SAMPLE DELIVERABLE

BOARD UPDATE · Q2 2026 · PARTNER MEETING JUL 8

The AI program, one quarter in

What we bought, what shipped, what it returned, what it risks — and the three decisions this meeting owns.

PREPARED FOR
The partners of
Meridian & Frost, PLLC

MEETING DATE
July 8, 2026

PREPARED BY
Brad Taylor
Fractional Chief AI Officer

SUMMARY

The quarter on one slide

Every number below has a comparison. Sources: enterprise admin consoles, workflow logs, and the June 30 spot survey.

64% WEEKLY ACTIVE USERS ▲ from 41% in April	**~560 h** HOURS RETURNED TO CLIENT WORK ▲ since May 12 — about 7 weeks	**$99K→** COMMITTED · VALUE RUN-RATE ▲ ≈ $106K/yr and growing	**41→12** SHADOW-AI USERS, OF 96 SURVEYED ▼ down two-thirds since March

The one-sentence verdict

The program is working: adoption is broad and measured, the first workflow already covers its own cost curve, and the firm's AI risk is lower today than the day before we started.

MARCH 2026 · ASSESSMENT

Where we started

The two-week assessment interviewed ten people and surveyed 96 of 118 staff. This is what it found before any policy existed.

41/96
ALREADY USING PUBLIC AI TOOLS
for client work, prior 90 days

9
HAD PASTED CLIENT DATA
3 personal accounts held it in history

0
WRITTEN AI POLICIES
no rules, no tiers, no audit trail

~40%
TAX TEMPLATES WITH 3+ VERSIONS
~290 duplicates on SharePoint

The point partners should hold on to

The risk was already here. Staff didn't wait for permission — nearly half were using ungoverned tools before the firm had a single rule. The program didn't introduce AI to Meridian & Frost; it made what was already happening visible and governed.

DELIVERY

What shipped

Three things went live this quarter. Each one is either producing measured results or on a dated plan.

Engagement-letter workflow

212 letters drafted through it since May 12. Drafting time fell from **3.5 hours to 45 minutes** — a partner reviews and signs every one. The $28K build pays for itself in about **5 months from the May 12 ship — mid-October**, on the strict count of saved hours.

● In production

Document-intake automation

The "doc chase" that eats **2,100 admin hours a year**. Build is on plan for an **August 3 go-live** — before extension season, which is the whole point of the date.

● On plan · Aug 3

The guardrails

Written policy (now v2.1) with three plain-language data tiers, DLP rules that actually block, six named champions, and usage measured weekly from the admin consoles — not from surveys.

● Live since April

What deliberately did not ship

The knowledge assistant is held until the template cleanup makes its answers trustworthy, and five ideas stay on the killed list — including the website chatbot and anything that drafts audit opinions. Saying no is part of the delivery record.

ADOPTION

Who is actually using it

Measured from licensed-seat activity in the enterprise admin consoles — behavior, not self-reporting. 92 licensed seats: Claude Enterprise 78, Copilot pilot 14.

Weekly active users — % of 92 licensed seats

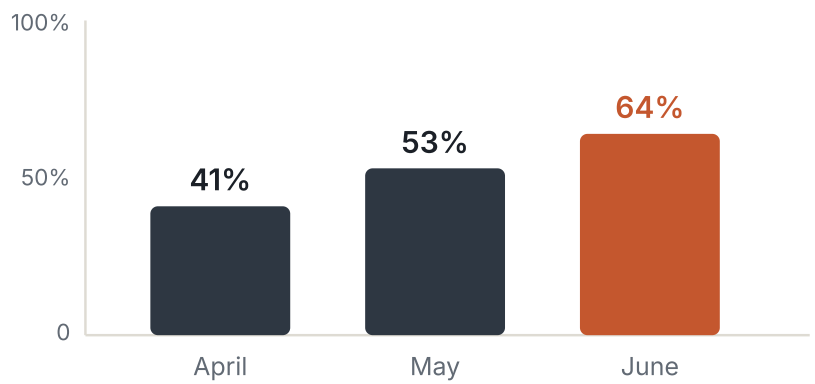

Target: 70% by September. On track — up 12 points in May, 11 more in June.

June weekly active users, by department

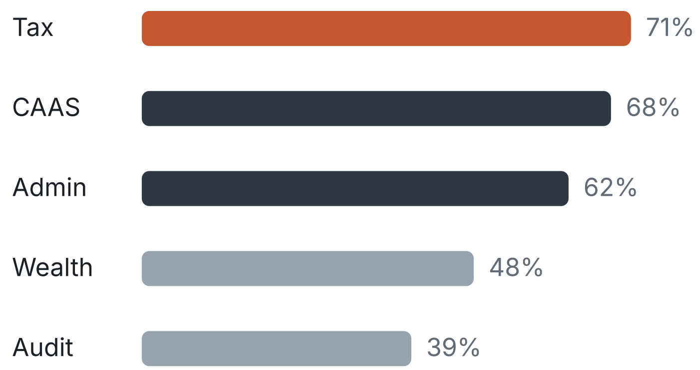

Tax leads because the first workflow landed there — a visible win recruits better than any memo.

Reading the low bar honestly

Audit's 39% is partly **correct behavior**: most audit work sits in Tier 3 data that today's approved scope keeps out. The July 18 session with Jordan Pike widens the safe-list, not the risk.

INVESTMENT & RETURN

What it costs, what it returns

Committed spend through June 30 against the value already running. Hours are valued at the firm's blended $85/hour.

$ Committed through Jun 30 — ≈ $99K

Advisory retainer, Apr–Jun	$45K
Engagement-letter workflow build (shipped)	$28K
Intake automation, progress to date ($52K SOW)	~$26K

Software seats run ≈ $61K/yr, expensed separately as licenses.

Value running today — ≈ $106K/yr

The letters workflow alone returns ≈ **1,250 hours a year** at the steady-state pace — about **$106K/yr** at $85/hour, from a $28K build. (May–June's 212-letter volume is a renewal-season peak, so the run-rate, not a straight-line of the quarter, is the honest number.)

Next step change: intake automation has **2,100 admin hours a year** in scope when it lands August 3.

June verdicts, by line

FUND / EXPAND Engagement letters
FUND Intake automation — on plan
HOLD Knowledge assistant — corpus not ready
HOLD Copilot pilot — decision Jul 15
KILLED Website chatbot — March, stays visible

The discipline behind the tags

Anything that can't name its P&L line gets held, not funded. That rule is why the killed list exists and why the spend above maps one-to-one to measured hours.

RISK & GOVERNANCE

What you're on the hook for — and what's controlled

Partners carry the legal and reputational exposure: §7216 on tax data, professional standards on audit work, and client trust on everything. Here is where each one stands.

EXPOSURE	CONTROL IN PLACE	STATUS
Client data in public AI tools	Zero confirmed instances since May. DLP blocked the only two attempts in June; both handled as coaching, and July's all-hands runs the anonymized cases.	● **Controlled**
No rules of the road	Written policy v2.1 with three data tiers, live since April; every approved tool carries a contractual no-training commitment.	● **Live**
§7216 — tax-return information	Consent-language refresh in flight with outside counsel; it gates the intake go-live, and the go-live waits for it — not the other way around.	● **In flight**
Files outside governed storage	One partner team's personal Dropbox migrates to SharePoint by Jul 31; those folders are excluded from every AI integration until then.	● **Jul 31 deadline**
"Prove what the AI touched"	Every governed interaction is logged in the enterprise consoles — an audit trail exists for the tools we sanction, and none exists for the ones we banned.	● **Logged**

DECISIONS

Three decisions this meeting owns

Each comes with my written recommendation, what a yes costs, and what waiting costs. Stalled AI programs die from unmade decisions, not bad technology.

1 · Copilot: consolidate DANA · BY JUL 15

Recommendation: cap Copilot at the 14 CAAS seats through September; consolidate everyone else on Claude. Revisit with Q3 data.

A yes costs: 20 requested seats not granted; CAAS keeps Copilot where it earns its keep (Excel-heavy close work).

Waiting costs: ~$33K/yr running both firm-wide, plus split training, split governance, and two precedent libraries.

2 · Disclosure language PARTNER VOTE · THIS MEETING

Recommendation: adopt counsel's middle draft — disclose the firm's AI governance posture once, in the engagement letter, not per-deliverable.

A yes costs: one paragraph in every engagement letter and the discipline to keep it true.

Waiting costs: every AI-assisted deliverable ships under an undefined posture — the worst position if a client or regulator asks first.

3 · E-signature chasing MARCUS · BY JUL 22

Recommendation: defer to v1.1 in November. Keep intake v1 scoped to the document chase.

A yes costs: e-signature follow-up stays manual through the fall.

Waiting costs: +$9K and +3 weeks pushes go-live from August 3 into extension season — the worst month to change how documents arrive.

Q3 2026

Next quarter, with dates and owners

The October report to this group scores this list — done, slipped, or dropped, with reasons.

WHEN	WHAT	OWNER
Jul 21–29	Intake automation UAT, then the go/no-go gate review. §7216 consent language must be signed off first.	Devon Marsh (UAT) · Brad (gate)
Aug 3	**Intake automation go-live**, then hypercare through extension season — daily checks in week 1, weekly through October.	Marcus Bell · Last Rev
Aug–Sep	Template cleanup passes the 90% gate → the knowledge-layer build starts (8–10 weeks). The build waits for the corpus, not the calendar.	Nicole Tran (corpus) · Brad (gate)
September	CAAS monthly-close workflow scoping — backlog #4, designed as draft-plus-mandatory-reviewer from the June session findings.	Elena Ruiz · Brad
Monthly	Full evaluation run against the 150-question golden dataset, plus a run at every vendor model update. Results published in the Monthly Report.	Nicole Tran · Brad reviews

What Q3 is really about

Q2 proved the program on one workflow. Q3 moves it from one win to a system: the biggest admin time sink automated, the knowledge layer gated on trustworthy data, and evaluation on a standing cadence.

THE EXTERNAL STORY

How to talk about this with clients

You will be asked — at closings, at conferences, by the client whose competitor just had an AI incident. These four sentences are accurate today and safe to use verbatim.

"We govern AI use under a written policy."

Named tools, three data tiers, technical controls that block mistakes, and a policy document we can show you — not a verbal understanding.

"Client data never enters public AI tools."

Only enterprise systems with contractual no-training commitments touch firm work, and data-loss-prevention rules enforce it automatically.

"AI drafts. Our professionals review and sign."

Nothing AI-assisted leaves the firm without professional review — the same signature standard that has always applied.

"Ask us anything about it."

The posture is a strength, not a secret. A firm that can answer AI questions calmly and specifically stands out in this market.

What not to say

Don't promise "we don't use AI" — it isn't true of any modern firm and it forfeits the story above. Don't improvise specifics beyond these four sentences; route detailed questions to Marcus or to me.

PRESENTER NOTES — DANA

Running the twenty minutes

The clock

Min 0–4 Slides 2–4, fast. The numbers carry themselves — don't narrate every card.

Min 4–10 **Slow down on 5 and 6.** Adoption and money are where skeptics engage; take questions here, not at the end.

Min 10–16 Slide 8 — the decisions. Get an actual yes/no on disclosure in the room; Copilot can close by Jul 15.

Min 16–20 Slides 9–10. End on the client narrative — it's the slide partners will use tomorrow.

Slides 3 and 7 are reference — skip unless asked. Brad attends for Q&A but Dana presents: the program should read as the firm's, not the consultant's.

The four questions to expect

"Are we exposed?" Less than in March: shadow use 41 → 12 of 96, DLP live, zero confirmed client data in public tools since May. Residual items are named on slide 7 with dates.

"What if the vendor changes terms?" The no-training addendum is contractual, the monthly eval run benchmarks alternatives, and the precedent library moves with us — switching is leverage we maintain, not a crisis.

"Why not cheaper seats?" The eval scored platforms against our own work, not a brochure; all licenses together run ~$61K/yr and the first workflow alone returns ~$106K/yr. A cheaper seat nobody uses returns zero.

"When do we see revenue, not hours?" Q4: post-filing advisory letters is the first revenue-side workflow; the first estimate lands next month after Priya's team samples 20 files.

About this sample

This is the board deck a fractional Chief AI Officer prepares for a partner meeting — the numbers, the decisions, and the notes that let the managing partner present it as the firm's own program. Meridian & Frost is a fictional composite; the format is exactly what you'd get.

bradtaylorai.com/advisory · brad@bradtaylorai.com

Related samples: Monthly AI Executive Report — the operating version of this deck | AI Investment & ROI Memo — the money slide's source | Governance & Risk Review — slide 7's source