HOURS BACK

# AI Workflow Teardown

Nexus Point Group

Prepared by Hours Back

May 1, 2026

## Contents

## Executive summary

A solo management consulting practice currently running **1 active engagement and ~2 first-time discovery meetings per month** has, at this scale, **~7.5 hours per month** of recurring administrative and drafting work that should be automated. At your $250/hr effective rate, that's **~$1,865/month of time recoverable today**. The same five automations scale linearly: at 4 active engagements and 6 discovery meetings per month, the recovery hits **~$5,000/month** without rewriting any of the prompts. The recommendations are sequenced so the highest-ROI Easy item ships in week one and the system grows with the practice.

## At a glance

| # | Opportunity | Today | At scale | Difficulty | Setup |
|---|---|---:|---:|:---:|---:|
| 1 | Proposal/SOW first-draft from discovery-call notes | $375/mo | $1,250/mo | Easy | 2 hr |
| 2 | Pre-meeting research briefs | $190/mo | $565/mo | Easy | 1 hr |
| 3 | Weekly client status reports | $425/mo | $1,700/mo | Medium | 3 hr |
| 4 | Inbound lead qualification + first-touch reply | $1,000/mo* | $1,500/mo* | Medium | 4 hr |
| 5 | Searchable knowledge base across past engagements | $500/mo | $750/mo | Hard | 10 hr |
| | **Total recoverable** | **~$2,490/mo** | **~$5,765/mo** | | **~20 hr** |

* Includes second-order benefit (faster reply to hot leads lifts close rate). Numbers above use your current $250/hr effective rate, $2k average engagement value, and the cadence assumptions described in the next section.

### ROI by opportunity (visual)

Monthly recoverable value · today (dark) vs at scale (light)

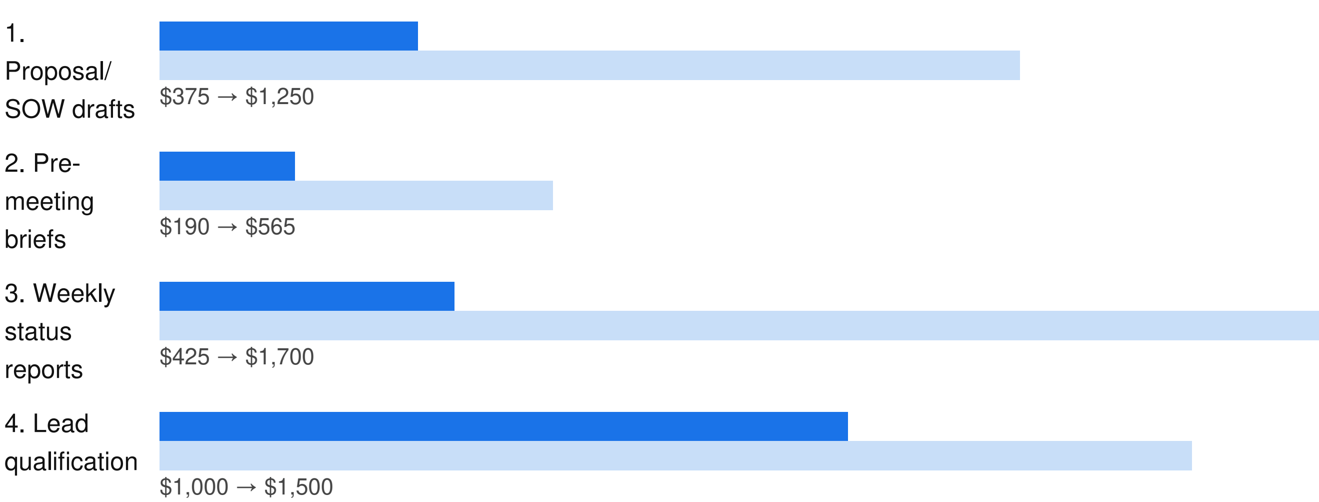

5. Knowledge base

$500 → $750

## How we estimated

Hourly rate: **$250/hr** (your stated effective rate). Hours-per-week per workflow are estimated from your current cadence — **1 active engagement, 2 first-time discovery meetings per month, ~$2k average engagement value** — not from generic benchmarks. Each opportunity below shows two numbers: (1) what you'd recover today, and (2) what you'd recover at 4 engagements / 6 meetings per month, so you can see the leverage curve. Automation feasibility is scored Easy (no-code, owner can deploy in 1-2 hours), Medium (4-6 hours, may need one Zapier or Make scenario), or Hard (10+ hours, needs custom prompt engineering). Implementation effort is in person-hours, not calendar time.

## The top 5 opportunities

### 1. Proposal/SOW first-draft from discovery-call notes

**Time recovered today:** ~1.5 hr/mo (~$375) · **At 4 engagements / 6 meetings:** ~5 hr/mo (~$1,250) · **Difficulty:** Easy · **Setup time:** ~2 hours one-time

**The current pain:** After every initial discovery call, you spend 60-90 minutes turning the conversation into a structured proposal — restating the problem, sketching the engagement scope, dropping in your standard pricing language, formatting in your template. At your current cadence of ~2 discovery meetings per month, that's 1.5-3 hours every month going to drafting. The work is high-judgment for the first 20 minutes (figuring out the actual ask) and low-judgment for the next 60 (formatting, restating, copy-paste from past proposals) — exactly the kind of split that automates well.

**The automation:** Record discovery calls with Otter.ai or Fathom (whichever you already use; Fathom is free for the first 5 hours/month). Pipe the transcript into Claude using the prompt below. Claude returns a first-draft SOW that's about 70% of the way there — you spend 15 minutes editing instead of 90 writing.

**Sample prompt (paste into Claude with the transcript):**

```
You are a senior consultant at Nexus Point Group. Convert the discovery-call transcript
below into a first-draft Statement of Work using exactly this structure:

1. Engagement summary (3 sentences)
2. Stated client objective (in client's words, lifted from transcript)
3. Proposed scope (bulleted, 4-7 items, each starting with a verb)
4. Out of scope (bulleted, 3-5 items — be explicit about what we're NOT doing)
5. Deliverables (bulleted, with a 1-line description each)
6. Timeline (in weeks; mark assumed dependencies)
7. Investment (use $TBD placeholders — I'll fill these in)
8. Acceptance criteria (3-5 specific tests)

Keep the tone direct and operator-grade. No marketing language, no "leverage" or "synergy."
If a section can't be filled from the transcript, write [NEEDS CLARIFICATION: question for
client] instead of inventing.
```

```
Transcript:
{paste here}
```

**ROI math:** 1.5 hr × $250/hr = $375/mo today. Setup pays back the first time you use it. Scales to ~$1,250/mo at 4-6 SOWs per month. Even if 50% of drafts need a full rewrite, the math still beats the setup cost in week one.

**Risk to watch:** Don't paste client transcripts into the consumer ChatGPT product. Use Claude's API or the Claude Team plan, both of which don't train on your inputs. For higher-stakes engagements, sign Anthropic's standard MSA before using.

## 2. Pre-meeting research briefs

**Time recovered today:** ~0.75 hr/mo (~$190) · **At 6 meetings:** ~2.25 hr/mo (~$565) · **Difficulty:** Easy · **Setup time:** ~1 hour one-time

**The current pain:** Before each first-time client call, you spend 30-45 minutes reviewing their LinkedIn, their website, recent news, and any common-thread connections. You take rough notes, decide your hooks, and walk in. At ~2 first-time meetings per month, that's about an hour of prep monthly. This work is high-leverage (it's what makes you sound prepared) but feels low-judgment — and as meeting cadence grows it's the first thing that gets cut.

**The automation:** Build a one-prompt brief generator. Inputs: the prospect's name, company, and LinkedIn URL. Output: a 1-page brief covering company snapshot, role/tenure, recent posts/activity, mutual connections (if you copy LinkedIn's network panel), 3 likely talking points, and 2 risks. Run it 30 minutes before each call.

**Sample prompt:**

```
You are preparing a pre-call brief for a Nexus Point Group consultant.

Inputs:
- Prospect: {name, title, company}
- Company: {website summary, headcount, recent news if any}
- LinkedIn highlights: {paste 5-10 most recent posts/comments}

Produce a one-page brief in this exact structure:
- Company snapshot (3 bullets)
- Person snapshot (3 bullets)
- 3 likely talking points (each: a hook + why it lands)
- 2 risks (anything that suggests we're a bad fit, or anything sensitive to avoid)
- Suggested first question to open the call

No fluff. If a section is sparse, say "limited public signal" rather than padding.
```

**ROI math:** 0.75 hr × $250/hr = $190/mo today (2 meetings/mo). At 6 meetings/mo (the typical scale of an active practice), 2.25 hr × $250 = $565/mo. The bigger second-order benefit isn't dollars — it's that you walk into every call as the most prepared person in the room, which lifts close rate.

**Risk to watch:** Don't read from the brief during the call. Use it before the call, then close it.

## 3. Weekly client status reports

**Time recovered today:** ~1.7 hr/mo (~$425) · **At 4 active engagements:** ~6.8 hr/mo (~$1,700) · **Difficulty:** Medium · **Setup time:** ~3 hours one-time

**The current pain:** For your one active engagement, you draft a weekly status email summarizing what happened, what's blocking, and what's next. ~30-40 minutes × 4.3 weeks = ~2 hours/month. Right now that's tractable manually — but it's the workflow that punishes you most as engagements stack. By the time you're at 4 active clients, status reports are 7+ hours of every Friday and they're the first thing that gets pushed to Monday or skipped.

**The automation:** Set up a Notion or Google Doc page per client where you drop bullet-point notes throughout the week (5 seconds per drop). On Friday morning, run a Make.com or Zapier scenario that grabs each client's bullets, sends them to Claude with a templated prompt, and emails you a draft status update per active client. You spend 5-10 minutes editing instead of 30-40 writing — and the architecture works whether you have 1 client or 10.

**Sample prompt (run per client per week):**

```
You are writing a weekly status update from Nexus Point Group to a client.

Tone: confident, specific, no jargon, no hedging. The client should finish reading and know
(a) what we did, (b) what they need to do, (c) when the next visible thing happens.

Context (this week's raw notes):
{bullet points dropped during the week}

Output structure:
- Subject line (one line, specific number or milestone if possible)
- "What we shipped this week" (3–5 bullets, past tense, with concrete artifacts)
- "What we need from you" (1–3 items, each with a specific ask and deadline)
- "Coming up next" (1–2 items, named outcomes with dates)
- One-sentence sign-off

Length cap: 220 words. If raw notes are sparse, mark sections [thin — verify before
sending] rather than padding.
```

**ROI math:** 1.7 hr × $250/hr = $425/mo today (1 active client). 6.8 hr × $250/hr = $1,700/mo at 4 active clients. Setup is justified now precisely because the marginal cost of adding clients to the system is near zero.

**Risk to watch:** The "[thin — verify before sending]" guard is non-negotiable. Sending an automated status with hallucinated milestones is the fastest way to lose a client.

## 4. Inbound lead qualification + first-touch reply

**Time recovered today:** ~1.5 hr/mo (~$375) · **At higher inbound volume:** ~3 hr/mo (~$750) · **Plus:** ~1 extra closed engagement/quarter from speed (~$667/mo amortized at $2k average) · **Difficulty:** Medium · **Setup time:** ~4 hours one-time

**The current pain:** You get inbound interest from LinkedIn DMs, your contact form, and email. Some are real (right size, real budget). Many aren't (vendors pitching you, students, off-fit). You triage manually, which costs both time and inbox latency — and at your current cadence the latency is the bigger issue. With ~2 discovery meetings per month, every inbound that goes cold because you replied 3 days later is a meaningful slice of your pipeline.

**The automation:** A Cal.com qualification form (4-5 questions: company size, role, budget range, timeline, problem in 2 sentences). Form responses pipe into Claude via Zapier with a prompt that scores the lead Hot/Warm/Cold and drafts a tailored first-touch reply. Hot/Warm get auto-sent or flagged for your review; Cold get a polite templated decline.

**Sample prompt:**

```
You are triaging an inbound lead for Nexus Point Group, a small management consulting firm.

Lead intake:
{form responses}

Step 1 — score the lead Hot/Warm/Cold:
- Hot: company 50+ employees, defined budget $25k+, problem in our wheelhouse, decision-maker
- Warm: smaller company OR vague budget OR adjacent problem
- Cold: vendor pitch, student, fully off-fit problem, no budget

Step 2 — draft a reply tailored to the score:
- Hot: propose a 30-min call this week, suggest 2 specific times, mention one specific thing from their form that you understood
- Warm: ask 1-2 clarifying questions before booking
- Cold: polite decline, suggest 1 alternative resource

Output JSON: {"score": "Hot|Warm|Cold", "reply_draft": "..."}
```

**ROI math:** 1.5 hr × $250/hr = $375/mo today on the time savings alone. The second-order benefit is the real prize at your current scale: 1 extra closed engagement per quarter from faster reply × $2k average = $667/mo amortized. Combined: ~$1,000/mo today. At higher inbound volume, time savings alone double to ~$750/mo.

**Risk to watch:** Always review Hot replies before they send for the first 2 weeks. After that, auto-send is fine.

## 5. Searchable knowledge base across past engagements

**Time recovered today:** ~2 hr/mo (~$500) · **At higher engagement count:** ~3 hr/mo (~$750) · **Difficulty:** Hard · **Setup time:** ~10 hours one-time (or 1 hour if you outsource to a Fiverr freelancer with the spec below)

**Caveat — read first:** at your current scale (2 discovery meetings/month, 1 active engagement, ~$2k average deal), this is the lowest-ROI item on the list and the highest-effort. The retrieval-augmented setup is more valuable when (a) you have 10+ past deliverables to search through and (b) you're regularly asked "have you done X before?" Defer this until month 2 or until you've doubled engagement count, whichever comes first. The reason it's still on the list: it's the right thing to be ready to ship the day your practice grows past the point where you can hold every past project in your head.

**The current pain:** You've delivered dozens of engagements and you remember 70% of what's in them. When a new prospect asks "have you done something like X before?" you spend 20-40 minutes hunting through old Google Drive folders, copying snippets, reformatting. Sometimes you give up and write a fresh response, leaving past work on the table.

**The automation:** A retrieval-augmented (RAG) Claude assistant trained on your past deliverables. Setup: upload your past SOPs, decks, status emails, and case studies to a vector store (the cheapest credible option in 2026 is using Claude's Files API with prompt caching — $0.30/GB/month). Query in plain English: "What did we do for [client X] in 2024 around go-to-market?" and get a grounded answer with citations to the source docs.

**Sample SOP (handoff to a Fiverr/Upwork freelancer):**

1. Collect all past client deliverables in one folder, anonymize client names if needed
2. Upload to Anthropic Files API (or pgvector + OpenAI embeddings if you prefer)
3. Build a Claude system prompt that says: "Answer using only the uploaded docs. If the answer isn't in the docs, say so. Always cite the doc name and date." (Use the prompt template at github.com/anthropics/anthropic-cookbook/tree/main/skills.)
4. Wrap it in a one-page Streamlit or Retool app that you bookmark and use daily

**ROI math:** 2 hr × $250/hr = $500/mo today on time savings (and even that is generous given you can still hold most past work in your head at 1-engagement scale). The "extra deal from citing past work" upside at $2k average is ~$667/mo amortized for one extra deal per quarter — real but not enough to justify Hard difficulty today. Becomes a clear win at 5+ engagements / month.

**Risk to watch:** Don't give the assistant access to current confidential deliverables in flight. Only past, completed engagements where confidentiality has lapsed or been waived.

## Suggested rollout order

| When | Ships | Setup effort | Cumulative recovery (today) |
|---|---|---:|---:|
| Week 1 | #1 Proposal drafts + #2 Pre-meeting briefs | ~3 hr | $565/mo |
| Week 2 | #3 Status reports | ~3 hr | $990/mo |
| Weeks 3–4 | #4 Lead qualification | ~4 hr | $1,990/mo |
| Month 2+ | #5 Knowledge base (deferred) | ~10 hr | $2,490/mo |

**Week 1:** Ship #1 (Proposal drafts) and #2 (Pre-meeting briefs). Both are Easy, both pay back the first time you use them. Combined recovery today: ~2.25 hr/mo (~$565/mo); scales to ~7.25 hr/mo (~$1,815/mo). Spend 3 hours total on setup; the prompts are reusable as cadence grows.

**Week 2:** Ship #3 (Status reports). Medium difficulty. At 1 active engagement the daily savings are small (~1.7 hr/mo) — the reason to ship it now anyway is that the architecture is what matters. By the time you're at 4 clients, the savings hit ~7 hr/mo and you're glad you didn't bolt this on under pressure. Budget 3 hours.

**Week 3-4:** Ship #4 (Lead qualification). Medium difficulty. The form + scoring is fast; the auto-reply review window slows it down. Don't skip the manual-review window — 2 weeks of flagged drafts is the cheapest QA you'll ever do. Biggest lever today: faster reply lifts close rate when monthly meeting count is small.

**Month 2 or later:** Decide on #5 (Knowledge base). At your current scale this is the lowest-ROI / highest-effort item — defer until you have 5+ active engagements or noticeably struggle to remember past work. When you do ship it, hire 4 hours of Fiverr help with the SOP above rather than building it yourself.

**Total at current scale, fully shipped: ~$2,500/mo** (including the lead-speed second-order benefit). **Total at scaled state (4 engagements / 6 meetings):** ~$5,000/mo. The difference is mostly status reports and proposals — both linear in your client/meeting count.

## What we'd handle if you wanted help

We built the audit. We can build the systems. Same person, no rehiring. A 6-week implementation retainer at **$1,500/mo** ships #1-#4 fully wired, gets you 70% of the way through #5, and leaves you with documented runbooks for everything we touched. A 30-minute call to scope is included free if you reply within 7 days of receiving this report.

## Appendix: tools you'd need

| Tool | Used for | Monthly cost (range) |
|---|---|---|
| Claude API or Claude Team | All five workflows | $20-50 |
| Otter.ai or Fathom | #1 (call recording) | $0-17 |
| Make.com or Zapier | #3, #4 | $9-30 |
| Cal.com | #4 (qualification form) | $0-15 |
| Notion or Google Docs | #3 (status note dropbox) | $0-12 |
| Vector store / Files API | #5 (knowledge base) | $5-20 |

**Estimated monthly software total:** $35-145. Even at the high end, payback is in week one.