SAMPLE DELIVERABLE

ADVISORY DELIVERABLE · ENABLEMENT PLAYBOOK · V1.1

The Commercial AI Enablement Playbook

Buying seats is procurement. This is the part that makes them get used — the 60-day sequence that took Meridian & Frost from 41% to 64% weekly active, in order, with the guardrails installed first.

PREPARED FOR

Meridian & Frost, PLLC
Dana Whitfield · Marcus Bell

COVERING

Rollout, Apr 7 – Jun 6, 2026
Claude Enterprise + Copilot pilot

PREPARED BY

Brad Taylor
Fractional Chief AI Officer

The 60-day plan at a glance

Most firms that "rolled out Copilot" or "bought ChatGPT" did exactly one of these eight weeks — the procurement one — and then wondered why nothing happened. The plan below is what actually ran at M&F, in the order it ran.

41→64%
WEEKLY ACTIVE USERS
▲ April → June, admin console

78
CLAUDE ENTERPRISE SEATS
▲ ramped from 40 in 60 days

12/96
STAFF ON UNSANCTIONED TOOLS (OF 96 SURVEYED)
▼ from 41/96 at assessment

212
LETTERS VIA THE WORKFLOW
▲ since May 12, as of Jun 30

The sequence, in order

Before day one · Mar 24 – Apr 6 — Guardrails before users
Acceptable-use policy v1.0 published Mar 24 with the three data tiers in plain language. M365 DLP rules live and tested. Six champions selected, one per department plus marketing — each with a named ownership.

Week 1 · Apr 7–13 — Procurement done right, then the kickoff
Contractual no-training addendum executed Apr 7, **before any user touched the tool**. Entra SSO + SCIM wired, so every seat maps to a person and offboarding is automatic. Seat plan committed: 40 to start, 78 by week 6. All-hands kickoff Apr 8 — 45 minutes, live demo on real firm work. First 40 seats go to tax and admin, where the first workflow lands.

Weeks 1–6 · Apr 9 – May 15 — Department trainings + the precedent library
Role-specific sessions in five waves, tax first (page 4 is the syllabus). Precedent library seeded with 60 approved engagement letters and 25 proposals, so the tool starts from the firm's own voice. Policy v2.0 lands Apr 28, adding the Copilot pilot to the approved list.

Week 6 — Ramp to 78 seats
Remaining seats provisioned once the first cohort's usage proved the training stuck. The engagement-letter workflow ships May 12, five weeks in.

Weeks 7–8 — Make it a habit
Weekly 30-minute office hours per department. Show-and-tell #1: staff demo their own wins to their own peers. Usage dashboard live — WAU by department, from the admin consoles, published monthly.

The sequence matters

The eight weeks aren't a checklist — they're an order. Every step exists to prevent a specific, common way that AI rollouts die. Reverse any of these and you get the failure it prevents.

Guardrails before users

41 of 96 staff were already using public AI tools when we started. You can't ask people to abandon the shadow path until a sanctioned one exists — so DLP, policy, and data tiers went live before the first login, the policy a full two weeks before.

Prevents: the ban-then-blink rollout, where policy arrives after habits form and every rule reads as punishment.

Quick win first

The engagement-letter workflow shipped May 12 — five weeks in, inside the first 45 days. 212 letters later, tax leads the firm at 71% weekly active. A visible result recruits better than any memo.

Prevents: the six-month "exploration phase" where the tool is licensed, discussed, and never load-bearing.

Champions before scale

Six champions were named before day one — before their departments had seats. Each owns something real: Sam Alvarez owns letter edge cases, Grace Liu owns the audit safe-list, Devon Marsh owns intake UAT. Peers teach peers; I teach the champions.

Prevents: the 15%-enthusiast plateau — rollouts that skip this step stall at the people who would have used it anyway.

Measurement from day one

Weekly active users comes from the Claude Enterprise and Copilot admin consoles, not from surveys. Surveys measure enthusiasm; consoles measure behavior. The 41% April baseline was captured in week 1, so every later number has a comparison.

Prevents: discovering in month four — from anecdote — that adoption stalled in month two.

The rule underneath all four

Unused tools have zero ROI. Every step in this playbook is aimed at one number — weekly active users doing real client work inside the guardrails — because every other benefit is downstream of it.

Department prompt patterns — what we actually taught

Generic prompt training doesn't survive contact with a busy season. Each department got two or three starter patterns tied to work they already do, each with its guardrail attached — the tier it runs at and who reviews the output.

DEPT	STARTER PATTERNS TAUGHT	GUARDRAIL
Tax	**Engagement-letter drafting from precedent** — start from the library's closest match, never from scratch. **Research memo first drafts** — with a mandatory citation check against primary sources before anything is filed.	YELLOW Precedent library only; partner review before anything reaches a client. Citations verified by the preparer, initialed.
CAAS	**Variance-narrative drafts** from close data. **Client email drafts** — status updates, document requests, explanation of adjustments.	YELLOW Anonymized or aggregated figures only until the governed Q3 close workflow ships; named-client financials are Red. Reviewer sign-off on every narrative.
Admin	**Intake follow-up drafts** — organizer-chase emails from the standard sequence. **Meeting summaries** — internal meetings only.	YELLOW No recording or summarizing of client calls — consent rules. Client-specific tax detail stays out of prompts.
Wealth	**Meeting-prep research briefs** — market context, planning-topic summaries, agenda drafts from the firm's templates.	YELLOW Public sources and firm templates only; client PII and account data are Red and stay out. Omar Haddad monitors for Tier 3 drift.
Audit	**Planning-memo drafts** and **PBC list drafts from prior-year files** — safe-list items only.	RED-ADJACENT The safe-list (owned by Grace Liu, approved by Jordan Pike) defines what's in. Workpapers and opinions: never.

Patterns are a living asset, not a handout

Every pattern lives in the precedent library next to the examples it draws on, and the set evolves monthly: show-and-tell surfaces what staff invented, champions promote the keepers, and anything that produced a near-miss gets rewritten. The training deck from week 4 is already two revisions out of date — on purpose.

Instrumentation — what we measure

METRIC	SOURCE	WHERE IT'S PUBLISHED
Weekly active users, by department	Claude Enterprise + Copilot admin consoles — behavior, not self-reporting	Monthly AI Executive Report, page 1
Workflow completions — 212 letters through the pipeline since May 12	Workflow logs (Last Rev build)	Monthly Report, shipped-work section
Time per letter — 3.5 h → 45 min against the pre-workflow baseline	Sampled timing, partner-review timestamps	Monthly Report + ROI memo
Near-miss reports — DLP blocks and self-reported almosts	M365 DLP console + the no-blame report channel	Monthly Report, governance section

What didn't stick — and the fix

A playbook that only reports wins is marketing. These are the four things that resisted, and what each one got.

STICKING POINT	THE FIX	STATUS
Wealth: deep but narrow. 48% WAU — a few heavy users, the rest watching.	Depth-vs-breadth watch: pair the heavy users with the holdouts in office hours; Omar Haddad monitors for Tier 3 drift while usage spreads.	● **Monitoring**
Audit: scope fear. 39% WAU — staff unsure what's allowed near workpapers, so they use nothing.	Partly correct behavior — most audit work *is* out of scope. Jul 18 session with Jordan Pike draws the safe-list line brightly so the allowed 20% actually gets used.	● **Jul 18 session**
Two DLP near-misses (Jun 11 CAAS, Jun 23 tax) — client data headed toward a personal account.	Coaching, not discipline, per policy — both staff lacked a known sanctioned path. July all-hands runs the anonymized cases live.	● **Closed**
Copilot pilot: unresolved. 14 CAAS seats; useful in Excel-heavy close work, redundant elsewhere.	Verdict due Jul 15 (Dana). My written recommendation: consolidate on Claude, keep Copilot capped to the pilot through September.	● **Decision pending**

The enablement checklist

Ten questions any firm can score itself against. If you bought seats six months ago and adoption is flat, the "no" answers below are usually why.

- [ ] **No-training addendum signed before the first seat went live?** Data terms are leverage before rollout and a plea after.
- [ ] **SSO and SCIM provisioning wired?** Every seat maps to a person; departures deprovision automatically.
- [ ] **DLP rules and a written policy live before users?** The sanctioned path must exist before you ask anyone to leave the shadow one.
- [ ] **Data tiers defined in plain language?** Staff can say what's green, yellow, and red without opening the policy.
- [ ] **Champions named per department, each owning something real?** Not "AI enthusiasts" — owners of edge cases, safe-lists, cleanup.
- [ ] **Kickoff demo ran on your own work?** A live engagement letter beats any vendor deck.
- [ ] **Quick win shipped in the first 45 days?** One workflow, visibly load-bearing, that staff can point to.
- [ ] **Usage measured from admin consoles, not surveys?** With a baseline captured in week 1 so every number has a comparison.
- [ ] **Standing office hours and a monthly show-and-tell?** Habits form on cadence, not on launch day.
- [ ] **A monthly report where the numbers get published?** Including the ones that didn't stick — or the plateau hides until it's expensive.

About this sample

This is the playbook that turns purchased seats into used ones — the 60-day sequence, the guardrails, the patterns, and the honest list of what resisted. Meridian & Frost is a fictional composite; the sequence and the discipline are exactly what an engagement gets.

bradtaylorai.com/advisory · brad@bradtaylorai.com

Related samples: 90-Day Adoption & Training Plan — the program behind weeks 3–8

Vendor & Model Evaluation Brief — how the platform was chosen

AI Acceptable Use Policy — the guardrails