

● SAMPLE DELIVERABLE

EXECUTIVE AI ADVISORY · MONTHLY REPORT Nº 3

The Monthly AI Executive Report

Adoption, shipped work, the opportunity backlog, blocked decisions, risk, and the next 30 days — the written record of what your AI executive owns.

PREPARED FOR

Meridian & Frost, PLLC
Dana Whitfield · Marcus Bell

PERIOD

June 1–30, 2026

PREPARED BY

Brad Taylor
Fractional Chief AI Officer

The month at a glance

Month three of the advisory engagement. Adoption is now past the tipping point in three of five departments, the engagement-letter workflow has paid for a fifth of its build cost in saved hours, and shadow AI use has dropped by two-thirds since March. Two things need your attention: the Copilot decision is now costing us focus, and the §7216 consent refresh is on the critical path for the intake automation going live in August.

64%

WEEKLY ACTIVE USERS

▲ from 53% in May

212LETTERS VIA AI
WORKFLOW

▲ 3.5 h → 45 min each

~560 hHOURS SAVED SINCE
MAY 12

▲ run-rate ≈ 1,250 h/yr

12/96STAFF ON
UNSANCTIONED
TOOLS

▼ from 41/96 in March

Three decisions are waiting on you (page 6)

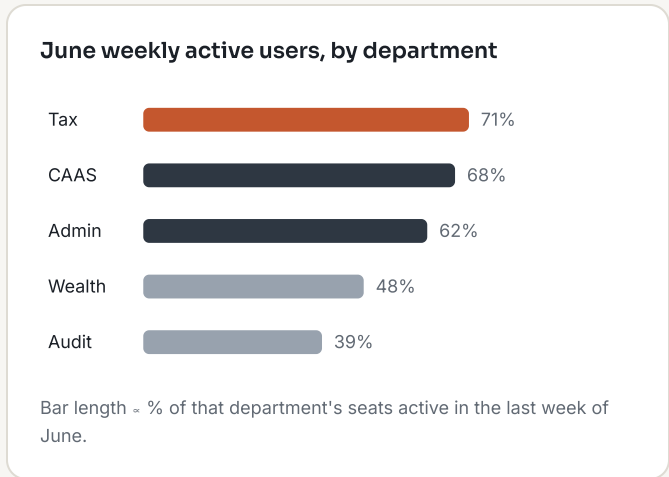
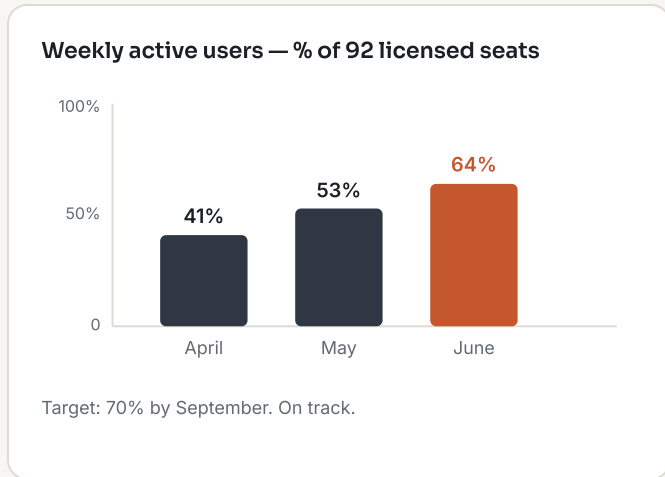
1 — Copilot: expand the CAAS pilot or consolidate on Claude (needed by Jul 15). 2 — AI-assistance disclosure language for client deliverables (partner vote Jul 8). 3 — Intake automation v1 scope: e-signature chasing in or out (my recommendation: out).

In this report

- 3 **Adoption** — who is actually using what, by department
- 4 **Shipped & in flight** — engagement letters, intake automation, knowledge assistant
- 5 **Opportunity backlog** — top five, movement, and the killed list
- 6 **Blocked decisions** — what only you can unblock, with my recommendation on each
- 7 **Governance, risk & vendor notes** — near-misses, model changes, exposure
- 8 **The next 30 days** — dated commitments you can hold me to

Adoption — measured, not assumed

Unused tools have zero ROI, so this section leads the report every month. Usage is licensed-seat activity from the Claude Enterprise and Copilot admin consoles, not self-reporting.



What the numbers say

SIGNAL	READING	ACTION
Tax leads at 71% — Priya Raman's team drove the engagement-letter workflow.	The quick win did its job: a visible result recruits better than any memo.	● Sustain
Audit trails at 39% — lowest WAU, and flat vs May.	Not a failure: audit's Tier 3 workpapers are mostly out of approved scope today. Some of this gap is correct behavior.	● July session
Wealth at 48% but highest depth-per-user.	A small team using it heavily for meeting-prep research. Watch for Tier 3 drift; Omar Haddad is the champion here.	● Monitor

Shipped work & work in flight

Engagement letters: six weeks in production

212 letters and proposals drafted through the workflow since May 12. Average partner-review time is holding at 45 minutes against the old 3.5-hour draft-from-scratch baseline — roughly **560 hours returned to client work**, a fifth of the \$28K build cost already recovered at your blended \$85/hour. Two letters required full rewrites (both unusual fee structures); both are now precedent examples in the library.

Workstream status



WORKSTREAM	OWNER / BUILDER	THIS MONTH	STATUS
Engagement-letter & proposal drafting Backlog #1 · shipped May 12	Priya Raman · Last Rev	Fee-structure edge cases added to the precedent library; house style guide now pins the drafting voice.	● In production
Client document intake automation Backlog #2 · \$52K SOW	Marcus Bell · Last Rev	Karbon and SafeSend integration built; ready for August go-live if the \$7216 consent refresh lands (page 7). Scope decision pending (page 6).	● On plan
Internal knowledge assistant Backlog #3 · \$45–65K budget	Nicole Tran · TBD	Corpus cleanup 40% done: 118 of ~290 duplicate tax templates retired, current-version owners assigned. I am holding the build until the corpus is trustworthy — a knowledge assistant over five versions of the truth answers with the wrong one.	● Held — by design
Copilot pilot (CAAS) 14 seats since April	Elena Ruiz	Useful for Excel narratives; redundant with Claude everywhere else. Decision needed — see page 6.	● Decision pending

Department session this month

June's deep-dive was the **CAAS monthly close** (backlog #4, scheduled Q3). Elena's team walked me through a live close for two clients. Finding: the checklist itself is automatable, but the variance narratives depend on knowing client history that currently lives in three people's heads — so the workflow design will pair drafting with a mandatory reviewer, not replace the reviewer. Session notes went to Elena and Marcus on June 19.

The opportunity backlog — top five of eighteen

Every candidate workflow is scored on impact (annual hours or dollars at stake), feasibility, data tier, and time-to-value. Scores are re-checked monthly; the full 18-item backlog with scoring math is maintained as a separate living document.

#	WORKFLOW	ANNUAL HOURS	TIER	SCORE	STATUS
1	Engagement letters & proposals	1,400 h	YELLOW	 87	● Shipped
2	Client document intake & organizer chase	2,100 h	RED	 84	● In build
3	Internal knowledge assistant	1,800 h	YELLOW	 81	● Corpus prep
4	CAAS close checklist + variance narratives	950 h	RED	 74	● Q3
5	Post-filing client advisory letters	600 h + revenue	RED	 71	● Q4

Movement this month

No rank changes. The June CAAS session (page 4) sharpened #4's design but didn't change its score — the constraint was never feasibility, it's that Elena's reviewers are underwater until the October close cycle eases. #5 gains a revenue estimate next month once Priya's team samples 20 post-filing files for planning opportunities.

⊗ Still dead, on purpose

The killed list stays visible so ideas don't resurrect by hallway conversation: public website chatbot (risk > value) · fully automated 1040 prep (\$7216 + quality) · meeting recorders on client calls (consent) · fine-tuned custom model (cost, no differentiator) · AI-drafted audit opinions (professional standards — never).

Blocked decisions — what only you can unblock

Each of these has an owner, a date, and my written recommendation. If a decision slips, the cost of the slip shows up here next month.

1 · Copilot: expand or consolidate DANA · BY JUL 15

Elena wants 20 more Copilot seats for CAAS. Usage data says Copilot earns its keep only in Excel-heavy close work; for everything else it duplicates Claude at \$30/seat. Running both across the firm costs ~\$33K/year more and splits training, governance, and precedent libraries in two.

My recommendation: keep Copilot capped at the 14 CAAS seats through September, consolidate everyone else on Claude, and revisit with Q3 usage data. Savings fund the knowledge-assistant build.

2 · AI-assistance disclosure in client deliverables PARTNER VOTE · JUL 8

Outside counsel's draft language is ready. Jordan wants disclosure on every AI-touched deliverable; Priya argues that overbroad disclosure implies the firm isn't reviewing its own work. Both are protecting the same thing — client trust.

My recommendation: adopt counsel's middle draft — disclose the firm's AI governance posture in the engagement letter once, not per-deliverable — and revisit if a client or regulator asks for more. This matches current AICPA guidance direction and what I see peer firms adopting.

3 · Intake automation v1 scope MARCUS · BY JUL 22

Adding e-signature chasing to v1 costs +\$9K and +3 weeks, pushing go-live into September — the start of extension season, the worst possible month to change the intake process.

My recommendation: defer to v1.1 in November. An August go-live on the document chase alone still removes the largest admin burden before the fall deadline wave.

Why this section exists

Most stalled AI programs don't die from bad technology — they die from unmade decisions. Publishing the queue, with dates and a recommendation attached, is how we keep decision latency measured in days instead of quarters.



Governance & risk review

ITEM	WHERE IT STANDS	STATUS
Two DLP near-misses — client financials pasted toward a personal AI account, blocked by the M365 rule both times.	Handled as coaching, not discipline, per policy. Both staff had a sanctioned path and didn't know it — that's a training gap, so July's all-hands demo covers exactly this case.	● Closed
\$7216 consent language refresh — required before intake automation touches client tax data.	Counsel engaged June 9; draft due Jul 3. This is the critical path to the August go-live.	● On critical path
Personal Dropbox on one partner team — client files outside governed storage.	Migration to SharePoint committed for Jul 31. Until then that team's folders are excluded from every AI integration.	● Jul 31 deadline
Template version sprawl — ~40% of tax templates had 3+ live versions at assessment.	Cleanup 40% complete (118 of ~290 retired). Blocks the knowledge-assistant build, by design.	● Improving



Vendor & model updates

Claude Enterprise	May model update shifted drafting tone in engagement letters — caught in partner review, not by clients. The house style guide now pins voice and format, and model updates get a 10-letter regression check before rollout. This is the operational reality of AI systems: the vendor will change the model under you; your process has to notice.
Microsoft Copilot	No material product changes affecting the pilot. Decision on page 6.
SafeSend	Vendor announced expanded AI features for organizer processing; I've requested their data-handling addendum before we evaluate. Nothing enabled yet.
Market note	Nothing announced this month changes the hybrid commercial-first architecture. The private-AI revisit triggers (RIA registration scale, residency terms in a client contract) remain unmet.



The next 30 days

Dated commitments. July's report opens by scoring June's list — done, slipped, or dropped, with reasons.

- **Jul 3 — §7216 consent draft from counsel**
I review same week; partner sign-off folded into the Jul 8 vote so we don't burn a second partner meeting.
- **Jul 8 — Partner vote: disclosure language**
Materials go out Jul 1 with my one-page recommendation attached.
- **Jul 10 — All-hands: "what the DLP rule caught"**
Fifteen minutes, real (anonymized) near-misses, the sanctioned path demonstrated live. Champions run department Q&A after.
- **Jul 15 — Copilot decision due (Dana)**
Whichever way it goes, training and the precedent library consolidate accordingly within two weeks.
- **Jul 18 — Audit department session with Jordan Pike**
Scope: what audit *can* use safely today (planning memos, PBC list drafting from prior-year files) versus what stays out (workpapers, opinions). Goal: move audit's WAU without moving its risk.
- **Jul 22 — Intake scope call (Marcus) · Jul 31 — Dropbox migration complete**
Both are go/no-go gates for the August intake go-live.
- **Rolling — corpus cleanup to 65%**
Nicole's team retires the next ~70 duplicate templates; knowledge-assistant build kicks off when trust in the corpus is earned, not before.

About this sample

This is the report every advisory client receives each month — the written record of adoption, shipped work, decisions, and risk that a \$15K/month seat buys. Meridian & Frost is a fictional composite; the format, the discipline, and the level of detail are exactly what you'd get.

bradtaylorai.com/advisory · brad@bradtaylorai.com