Iterategy

ILLUSTRATIVE SAMPLE, FICTIONAL BUSINESS

AI Operations Audit: Sample

Prepared for

Bramblewood Bookkeeping and Advisory, LLC

Prepared **2026-09-01**
Order **SAMPLE-BRAMBLEWOOD**

Contents

Every section starts on its own page. Page numbers below are the numbers printed at the foot of each page.

26 pages including the cover. The cover is not numbered, so the numbers above are the ones printed at the foot of each page.

1. Read this first

This sample audit describes a business that does not exist. Bramblewood Bookkeeping and Advisory, its owner Renata Kessler, and her bookkeeper Deshawn are invented, and every number below is fabricated but internally consistent. It is published to show what the real thing looks like. No real client is referenced anywhere in it.

What this is. You bought a written report that reads your intake answers, counts the hours in your week, and puts a dollar figure and the arithmetic behind every finding. This section is the whole audit in plain language. It takes about five minutes and it stands on its own.

What we found. Six things, in the order they pay you back.

1. **Chasing clients for documents is not unfixable. It is unautomated.** You told us it is the task everyone hates most and that you do not think it is fixable. It is worth **$4,900 to $9,850 a year**, and you are partly right about why.
2. **The line between what a rule can do and what your judgment has to do.** Categorization and reconciliation is the biggest block of hours in the firm, yours and Deshawn's together. Two thirds of the transactions on your messiest file are vendors that were there last month. Worth **$4,395 to $9,390 a year**.
3. **Thirty-eight close summaries written from a blank page every month.** Worth **$2,775 to $6,300 a year**, and part of that gain you have already banked yourself.
4. **New client details typed into three systems.** Worth about **$102 a year**. We are counting it at zero and telling you to leave it alone. You were right.
5. **What you actually pay for software.** You said a couple hundred a month. It is **$417**. There is no saving in that sentence, but there is something you should see.
6. **The three or four more clients you asked about.** We are counting this at **zero dollars** and answering it in hours instead, because hours are the thing we can check.

What matters most, and why it is first. Start with the document chasing, not with the bigger number underneath it. Not because it is worth more, but because the first move takes about twenty minutes, costs nothing, and uses a feature you are already paying for and have switched on for four clients out of fifteen. The categorization finding is worth about the same and takes a couple of evenings of setup. Do the cheap one first.

What it adds up to. Between **7.0 and 13.4 hours a week**, worth **$12,070 to $25,540 a year**. Even counting only what we're certain of: $7,170/yr.

Now the part you should read twice. You asked for ten hours a week. **The low end of what we found is 7.0 hours, and only 4.8 of those are yours.** The rest are Deshawn's, and Deshawn's hours cannot review a new client's books. So the honest answer to your question is that the low end does not reach ten, and the top end comfortably does, and where you land inside that range depends almost entirely on how far you take the second finding.

And the Saturdays. You said you would take some of the time back for yourself and stop working Saturday mornings in tax season, and put the rest toward three or four more clients. Those two things draw on the same recovered hours. Averaged over the year the low end pays for both, with about sixty hours to spare. **Inside tax season it does not.** Your twelve Saturdays are 48 hours falling in twelve consecutive weeks, and those are the same weeks your volume peaks and Deshawn is already at thirty hours. So it is a real choice, not a free one, and you should make it deliberately rather than discover it in February.

What happens next. Section 6 is a thirty-day plan in the order we would do it, free and fast first. Your questions are open for two weeks from today: reply to the email this came with and we answer. If you find a factual error, tell us and a corrected copy is back to you inside two business days at no charge.

2. What we looked at, and how

What you gave us. Your intake form, submitted 2026-09-01, all thirty-four answers. Then a follow-up set of twenty-six questions, because the form named every task in your week and quantified none of them. Your answers to those questions are the backbone of this report, and everywhere a number came from you rather than from a system we have said so in the words “you told us”.

What you gave us access to. You granted read-only access to your QuickBooks Online Accountant view and your Financial Cents admin view. Both are your own screens rather than any client’s books, and that distinction runs through every finding in this report. The transaction counts below came out of the QuickBooks Accountant view. Nothing in this audit opened a client file.

What we counted, and how you can re-run it.

What	Value	How to get it yourself
Transactions a month, whole book	about 9,500	QuickBooks Online Accountant, your client list, transaction volume by client
Transactions on the three messiest files	700 to 900 each	same view, filtered to those three
Bank rules that exist today	about 12, across 5 or 6 of 34 files	in each client file, Transactions, Banking, Rules
Clients on Dext	15 of 38, on a tier covering 25	Dext practice dashboard, client list
Clients with Dext client requests switched on	4 of the 15	Dext, per client, notification settings
Receipts a month through Dext	about 600	Dext, document count, last full month
Financial Cents seats	2, on the Team plan	your Financial Cents billing page

How a month moves through your firm today. A client signs after a twenty-minute Calendly call and a PandaDoc engagement letter. Their details get typed into PandaDoc, then again into QuickBooks to open the company file, then again into your Client Tracker sheet. Then the monthly cycle starts and it is the same for all thirty-eight: transactions categorized weekly or every other week, reconciliation at month end, anything unclear chased down, and a close summary email out by the tenth of the following month. Invoicing runs on autopay through QuickBooks Payments. In your words, that part runs on rails.

Where it stops running on rails. Two places, and you named both.

- **Waiting on the client.** A missing bank statement, an unanswered question about a $4,200 charge, a receipt nobody uploaded. This is what pushes closes past the tenth.
- **The priority list.** Categorization sits in a queue longer than it should because you and Deshawn are not always working from the same one.

Hand re-entry points. Three that you named, and two more that Deshawn added when we asked him: new client details across PandaDoc, QuickBooks and the Client Tracker; client W-9s and vendor details typed into QuickBooks and then again into Gusto for the nine payroll clients; and the monthly Dext export matched by hand to QuickBooks on the four clients whose bank feeds lag.

The ten most-repeated manual workflows

#	Process	Trigger	Current tool	Hrs/wk	Maps to
1	Categorize transactions, harder files	daily feed	QuickBooks Online	12 (you)	Finding 3.2
2	Categorize transactions, easier files	daily feed	QuickBooks Online	15 (Deshawn)	Finding 3.2
3	Chase statements, receipts and answers	month end, then repeatedly	Gmail, by hand	6 to 8 (you)	Finding 3.1
4	First-round chase on Deshawn's files	month end	Gmail, by hand	2 (Deshawn)	Finding 3.1
5	Write 38 close summary emails	by the 10th	Gmail, sometimes ChatGPT	5 (you)	Finding 3.3
6	Prepare for check-in calls	before each call	reading, by hand	1 (you)	Finding 3.3
7	Run payroll for 9 clients	each pay date	Gusto	2 (both)	Section 4, inherent work
8	Onboard a new client into 3 systems	11 times a year	PandaDoc, QBO, Sheets	under 0.1	Finding 3.4
9	Match the Dext export to QuickBooks	monthly, 4 clients	by hand	not measured	Finding 3.4
10	Reconcile at month end	month end	QuickBooks Online	inside rows 1 and 2	Finding 3.2
	Total quantified			**43 hrs/wk**	

That 43 hours a week across 38 clients is about 1.13 hours per client per week, and it is the number Finding 3.6 is built on.

One thing we could not count, and are not pretending we did. Row 9, the Dext-to-QuickBooks matching Deshawn named, has no hours figure because nobody has ever timed it. It is not in any total in this report. Tell us what it runs to and we will recompute at no charge inside your question window.

3. The six findings, in payback order

Ranked by what you get back per hour of your own effort, not by size. The twenty-minute fix that returns $4,900 comes before the two-evening one that returns about the same.

A note that applies to all six. You told us your rule: nothing goes to a client without your review, because your name and your license are on the work. Every fix below is built that way. None of them sends a client anything you have not seen, and none of them gives any new company access to any client's books. That last point is not an accident; it is a constraint we set before writing a word, and two otherwise attractive ideas were cut by it. They are in Section 4.

3.1 Chasing clients for documents

Your most hated task, and the one you told us cannot be fixed.

What you told us. "Chasing late clients for bank statements and receipts every single month. It's the same conversation on repeat and it's the reason our closes slip past the 10th." And then, about the same task: "that one I don't think is fixable, it's just what clients are like."

You are partly right, and the finding has to concede it before it argues. A reminder does not make a client answer the question about the $4,200 charge. Nothing in this report makes a disorganized client organized. What a reminder system removes is not the client's slowness. **It removes you being the reminder.** If this finding wins, it wins on that distinction and on nothing else.

Current cost.

```
you       6 to 8 hr/wk x 50 weeks x $37 to $42/hr  =  $11,100 to $16,800
Deshawn   2 hr/wk      x 50 weeks x $29/hr         =  $ 2,900
                                             total  =  $14,000 to $19,700 a year
```

What is recoverable. 35% at the low end, 50% at the high end. **$4,900 to $9,850 a year**, and **2.8 to 5.0 hours a week**, of which 2.1 to 4.0 are yours.

Why not more than 50%. Two reasons, both from your own answers. First, you told us the split of what you chase: 40% bank and card statements, 35% receipts, and **25% questions about what a charge was for**. That last quarter is not a document request. It is a question only the client can answer and only you can ask well, and no system touches it. Second, this finding is Medium confidence, so 50% is its ceiling and we have not gone near it.

What does it today. You, from your own inbox, one client at a time, from memory of who owes what. Deshawn sends the first round on his own files and it comes back to you when it is the second or third ask.

What could. Three things you already pay for, in the order they cost you least.

1. **Dext, Request paperwork, on the other 11 of your 15 Dext clients.** It sends the client a notification in the Dext app naming the specific missing receipt, and it carries an in-app chat for document queries. You have it on for 4 of 15 because you were worried it would annoy people. You have now had it running on four clients for long enough to know whether it does. **Cost: $0. Time: about twenty minutes.** This is the first thing to do in this whole report.
2. **Dext, Fetch, on the clients whose feeds break.** Fetch connects to a supplier's portal and collects the bills itself, checking each connected supplier weekly. Statements are 40% of what you chase and you told us it is mostly clients whose bank feeds break or who have an account that is not connected. This is the category that stops being a conversation at all. **Cost: $0 on your tier. Time: one client at a time, about ten minutes each.**
3. **Financial Cents client tasks and requests, for the rest.** Your Team plan carries the client portal, client tasks and requests, and secure file sharing. That gives you one place where a request lives with a due date instead of a sent email you have to remember. **Cost: $0.**

Start with the nine. You told us 9 clients are late every single month, 12 occasionally, and 17 never. Do not build a system for 38. Switch it on for the 9, watch one close cycle, then decide about the 12.

One correction, and it changes what this costs you. You told us your Financial Cents Team plan "does have client requests and a client portal with automatic reminders". **Financial Cents' published plan table says the automatic part is not on Team.** Team carries the portal, the client tasks and the requests. The automated follow-ups are named "Auto-follow ups for client tasks" and they are the first feature of the **Scale** plan, at $69 a seat a month against Team's $49, both on annual billing (checked 2026-09-01). Your invoice of $98 a month for two seats confirms you are on Team annual.

So: **check your own plan page before you rely on the automated half.** It takes a minute and you have the admin view. If we are right, the arithmetic on the upgrade is:

```
Scale $69 - Team $49 = $20 a seat a month x 2 seats = $40 a month = $480 a year
```

We have not counted a single dollar of that upgrade in any total in this report. The $4,900 to $9,850 above is what the free path returns. The upgrade is a decision you make after you have run the free path for a month and know whether the manual send is the part that hurts. **The do-nothing option is real here**: the free path plus your existing habit of sending the reminders yourself is still better than today, because today the request lives in your head.

The data boundary. Everything above is an outbound request from a system you already own to a client who already uses it. No new company gets access to any client's books. Your engagement letter already names Dext, so nothing here needs a conversation with anyone.

Verdict: habit. The capability is bought and switched off. Sources: Dext Help Centre on Request paperwork and Fetch, and the Financial Cents published plan table, both checked 2026-09-01.

Confidence: medium (it depends on the 9 clients responding to an app notification better than they respond to you, which we cannot know until you try it on a cycle).

3.2 Categorization and reconciliation

The biggest block of hours in the firm, and the one place you have already tried the obvious fix and rejected it.

What you told us. Twelve hours a week for you on the harder files, fifteen of Deshawn's twenty-five on the easier ones. About 9,500 transactions a month across the book, roughly 250 per client, with the three messiest running 700 to 900 each. On one of those 900-a-month files, **about two thirds are vendors that were there last month.** The gas stations, the supply house, the phone bill, the same subcontractors. The other third needs a human.

And: about a dozen bank rules exist in total, across 5 or 6 of your 34 QuickBooks files, all set up years ago at onboarding. **Deshawn has never created one.**

Current cost.

```
you       12 hr/wk x 50 weeks x $37 to $42/hr  =  $22,200 to $25,200
Deshawn   15 hr/wk x 50 weeks x $29/hr         =  $21,750
                                         total  =  $43,950 to $46,950 a year
```

What is recoverable. 10% at the low end, 20% at the high end. **$4,395 to $9,390 a year**, and **2.7 to 5.4 hours a week**, of which 1.2 to 2.4 are yours.

Why so low, when the number underneath it is so large. This is the finding most likely to be overclaimed, so here is the ceiling we are allowed and the number we chose. Two thirds of transactions on a messy file recur, so two thirds is the absolute maximum a deterministic rule could ever touch. Our own house limit for a high-confidence labor finding is 80% of the addressable share, which would be 80% of 66.7%, or **53%**. At 53% this line would read about $23,400 to $25,000 a year. **We are claiming 10% to 20%, between a fifth and a third of what we are permitted.** The reason is that the two-thirds figure is your eyeball estimate on one file, and the entire finding would move if it turned out to be half or three quarters on the others. Correct that number and we recompute, free, inside your question window.

Now the objection, because it is the right one.

You turned on Intuit Assist's auto-categorization for one client as a test, it miscategorized, and you turned it off. When we asked what it actually got wrong you said: it put a contractor's Home Depot purchases in Office Supplies instead of job materials, and it coded an owner draw as a vendor payment, both without hesitation. The routine things, gas, phone, the same subcontractor every month, it got right. What scared you was that **it was confident about the judgment calls it had no business making.**

That is not an argument against what we are recommending. It is the argument for it.

A bank rule and an AI suggestion are different kinds of thing, and Intuit's own documentation says so. A rule is a condition you wrote and a category you chose. It tests the condition and applies the category, identically, every month, and it does not estimate. The AI suggestions in QuickBooks are built the other way: Intuit describes them as based on “your past work and transaction details” and on “info in the full bank description and your transaction history”, which is a description of inference from pattern.

The clearest evidence is a feature Intuit ships in the product. **Review Signals** are color-coded icons on the bank transactions page that show, in Intuit’s words, “the quality of data behind each QuickBooks suggestion”: whether it rests on a strong pattern in your books, on limited data “worth a second look”, or on very little history that “needs your attention before you post”. Intuit also states plainly that QuickBooks “may not have access to all the necessary information to categorize certain transactions accurately” and “relies on patterns and algorithms to make suggestions, but sometimes, additional context is needed.”

A bank rule has no confidence indicator, because it is not guessing. Your two failures were both “additional context” cases: knowing that this client is a contractor and therefore that Home Depot is job materials, and knowing that this particular transfer is the owner taking money out. You have that context. The suggestion engine does not. **You were right to turn it off, and you should leave it off.** The recommendation is the other tool.

What could.

1. **Write rules for the recurring two thirds on the three messiest files first.** Those three carry 700 to 900 transactions a month each. A QuickBooks company file holds up to 2,000 rules and a single rule can carry up to 5 conditions, so nothing about your volume is near a product limit. **Cost: $0.** Time: an evening for the first file, much less for the next two.
2. **Then stop writing them one file at a time.** You told us you have always done rules per file and had heard there is an export and import but never tried it. **There is.** In the file that has the rules: Transactions, then Banking, then Rules, then the arrow beside New rule, then **Export rules**, which saves a spreadsheet. In the target file: the same menu, then **Import rules**. Your clients are contractors, small retail and similar firms, so the vendor overlap across files is real: the same fuel, phone and supply-house vendors recur. **Cost: $0.** This is the difference between the fix being 34 jobs and being about 4.
3. **Hand the rules to Deshawn.** He has never made one, and he is carrying fifteen hours a week of the easier files, which is exactly where recurring vendors dominate. Half of this finding’s hours are his.

Where the real limit is, and we are not going to pretend otherwise. The other third is the part that needs a human, and after every rule you can write, it still needs one. This finding does not touch it, no tool in this report touches it, and the judgment on the harder files is the thing you are actually paid for. **The recoverable share stops where your judgment starts**, and that is why this is a 10% to 20% finding rather than a 53% one.

The data boundary. Bank rules live inside your own QuickBooks admin view, in files you already administer as the ProAdvisor of record. Nothing is shared with anyone, nothing new is connected, and no third party sees a transaction. Your engagement letter already names QuickBooks.

Verdict: setting. The capability is inside a subscription you already pay for and it is switched off in 28 of 34 files. Sources: Intuit help on bank rules, on importing existing bank rules, on AI suggestions and on the AI-powered banking page, all checked 2026-09-01. *Confidence: high on the labor figures, which come from your own hours and your own transaction counts.*

3.3 Thirty-eight monthly close summaries

A discrete, countable, monthly task, and the one place in this firm where drafting-for-review fits your rule exactly.

What you told us. About twenty hours a month across all thirty-eight, so five hours a week. Roughly twenty-five minutes each on average, but not evenly: **8 clients take about half the time**, the ones with two entities, a loan, or an owner who asks questions. The other 30 run 10 to 15 minutes each. Separately, about an hour a week of call prep, where you open the P&L and the summary before a check-in.

Current cost.

```
5 to 6 hr/wk x 50 weeks x $37 to $42/hr  =  $9,250 to $12,600 a year
```

What is recoverable. 30% at the low end, 50% at the high end. **$2,775 to $6,300 a year**, and **1.5 to 3.0 hours a week**, all of them yours.

Why we are holding this below its ceiling, and it is a concession rather than caution. You told us you already use ChatGPT to draft client emails. **Part of this gain is therefore already banked, and we are not going to sell it to you twice.** This finding is High confidence and its ceiling is 80%. We are at 30% to 50% because the drafting-from-a-blank-page saving is largely yours already.

What is not banked, and this is the actual finding. Two things.

- **Drafting from the numbers instead of from nothing.** Today you open the P&L, read it, and write prose. The version that saves time starts from a short export of the month's figures and the two or three movements worth mentioning, so the draft arrives with the content already in it and you are editing rather than composing.
- **Doing it the same way thirty-eight times.** Twenty-five minutes average across a range of ten minutes to an hour is a sign that each one is being approached fresh. A fixed shape for the 30 simpler clients is most of the win. **The 8 harder ones should stay handwritten**, because those are the two-entity, loan-carrying, question-asking clients and that is the paragraph they are paying you for.

And it doubles as your call prep. The hour a week you spend reading before a check-in is reading the same two artifacts. If the summary already leads with the two or three things that moved, the prep is done when the summary is.

What could, and where it should run. You told us you are on Google Workspace Business Starter, that Gemini shows up in Gmail now, and that you have not read the terms. We read them, and this matters more than the hours do.

Google's position on Workspace data. Google states that interactions with Gemini in Workspace "stay within your organization", that your content "is not shared with or used by any other customers", and that submissions are "not used to train models and never reviewed by humans" under a Workspace with Gemini license. Your existing Workspace protections apply to it automatically. Gemini is included in all paid Workspace plans at no extra charge; the separate add-on was discontinued and base prices absorbed it. On Business Starter that means Gemini in the Gmail side panel, at a lower capability tier than the more expensive plans, for **$0 more than you pay today**.

And the honest version of the comparison, which is narrower than it sounds. For personal ChatGPT accounts on the Free, Plus and Pro plans, conversations are used to improve OpenAI's models **by default**, and that default can be turned off at Settings, then Data Controls, then "Improve the model for everyone". So the difference is not that one company uses your text and the other does not. **The difference is that the Workspace protection is a term you have already accepted, applied automatically, with no per-person setting anyone can forget.** The consumer protection depends on a toggle that starts in the wrong position and that each person has to find. In a two-person firm where one of you is on the account and one is not, that distinction is the whole thing.

Cost: $0. It is the mail system you already run the firm on.

The data boundary. Draft from a sanitized export: the month's figures and movements, with the client's name and account identifiers left out, which the draft does not need in order to be useful. Nothing goes to a client without you reading it, which is your rule and not our concession. And Google Workspace is not a new provider on your engagement letter's list, so nothing here triggers the conversation you told us you would want to have first.

Verdict: habit. Nothing external is missing. The capability and the terms are already yours. Sources: Google's Generative AI in Google Workspace Privacy Hub and its Workspace with Gemini business FAQ, and OpenAI's Data Controls FAQ, all checked 2026-09-01. *Confidence: high on the labor figures, which are your own hours.*

3.4 New client details typed into three systems

You said it is fifteen minutes and you have never treated it as a priority. You were right.

What you told us. New client information goes into PandaDoc for the engagement letter, then into QuickBooks to open the company file, then into your Client Tracker sheet. "It's maybe fifteen minutes per client so I've never treated it as a priority." And you signed 11 new clients in the last twelve months.

Current cost.

```
11 new clients x 15 minutes = 2.75 hours a year x $37/hr = $101.75 a year
```

What is recoverable: nothing worth your attention. We are counting this at $0 and recommending you leave it exactly as it is. A hundred dollars a year is under three hours of your time annually. Any integration that removed it would take longer to set up than it will ever give back, and it would add a moving part to an onboarding that currently works. Your instinct on this was correct and it is the only finding in this report where the recommendation is to do nothing.

One thing in the same category that might not be small, and we do not have the number. When we asked Deshawn what he would add to your retyping list, he named two things you did not: client W-9s and vendor details typed into QuickBooks and then again into Gusto for the payroll clients, and **the**

monthly Dext export matched by hand to QuickBooks on the four clients whose bank feeds lag. That second one is monthly rather than eleven times a year, and it is a matching job rather than a typing job, which usually means it is bigger than it looks.

Nobody has timed it, so it is not in any total in this report. Time it for one month and tell us during your question window and we will recompute at no charge. If it runs to even an hour a month across the four clients it outranks this finding by a factor of ten, and the fix is likely to be inside the same bank-rule work as Finding 3.2.

The data boundary. Your own onboarding, in your own systems. Nothing external.

Verdict: habit, and a cheap one. Leave it. *Confidence: high. The arithmetic is small and simple and it does not depend on anything we could be wrong about.*

3.5 What you actually pay for software

You said a couple hundred a month.

What your invoices say. We asked for them with plan names, and the plan names turned out to matter more than the amounts.

Tool	Plan	Monthly	Whose cost
Dext	partner tier, up to 25 clients	$255.00	**firm**
Financial Cents	Team, 2 seats	$98.00	**firm**
PandaDoc	Essentials, 1 seat, annual	$35.00	**firm**
Google Workspace	Business Starter, 2 users	$16.80	**firm**
Calendly	Standard	$12.00	**firm**
HubSpot	free tier	$0.00	firm
Mailchimp	free plan	$0.00	firm
Gusto	accountant program, billed to clients	$0.00	clients
Firm total		**$416.80**	
QuickBooks Online	34 clients, ProAdvisor wholesale	$1,180.00	**rebilled to clients**
Xero	4 clients, partner pricing	$130.00	**rebilled to clients**

The split is the finding. $1,310 a month of that is client subscriptions you rebill inside each retainer. It moves through the firm and out again and it is not your cost. **What you actually pay to run the firm is $416.80 a month, or $5,002 a year.** You had it in your head as a couple hundred.

That is roughly double, and it is your discovery rather than ours: we only added up the invoices you already had. It is worth knowing for one specific reason. You told us your comfort for new tools is $100 to

$250 a month total, and that ceiling is a lot easier to reason about once you know the base it sits on top of. The only new spend anywhere in this report is the optional $40 a month in Finding 3.1, which fits inside $250 with room.

There is no saving in this finding and we are not going to manufacture one. We counted it at **$0**. HubSpot and Mailchimp are both free, so cancelling them saves nothing and is a two-minute tidy-up rather than a finding. Here is everything we did find, with the arithmetic, none of it in any total:

- **You are paying for ten Dext client slots that sit empty.** Your tier covers 25 clients and you have 15 in it. At $255 that is $10.20 per slot across what you buy, or $17.00 per client across what you use. **This is not a saving, it is an asset**: the ten empty slots are why extending Dext to more clients in Finding 3.1 costs you nothing. Dext does not publish partner pricing, so we cannot tell you what a smaller tier would cost and we are not going to guess.
- **PandaDoc, $35 a month for one Essentials seat on annual billing.** The published Essentials annual price is around $19 a seat (checked 2026-09-01). We cannot see your account, so your invoice is what governs here and the gap is a question rather than a finding: it may be a legacy plan, an add-on, or a plan that was renamed underneath you. **Ask them what the $35 is for.** If it comes back as $19 that is $192 a year, which we have not counted because we do not know that it will.
- **Calendly and Google Workspace are both on month-to-month billing.** Calendly Standard is $12 monthly and $10 annually. Workspace Business Starter is $8.40 per user on flexible billing and $7 annually, and you have two users. Switching both to annual is $24 plus $33.60, so **$57.60 a year**. We have not counted it either, because it buys a saving with a twelve-month commitment and that is a trade rather than a free win.
- **Financial Cents, paid and half used.** You called this out yourself. It is not a saving at all, it is the delivery mechanism for Finding 3.1 and for the priority-list problem you named in your intake. The money is already spent. The question is whether the plan does what you think it does, which is in Finding 3.1.

The data boundary. Your own invoices. Nothing external.

Verdict: setting. The Dext capacity is bought and empty and the Financial Cents capability is bought and unused. Sources: your invoices dated 2026-09-01, plus PandaDoc, Calendly and Google Workspace published pricing checked the same day. *Confidence: medium, because two of the four items above end in a question to a vendor rather than an answer from one.*

3.6 The three or four more clients

The question you actually asked, answered in hours rather than in dollars.

What you told us. Given ten hours a week back: some of it for yourself, to stop working Saturday mornings during tax season, and the rest toward three or four more clients without hiring. A new client of the size you would want is $450 to $550 a month. Your book is 38 monthly clients at about $450 average, and 30 files get touched in a typical week.

We are counting this at $0. It is not a separate saving, it is what the first three findings buy if you spend them this way. Counting it again would be counting the same hours twice, and putting the biggest number in the audit on an assumption we cannot check would be the least honest thing in this document. So here is the arithmetic instead.

```
your current delivery load, from your own figures:
   categorization and reconciliation   27 hr/wk  (12 yours + 15 Deshawn’s)
   chasing                              8 hr/wk  (6 yours + 2 Deshawn’s)
   close summaries                      5 hr/wk
   call prep                            1 hr/wk
   payroll on Gusto                     2 hr/wk
                              total    43 hr/wk  across 38 clients
                                             =  1.13 hr/wk per client
   your share of that load: 24.7 of 43   =  57%

four more clients   4 x 1.13   =  4.5 hr/wk of firm time, of which about 2.6 is yours
what 3.1 + 3.2 + 3.3 return    =  7.0 to 13.4 hr/wk, of which 4.8 to 9.4 is yours
your Saturdays      12 x 4 hr  =  48 hr/yr, about 1.0 hr/wk averaged across the year
```

The answer, and it has two halves.

The half that works. At the low end you get back 4.8 hours a week of your own time, which is 240 hours a year. Four more clients need about 130 hours of your year and the Saturdays need 48. That is 178 against 240. **On the annual average, the low end pays for both with about sixty hours to spare**, and the high end is not close. Deshawn’s side works too, but only just: four clients need about 1.9 hours a week of his and the low end returns 2.2.

The half that does not. Those 48 Saturday hours are not spread across the year. They are four hours on each of twelve consecutive Saturdays between mid-January and April 15, and those are precisely the weeks when your volume peaks, when Deshawn goes to 30 hours, and when the recovered time is least available because everything else is at its heaviest. **An annual average is the wrong unit for a seasonal problem.** The hours are found across 50 weeks and the Saturdays need to be paid for inside 12 of them.

So the choice is real. If the four clients arrive in the autumn, they are consuming the same capacity you were going to spend on tax season, and February decides which one actually happened. Our suggestion is to sequence it: take the Saturdays back first, in the tax season after you do the work, and add clients from the spring once you have seen what the hours actually came to.

And the constraint that is not about total hours at all. Between a quarter and a third of everything found here is Deshawn’s time, not yours. New clients need your review, because your name and license are on the work and you told us that is not negotiable. So the binding limit on growth is your hours, not the firm’s, and Finding 3.3 matters more to this question than its dollar figure suggests: it is the one that returns time exclusively to you.

The revenue figure, stated and not claimed. Four clients at $450 a month is **$21,600 a year**. It is in this report once, here, in this sentence, and it is in no total anywhere. It requires you to go and sell four clients, which is a thing this audit does not do for you and which your close rate of about one in three on three proposals a month says would take a while.

The data boundary. Arithmetic on your own figures. Nothing external.

Verdict: not applicable. This is a consequence of the three findings above rather than a finding of its own. *Confidence: medium, and it rests entirely on the hours figures you gave us.*

4. What we deliberately did not pursue

Everything we looked at and left, so you can see it was decided rather than missed.

- **Saturday work in tax season as its own finding.** It is the consequence of Findings 3.1, 3.2 and 3.3, and counting it separately would count the same hours twice. It appears in 3.6 as the thing the hours buy back, which is how you framed it.
- **Closes slipping past the tenth as its own finding.** Same reason. You told us yourself it is caused by the chasing, so it is Finding 3.1's consequence and it is already counted there.
- **Moving off QuickBooks Online.** You are a Certified ProAdvisor with 34 files on it. Migration would cost more than every finding in this report combined and nothing in your intake asks for it.
- **Migrating the four Xero clients.** A real annoyance, small money, and a client-relationship decision rather than an operations one. If they leave, let it happen at renewal.
- **Any AI tool with access to client books.** Your read-only answer was "maybe, tell me more", and the more turned out to be that the two views worth having are your own QuickBooks Accountant view and your own Financial Cents view. **We ruled out anything beyond that at the outline stage, not in the writing.** Two ideas died there, and both would have looked good in this report.
- **The Client Tracker sheet as a single point of failure.** You flagged it yourself: "If that sheet disappeared I'd genuinely be stuck for a week." That is a real risk and it is worth one sentence, which is this one. The honest fix is inside Financial Cents, which you already pay for, rather than a new tool, and it belongs after the thirty days in Section 6 rather than inside them.
- **HubSpot and Mailchimp as growth levers.** Finding 3.6 is about capacity to serve more clients, not about generating them. A marketing recommendation here would be us changing the subject.
- **Anything above $250 a month.** Your ceiling, and we treated it as hard rather than as a starting position.
- **Answering routine "what does this mean" client emails**, which is the third item on Deshawn's list. It is real and we are not making it a finding, because it is the same quarter of the chase that Finding 3.1 already concedes it cannot automate: a client asking what a charge was for, or what a line on their P&L means, is a question that needs one of you. The part of it that is answerable in advance is the close summary, which is Finding 3.3, and a better summary is the only lever on this we would trust.
- **The payroll work on Gusto.** Nine clients, about two hours a week between you and Deshawn, and it costs the firm nothing because Gusto bills the clients directly. It is two hours of genuinely inherent work in the two days before each pay date. We looked and found nothing to remove.

Your three most-hated tasks, and where each one went. You named transaction categorization on the harder files (Finding 3.2), chasing clients (Finding 3.1), and writing close summaries (Finding 3.3). All three are findings with numbers. Nothing you named was quietly dropped.

5. What this is worth, and the bar we hold ourselves to

Total identified opportunity: 7.0 to 13.4 hours a week, and $12,070 to $25,540 a year.

Finding	Recoverable $/yr	Hours/wk	Confidence
3.1 Chasing clients for documents	$4,900 to $9,850	2.8 to 5.0	Medium
3.2 Categorization and reconciliation	$4,395 to $9,390	2.7 to 5.4	**High**
3.3 Thirty-eight close summaries	$2,775 to $6,300	1.5 to 3.0	**High**
3.4 New client details typed three times	$0 (sized at $102)	0	High
3.5 What you actually pay for software	$0	0	Medium
3.6 Three or four more clients	$0 ($21,600 shown, not claimed)	0	Medium
Total	**$12,070 to $25,540**	**7.0 to 13.4**	

New tool spend: $0 a year. Every counted finding runs on tools you already pay for. The one optional purchase in this report, the Financial Cents upgrade at $480 a year, is excluded from every figure above.

Net: $12,070 to $25,540 a year.

Even counting only what we're certain of: $7,170/yr. That is Findings 3.2 and 3.3 alone, at the bottom of their ranges, which are the two we would stand behind with no benefit of the doubt because both rest on your own hours and neither depends on a client changing their behavior.

The bar we hold ourselves to. We don't send an audit that finds less than $4,500 a year against a $1,500 fee. This one finds $12,070 at the bottom of its range, and $7,170 counting only the high-confidence findings. That is a standard for our own work, checked before you saw this, and it is the reason three of the six findings above are counted at zero rather than padded into the total.

What would move these numbers, in the order it would move them. Every figure here rests on hours you gave us in a follow-up conversation rather than on a system that measured them, and we would rather say that plainly than let the precision of the arithmetic imply a precision the inputs do not have.

1. **The recurring-vendor share.** We used two thirds, from your eyeball estimate on one file. It is the single number that moves this audit most.
2. **Your hours on categorization.** Twelve a week is the largest input in the report.
3. **The Dext-to-QuickBooks matching nobody has timed.** It could be larger than Finding 3.4.
4. **Your Financial Cents plan's actual reminder capability**, which decides whether Finding 3.1 costs nothing or $40 a month.

Correct any of them during your question window and we recompute at no charge.

6. The first 30 days, in order

Free and fast first. Nothing in weeks 1 and 2 costs a dollar, and the one decision that costs anything waits until week 4, by which time you will have the evidence to make it.

Week	Do	Why this week	Who	Time and cost	Return (annualized)
1	Switch Dext Request paperwork on for the other 11 of your 15 Dext clients, starting with any of the 9 always-late ones already in Dext	Cheapest move in the report and it starts working on this month’s close	You	20 min, $0	part of $4,900 to $9,850
1	Open your Financial Cents billing page and confirm whether Team carries automated follow-ups	It decides the week-4 decision and it takes a minute	You	5 min, $0	decides $480/yr
1	Write bank rules for the recurring vendors on your single messiest file (the 900-a-month one)	Biggest concentration of recurring vendors in the book	You	2 hrs, $0	part of $4,395 to $9,390
2	Export those rules and import them into the other two messy files, then into 3 or 4 similar clients	Proves the export path works before you invest in it	You	1 hr, $0	part of $4,395 to $9,390
2	Set up Dext Fetch for the clients whose bank feeds break	Statements are 40% of the chase and this is the part that stops being a conversation	You	10 min each, $0	part of $4,900 to $9,850
2	Sit down with Deshawn over writing and importing a rule, then hand him the easier files	Half of Finding 3.2’s hours are his and he has never made one	Both	1 hr, $0	part of $4,395 to $9,390
3	Draft one month of close summaries from an export in the Gemini side panel in Gmail, for 5 of the 30 simpler clients	Small enough to abandon, big enough to judge	You	1 hr, $0	part of $2,775 to $6,300
3	Move the 9 always-late clients onto Financial Cents client tasks with due dates	Requests stop living in your head	You	1 hr, $0	part of $4,900 to $9,850
3	Find out from PandaDoc what the $35 is for	It is a two-line email and it may be worth $192/yr	You	5 min, $0	not counted
4	Run one full close cycle on the new setup and count what actually changed	The only way to know which end of our ranges you are at	Both	no extra time, $0	tells you where in the ranges you land

Week	Do	Why this week	Who	Time and cost	Return (annualized)
4	Then decide on the Financial Cents Scale upgrade	You now have a month of evidence about whether manual sending is the part that hurts	You	15 min, then $40/mo if yes	not counted

Total cost for the thirty days: $0. Total return at the low end of the three findings the plan touches: **$12,070 a year**. The one paid decision sits at the end on purpose.

If you only do three things. Switch Dext Request paperwork on for the other eleven clients. Write the rules on the messiest file and import them into the other two. Check your Financial Cents plan page. That is under four hours, it costs nothing, and it is most of the low end.

Month two, once the wins are funding it. Give the Client Tracker sheet a second home inside Financial Cents, where the due dates and per-client notes already have somewhere to live. You told us you would be stuck for a week without it, and it is the one risk in this report that is not about hours.

7. Appendix: our assumptions

Rate assumptions. Two loaded rates, both sourced and dated, both disclosed against your own figures with the lower one governing.

Role	Rate used	How we got there
You	**$37 to $42/hr**	Your own figure, burdened and rounded down
Deshawn	**$29/hr**	BLS OEWS state survey, burdened and rounded down

Both rates in full. Your **$37 to $42** is your own answer: a senior bookkeeper in Tulsa you would trust with the harder files is $28 to $32 an hour W-2, which burdens by 1.3365 to $37.42 to $42.77, rounded down at both ends. Deshawn's **$29** is BLS Occupational Employment and Wage Statistics, Oklahoma, Bookkeeping, Accounting and Auditing Clerks, mean $22.31 an hour (May 2024 survey); burdened by 1.3365 that is $29.82, rounded down.

Where the burden multiplier comes from. BLS Employer Costs for Employee Compensation, September 2024 reference period, private industry, establishments with 1 to 49 workers: total compensation $35.27 an hour against wages and salaries of $26.39. That ratio is 1.3365. We used the small-establishment cut rather than the all-industry one because your firm is two people, and because it is the lower of the two available.

Three disclosures about those rates, because they all push the numbers down.

1. **A loaded rate normally means wage plus taxes plus benefits plus overhead. The federal series we used contains no overhead component at all.** It covers wages, employer payroll taxes and employer-paid benefits and stops. So both rates above sit below a true loaded cost, and every dollar figure in this report is smaller than it would be on a fuller definition.
2. **Deshawn's own rate is higher than the one we used.** He is W-2 at $24 an hour, which burdens to $32.08. The sourced replacement figure is $29.82. We used the lower one, which under-claims his hour by $3.08 against what you actually pay for it.
3. **We cross-checked your rate a second way and it agreed.** A 50/50 blend of the bookkeeping-clerk series and the accountant-and-auditor series for Oklahoma, both burdened, lands at $42.11, against the $42.77 top of your own band. Two independent methods within seventy cents is the best corroboration available here. We used the lower of them.

Your unpaid Saturdays are counted at the same rate as your paid hours, and that is a judgment call worth naming. You do not invoice a Saturday morning, so counting it at replacement cost prices unpaid owner overtime as though it were payroll. We counted it anyway, because the alternative is to value your own hours at less than everyone else's, and that is the exact reasoning that produces Saturday mornings in the first place.

Working-week and volume assumptions. Fifty working weeks a year, which allows two weeks off and is the conservative choice: at fifty-two every annual figure here would be 4% larger.

Which volumes came from where. Everything marked "you told us" came from your intake form or your answers to our follow-up questions, and is your figure rather than ours. The transaction counts, the bank rule count and the Dext client counts came out of your own QuickBooks Accountant view and Financial Cents admin view under the read-only access you granted. **The single most important number in this report, the two-thirds recurring-vendor share, is your eyeball estimate on one file**, and it is labeled that way in Finding 3.2 rather than presented as measured.

What we did not measure and did not pretend to. Deshawn's monthly Dext-to-QuickBooks matching on the four lagging clients. It is named in Finding 3.4 and it is in no total in this report.

Methodology. Findings were built from your intake form, a follow-up question set, and a research pass over the tools you named. Every finding carries a number and shows its calculation. Every capability and price claim was checked against the vendor's own current documentation on 2026-09-01 and is dated in the finding that uses it. Where a vendor's published plan table disagreed with what you told us, both are in the report and the disagreement is named rather than resolved silently. Where a vendor does not publish its pricing at all, which is the case for Dext's partner tier, your invoice governs and we say we cannot compare it.

Confidence, and what it means here. High means the arithmetic rests on your own hours or your own system counts and does not depend on anyone changing behavior. Medium means something in it depends on a person or a vendor answer we do not have. Recoverable percentages are capped at 80% for high-confidence labor findings and 50% for medium-confidence ones. **No finding in this report sits at its cap and three are counted at zero.**

Delivery timeline. Day 0, intake received and parsed. Day 1, findings outline and the modeling decisions behind it, reviewed before any prose was written. Day 2, follow-up questions to you. Day 3, research against vendor documentation, every claim dated. Day 4, the numbers worksheet, where every figure in this report was computed and checked. Days 5 and 6, drafting and quality checks. Day 7, review and delivery.

Data handling. This report was produced for a fictional business created to demonstrate the work, so no client data of any kind was accessed. In a live audit, the read-only access you grant is used for counting only, it is never used to open a client file, and it is revoked on delivery. Working papers are deleted within thirty days and you get a confirmation when they are.

How to get help. Your questions are open by email for two weeks from the day this arrived. Reply to the delivery email and we answer. We check in on day ten in case something has come up that you have not asked about. And a live 30-minute readout call is included with your audit: book it with the link in the delivery email and we walk the findings together.

Bonus A: Getting named when a customer asks an AI assistant

More of your future clients will find you by asking an assistant than by typing into a search box, and the shape of that question is different. Nobody types “bookkeeper Tulsa” into an assistant. They type “I run a small construction company in Tulsa and I need someone to clean up two years of QuickBooks, who should I call”.

What we checked, and what we did not. We did not run assistant queries about your business, because **your business does not exist**: this is a fictional sample, and any assistant asked about Bramblewood Bookkeeping and Advisory would either say nothing or invent something. That is the honest expected result and reporting it is more useful than a fabricated screenshot. What follows is the input checklist, which is the part that is the same for every firm and the part you can act on. In a live audit this section carries a dated snapshot of what named assistants actually said, with the date and the number of questions stated on the page.

Who you are actually up against. Not the national franchises. When an assistant answers a question like the one above it tends to name two or three local firms with a real web presence, plus whichever directories it can read. In a market like Tulsa that is a short list, and being on it is mostly a matter of being legible rather than being loud.

Can the assistants even read your site? This is the part most firms fail and it is the cheapest to fix.

Check	What it means	How to test it
robots.txt	Whether the crawlers named GPTBot, ClaudeBot and PerplexityBot may read your pages at all	Open yoursite.com/robots.txt and read it. If those names appear beside Disallow, you are invisible on purpose
Structured data	Whether a machine can tell you are a bookkeeping firm in Tulsa rather than guessing from prose	Google’s Rich Results Test. You want LocalBusiness or Organization markup
Answer-first pages	Whether a page answers a question in its first sentence, or winds up to it	Read your own first paragraph. If it starts “In today’s business environment”, an assistant has nothing to quote
Name, address, phone	Whether the same three facts appear identically on your site, your Google Business Profile and the directories	Search your firm name and compare. Assistants downweight what they cannot corroborate
Freshness	Whether anything has changed on the site in a year	Your own judgment
Google Business Profile	Complete, categorized, with current hours and services	Sign in and read it as a stranger would

Check	What it means	How to test it
Review breadth	Whether reviews exist anywhere other than Google	Search your firm name plus “reviews”

What to do about it. In order of return.

1. **Check robots.txt.** Five minutes. If it blocks assistant crawlers, unblocking is a one-line change and it is the difference between being unfindable and being findable.
2. **Add LocalBusiness structured data** naming the firm, the city, the services and the credential. You are a Certified ProAdvisor and that is exactly the kind of fact an assistant will repeat if it can read it and will omit if it cannot. Half an hour with a generator, or ask whoever built the site.
3. **Write three answer-first pages** for the three questions your best clients actually arrived with. QuickBooks cleanup, monthly bookkeeping for contractors, and getting your numbers ready for a loan application are the three your intake describes. Lead each with the answer in the first two sentences.
4. **Make the name, address and phone identical everywhere**, including the Google Business Profile.
5. **Get reviews somewhere other than Google.** Breadth of source matters more than volume on any one.

What this does and does not tell you. Assistant answers vary between runs and between people, so a single check is a snapshot rather than a measurement, and anyone who promises you placement in an assistant’s answer is selling something nobody can deliver. What the checklist above does is make you legible to the systems that are already reading. That is ordinary work with an ordinary payoff, and it is the honest version of this.

Bonus B: A starter AI policy

A starting point, not a legal document. Two people means you do not need much, but you do need the part that says what never goes into a chat box. Have your attorney read it before it becomes official, especially the client-data clause, since your engagement letters carry a confidentiality term this has to sit underneath.

AI Usage Policy, Bramblewood Bookkeeping and Advisory, LLC

1. Approved tools

Tool	Approved for	Not approved for
Gemini in Google Workspace	Drafting client emails and close summaries from sanitized figures, summarizing your own notes	Anything containing a client's name, account numbers or identifying details
QuickBooks Online bank rules	Categorizing recurring transactions you have defined	Judgment calls: owner draws, job materials, anything needing context
QuickBooks AI suggestions	Nothing at present, by owner decision after testing	Auto-applying categories to client books
ChatGPT or any personal AI account	General questions with no client information in them	Any client work, any client data, any document
Dext, Financial Cents, Gusto	As sold, under their own terms, all named in the engagement letter	Adding a new integration without owner approval

2. What never goes into an AI tool that is not on the approved list

Client names. Employer identification numbers or Social Security numbers. Bank account or card numbers. Login details for anything. Payroll figures tied to a named person. Anything from a client that arrived marked confidential. If you are unsure, the answer is no and you ask Renata.

3. Client data

Client information may only be processed by providers named in the current engagement letter. Today that is QuickBooks, Dext and Gusto, and adding a provider means updating the list at renewal and telling the client, and nobody adds one without the owner's approval.

4. Acceptable uses

Drafting anything a person will read and edit before it goes out. Summarizing documents the firm already holds. Explaining an accounting concept in plainer words. Checking your own reasoning.

5. Prohibited uses

Sending anything to a client that a person has not read first. Anything that reads as tax or legal advice. Letting a tool categorize transactions on a client file without review. Uploading a client document to any tool not on the approved list.

6. The review rule, which is the whole policy in one line

Final review before anything goes to a client. Always. The firm’s name and license are on the work, and no tool changes that.

7. If something goes wrong

Say so to Renata the same day. Not the next week and not after you have tried to fix it. Say what went in, to which tool, and when. Nobody is in trouble for reporting it; the only mistake that gets worse with time is the one nobody mentioned.

8. Review schedule

Reread this every six months and whenever a tool is added or a plan changes. Tool capabilities move faster than policies do, and a policy naming a feature that no longer works the way it is described is worse than no policy.

End of sample audit.