killed list (website chatbot, automated 1040 prep, meeting bots on client calls, a fine-tuned custom model, AI-drafted audit opinions) stayed killed. The companion sample, the Monthly AI Executive Report, shows what month three of that engagement looks like.

About this sample

This is the report a \$5,000–\$7,500 AI Executive Assessment buys: two weeks of evidence, a scored roadmap, and a recommendation your partner group can interrogate. Meridian & Frost is a fictional composite client — the method, the math, and the level of candor are exactly what you'd get.

bradtaylorai.com/how-i-work · brad@bradtaylorai.com