AI Bez Obaw

AI READINESS AUDIT · TIER 1

Sample AI Audit *Report*

How a 30-person B2B services firm goes from scattered, individual AI use to a shared and governed capability in 90 days.

Composite subject · Meridian Partners (name changed) · distributed team, three time zones

SAMPLE REPORT · ANONYMIZED COMPOSITE

Prepared by
Michał and Chris Ścibiorek

2026

SECTION 01 · EXECUTIVE SUMMARY

What we found, and the *three moves* that matter most

Meridian already uses AI, but only as individuals pasting into chat. The value is real and trapped in a few heads. Over 90 days, three moves turn that into a shared, governed capability: a team workspace with a prompt library, AI wired into proposals and client meetings, and a content engine. Conservatively this reclaims 55 to 60 hours a month across sales, delivery, and marketing, with one human sign-off on everything that reaches a client.

THE 3 HIGHEST-LEVERAGE MOVES OVER THE NEXT 90 DAYS

1. **Stand up a shared, governed team workspace and ship the ProposalDraft module.** Move people off scattered personal accounts, capture the best prompts in one library, and cut proposal drafting time by roughly half.

2. **Standardize one meeting notetaker and wire in MeetingLedger.** Every client call ends with a summary, a decision log, and an owner-tagged action list, routed into the project workspace automatically.

3. **Turn delivered work into content with ContentRepurposer.** Lift marketing output from about two to about six published assets a month with the same two-person team.

WHERE THE HOURS ARE TODAY

WORKFLOW	NOW (HRS / MONTH)	RECLAIMABLE (HRS / MONTH)
Proposals and scoping	~38	~22
Client meeting follow-up	~32	~20
Marketing content production	~30	~15
Total across the three workflows	**~100**	**~57**

ROUGH TIMING

Foundation and the ProposalDraft module land in the first 30 days. Meetings, plus the automation and knowledge layer, come by day 60. Content and the first governance review close out day 90. Most of the reclaimed hours arrive in months two and three, once the workspace and the first module are in daily use. Figures here are sample values for the composite, not a promise for any one firm.

SECTION 02 · MATURITY ASSESSMENT

Where the firm *stands today*

AI use at Meridian is genuine but individual and manual. A handful of people get real value from chat drafting. Nothing is shared, connected to the firm's own knowledge, or automated. The firm sits at an early but real stage: past curiosity, short of a system.

TOOLS IN USE

ChatGPT, individual accounts, a few paid. Drafting and brainstorming.
Fathom notetaker, on some calls, used by two people.
Grammarly for editing. **Canva** for marketing visuals.

RELEVANT, NOT YET ADOPTED

A **shared team AI workspace** with saved prompts and projects. A **retrieval layer** over the firm's own proposals and playbooks. **Automation glue** (Zapier or Make) between CRM, docs, and the notetaker. A written **AI use policy**.

GAP MAP

The core gap is the jump from "some people paste into chat" to "the team's repeated workflows run through a shared, governed setup." Value lives in single heads with no prompt library. AI is not connected to Meridian's own material, so a new proposal starts from a blank page rather than from the last ten won ones. No workflow is automated, and there are no data-handling rules, which quietly caps how far anyone will push the tools.

AI USAGE PROFILE · AEI SCORECARD

DIMENSION	SAMPLE FINDING	READ
Artifact mix	Almost entirely first-draft prose: proposals, emails, summaries. Little structured-data work and no automation or build artifacts.	Narrow
Work / personal split	About 80 percent work, 20 percent personal. Use is individual; nothing is shared team-wide.	80 / 20
Automation share	Under 10 percent of repeatable workflows touch AI, and those are manual copy-paste, not wired into any tool.	< 10%
Autonomy level	AI drafts, a human reviews and sends every time. No supervised autonomous runs, no scheduled jobs.	2 / 5
Highest-value workflow	Proposal and scope drafting. Senior staff save two to three hours per proposal today, ad hoc and unshared.	Proposals

Overall maturity · Level 2 of 5 · Emerging

SECTION 03 · LEVERAGE OPPORTUNITIES

Three named modules, in *detail*

Each module has defined inputs, a defined output, and one named owner. AI does the repeatable draft. A person keeps judgment and the client-facing decision.

1 Proposal and scope engine

Effort · Medium

module · ProposalDraft

Every proposal and statement of work is written from scratch. About nine go out a month, roughly four hours of senior time each.

INPUTS	Discovery-call notes, two or three past won proposals, current rate card.	OUTPUT	First-draft proposal, scope of work, and a pricing table in the firm template.
OWNER	Business development lead.	HUMAN KEEPS	Final scope, pricing judgment, and sign-off.

Expected outcome: drafting drops from about four hours to about ninety minutes. Roughly 22 hours a month reclaimed across nine proposals, and a more consistent house style.

2 Meeting to action

Effort · Small

module · MeetingLedger

Notes and follow-ups are written by hand after client calls. Decisions and actions scatter across inboxes, and knowledge is hard to find later.

INPUTS	One standard notetaker transcript plus the project context.	OUTPUT	Summary, decision log, owner-tagged action list, and a draft client update.
OWNER	Engagement manager.	HUMAN KEEPS	Sending the update and judgment on sensitive items.

Expected outcome: about 20 minutes saved per client meeting, near 20 hours a month, plus a searchable record that speeds new-hire onboarding.

3 Content repurposing

Effort · Medium

module · ContentRepurposer

Thought leadership and case studies are produced slowly. One long-form piece takes close to a day, so cadence stays low at about two assets a month.

INPUTS	A recorded webinar or a delivered-project retro, plus SME voice notes.	OUTPUT	Draft article, five social posts, and a newsletter section, checked against brand rules.
OWNER	Marketing lead.	HUMAN KEEPS	Final edit and the publish decision.

Expected outcome: cadence rises from about two to about six published assets a month with the same two-person team, roughly 15 hours a month freed for higher-value work.

SECTION 04 · 90-DAY ROADMAP

One quarter, three *phases*

Days 1 to 30 Foundation

Stand up the shared team workspace and move people off personal accounts. Write a one-page AI use policy and data-handling rules. Build the first prompt library from what top performers already do. Ship ProposalDraft as a template and pilot it on three real proposals.

Days 31 to 60 Wire it in

Standardize one notetaker across client calls and route summaries and action lists into the project space (MeetingLedger). Connect the notetaker, CRM, and docs with a light automation. Begin the retrieval layer by indexing past won proposals and playbooks so ProposalDraft pulls from real firm material.

Days 61 to 90 Scale and govern

Bring ContentRepurposer live and lift cadence toward six assets a month. Run two short training sessions with the prompt library as the reference. Hold the first governance review: what worked, tool spend, what to keep or kill, and where to allow more autonomy.

MILESTONES AND OWNERS

WEEK	MILESTONE	OWNER
W1	Team workspace live, AI use policy drafted	Operations lead
W2	Prompt library v1 (10 prompts) published	Delivery lead
W3–4	ProposalDraft piloted on 3 live proposals	Business development lead
W5–6	One notetaker standard, summaries auto-routed	Engagement managers
W7	CRM, docs, and notetaker automation live	Operations lead
W8	Proposal knowledge base indexed	Delivery lead
W9–10	ContentRepurposer live, first 4 assets shipped	Marketing lead
W11	Team training done, prompt library handed off	Delivery lead
W12	90-day review and governance checkpoint	Managing partner

SECTION 05 · RISK AND GOVERNANCE

Guardrails for the *first quarter*

DATA HANDLING

Client PII, signed contracts, and credentials never go into a consumer chat account. Use the team plan with training turned off and data controls on. Redact identifying details before pasting. Publish the one-page rule and put it in onboarding so it is understood, not assumed.

VENDOR LOCK-IN

Keep prompts and the knowledge base portable, in plain text and in a store the firm owns, rather than only inside one vendor's proprietary features. Standardize on one notetaker for consistency but keep transcripts exportable. The goal is that switching a tool later costs days, not a rebuild.

ESCALATION AND REVIEW

Each module has one named owner. Anything client-facing gets a human sign-off before it leaves the building. When output is wrong or uncertain, the owner catches it and the case feeds back into the prompt library. Define upfront what never ships without review: pricing, legal wording, and anything sent under a partner's name.

TEAM TRAINING

Two short sessions in the first quarter, with the prompt library as a living reference rather than a slide deck read once. New hires get an onboarding checklist that includes the AI policy and the three modules. Training is measured by whether people actually use the shared workspace, not by attendance.

WHAT WE WOULD AVOID IN THE FIRST 90 DAYS

Do not automate a messy process; fix the step first. Do not give autonomy on client-facing sends before the human review habit is set. Do not roll out to all 30 people before the pilot proves out on three proposals. Do not buy five tools when one team workspace and one automation tool cover most of the value.

APPENDIX · METHODOLOGY AND SOURCES

How this audit was *produced*

> *This sample mirrors how we run a real audit. Discovery interviews across roles, a review of real work artifacts, a tool inventory, then an adversarial verification pass where every claim is checked against a source before it reaches you. AI accelerates the work. A person verifies it against the evidence.*

WHAT WE CHECKED AGAINST WHAT

- Interview claims cross-checked against real artifacts: did a proposal actually take that long, timed against three recent ones.
- The tool list confirmed against the firm's billing and admin console, not against memory.
- Hours estimates confirmed against a small sample of recent, dated work rather than round-number guesses.

WHAT VERIFICATION CHANGED

→ An early read said the firm has **no notetaker**. Two people use Fathom inconsistently, so it moved from "not adopted" to "adopted but not standardized."

→ A first estimate put proposal time at **six hours**. Timed against three real proposals it was closer to four, so the savings claim was lowered to match.

→ We dropped an unsupported line that AI would **cut proposal time by 70 percent**. The measured, conservative figure is about half, and only on the drafting portion.

HONEST NOTE ON CERTAINTY

High confidence on the workflow map and on where the hours go, since both are grounded in the artifacts reviewed. Lower confidence on the exact monthly hours, which come from a few weeks of sample data, and on tooling cost, which follows list prices and will move with team size and plan. Where a number is an estimate, it is written as a range, not a promise.

ILLUSTRATIVE SOURCES FOR A REAL ENGAGEMENT

Discovery interviews. BD, delivery, marketing, and operations leads.

Work artifacts. Recent proposals, meeting notes, content pieces.

Tool and billing review. Admin console and subscription list.

Time sample. Dated logs from a few recent weeks.

ABOUT THIS SAMPLE
Meridian Partners is a fictional but realistic composite drawn from real Tier 1 audits, with all identifying detail removed. Every number here is a sample value chosen to be modest and internally consistent, not a result claimed for any specific client.

Prepared by **Michał and Chris Ścibiorek** · sample deliverable · anonymized composite · aibezobaw.pl