

● SAMPLE DELIVERABLE

EXECUTIVE AI ADVISORY · VENDOR DECISION

# Vendor & Model Evaluation Brief

Three enterprise AI platforms, scored against your data tiers  
and workflows — never a reseller's quota.

PREPARED FOR

Meridian & Frost, PLLC  
Marcus Bell, COO

DATE

April 10, 2026

PREPARED BY

Brad Taylor  
Fractional Chief AI Officer

## The criteria, weighted for M&F

Three platforms made the shortlist after the March assessment: Claude Enterprise, ChatGPT Enterprise, and Microsoft 365 Copilot. The criteria and weights come from M&F's own facts — the data tiers, the 100-plus-page client files, the precedent-library workflow the engagement-letter build runs on, and a two-person IT team. I hold no reseller agreements and take no commission from any vendor in this brief.

CRITERION	WHY IT MATTERS AT M&F	WEIGHT
<b>Contractual no-training commitment</b>	Tier 2–3 work is impossible without it. Firm policy prohibits any tool that lacks one.	<b>GATE · PASS/FAIL</b>
<b>Long-document handling</b>	Client files run 100+ pages — returns, workpapers, agreements. Summarize-and-truncate fails quietly.	20%
<b>Precedent-library workflow fit</b>	Backlog #1, engagement-letter drafting from precedent, is the highest-value use in the firm.	20%
<b>SSO/SCIM + admin controls</b>	Two-person IT. Provisioning must ride the existing M365 identity stack, not add a user list to babysit.	15%
<b>Vendor stability &amp; roadmap</b>	Training, governance, and the precedent library standardize on this choice. Switching costs grow monthly.	15%
<b>Audit logs &amp; usage visibility</b>	Adoption is measured, not assumed — and \$7216 exposure questions need answerable logs.	10%
<b>M365 integration</b>	Staff live in Outlook, Teams, SharePoint, and Excel. Distance from the work is adoption drag.	10%
<b>Per-seat price</b>	Scored last, on purpose — see the method note below.	10%


### Method: gates first, scores second, price last

The no-training commitment is a gate, not a score — a vendor that fails it is out regardless of capability. Everything else is weighted to M&F's workflows and re-scored on evidence. Price carries the lowest weight because at these seat counts the spread between vendors is under \$30K a year, and a tool nobody uses is expensive at any price.

## The comparison

CRITERION	CLAUDE ENTERPRISE	CHATGPT ENTERPRISE	M365 COPILOT
<b>No-training commitment</b> (gate)	✓ Pass — contractual	✓ Pass — contractual	✓ Pass — contractual
<b>Long-document handling</b>	✓ Full 100+ page files in one pass	🟡 Good; thinner at the top end	✗ Truncates long files
<b>Precedent-library fit</b>	✓ Shared projects & curated prompts	🟡 Workable; weaker curation	✗ Not built for this
<b>SSO/SCIM + admin</b>	✓ Entra ID supported	✓ Entra ID supported	✓ Native to the tenant
<b>Vendor stability</b>	✓ Established enterprise terms	✓ Established enterprise terms	✓ Microsoft
<b>Audit logs</b>	✓ Admin console + export	✓ Compliance API	✓ Purview-native
<b>M365 integration</b>	🟡 Add-in, not native	🟡 Add-in, not native	✓ Lives inside Office
<b>Per-seat price</b>	🟡 \$60/seat/mo	🟡 ≈ \$60/seat/mo	✓ \$30/seat/mo

### Weighted scores



Composite of the seven scored criteria; all three pass the no-training gate. Copilot wins the Excel-native row outright and loses on document depth and precedent work — which is exactly why it stays, capped, where Excel dominates.

### Price at M&F's seat counts

PLATFORM	SEATS	PER SEAT	ANNUAL
<b>Claude Enterprise</b> — firm-wide at full ramp	78	\$60/mo	\$56,160
<b>Microsoft 365 Copilot</b> — CAAS pilot	14	\$30/mo	\$5,040
<b>ChatGPT Enterprise</b> — not selected; shown at the equivalent count	78	≈ \$60/mo	≈ \$56,160
<b>Total as recommended</b> — against the 1,400 assessed hours/yr in backlog #1 alone			<b>\$61,200</b>

## The recommendation

### Claude Enterprise — primary platform SELECTED

Primary for drafting, research, and precedent work across tax, CAAS, wealth, and audit's approved scope. 78 seats at full ramp; the engagement-letter workflow (backlog #1, in build now) runs on it — the only platform that handles M&F's long client files whole.

### Microsoft 365 Copilot — capped 14 SEATS · CAAS

Copilot wins the Excel-native row, and CAAS close work is Excel-native. Capped to Elena Ruiz's 14-seat pilot; expansion is a Q3 call, on usage data, not enthusiasm.

### ChatGPT Enterprise — not selected NOT SELECTED

It passed every gate and scored within four points of Claude — capability parity for M&F's workflows. That is precisely why it's out: a third vendor means a third governance surface, a third training track, and a third precedent library. Parity is not a reason to add overhead.

## Negotiation notes

No-training addendum first	Signed before any seat goes live. Owner: Marcus Bell, by Apr 17.
Seat ramp 40 → 78 over 60 days	Pay for adoption, not ambition. Seats follow the training rollout, department by department.
Quarterly usage true-up	Unused seats surface in a quarterly review, not at renewal. Shelfware gets cut or reassigned.

## Re-evaluation cadence

Re-scored every six months, or immediately on a material model or term change from any vendor — models change under you, without notice.

### Addendum — June 2026 review copy

In May, a Claude model update shifted engagement-letter tone. Partner review caught it; a house style guide now pins the voice, and updates get a 10-letter regression check. The cadence clause earned its keep in two months.

## About this sample

The vendor decision as an advisory client receives it — weighted criteria, auditable scores, and no reseller margin in the outcome. The method is exactly what you'd get.

[bradtaylorai.com/advisory](http://bradtaylorai.com/advisory) · [brad@bradtaylorai.com](mailto:brad@bradtaylorai.com)