McDole Group

TECHNOLOGY, HANDLED

DISCOVERY AUDIT · OPPORTUNITY REPORT

AI Opportunity Report

PREPARED FOR

Acalanes Insurance Partners

Lafayette, CA · Independent P&C agency · Second-generation family-owned, established 1987

PREPARED BY
Lee McDole · McDole Group
lee@mcdolegroup.com · (510) 667-2020
mcdolegroup.com

PERIOD
One week on-site, May 18 – 22, 2026
Synthesized and drafted May 25, 2026

ABOUT THIS SAMPLE

Acalanes Insurance Partners is a fictional independent insurance agency. All names, figures, scenarios, and recommendations are illustrative, modeled on the kinds of operations, software, and pain points common to East Bay independent P&C agencies of 8 to 25 staff. No real agency, employee, carrier relationship, or revenue figure is represented.

Contents

1. Executive summary

Acalanes Insurance Partners is a twelve-person family-owned agency with roughly 1,400 active personal- and small-commercial policies on HawkSoft, generating around $3.2 million in annual revenue. The owner-producer, his two producers, four CSRs, an operations manager, and four administrative staff run the operation. They're leaving operational capacity on the table that the team itself can describe without prompting.

The morning carrier-email triage is the loudest example. Across four CSRs, the agency spends roughly 50 hours per week filing carrier downloads into HawkSoft and notifying clients of changes. That's a full-time employee's worth of repeatable, rules-bound work being done by hand, every week, every CSR, before any of them touch a client conversation.

Four agent patterns, tailored to this agency's stack and workflows, would recover an estimated **$152,000 per year** in CSR and producer time plus modest revenue recovery from lost leads and slipped renewals. A fifth pattern, a voice receptionist, surfaced during the week but isn't worth building at the agency's current call volume. The honest call there is *not yet*.

The five opportunities, ranked

#	OPPORTUNITY	HEADLINE VALUE	RECOMMENDATION
1	Carrier document parser	~40 hrs/week recovered · ~$73K/year	**Build first**
2	After-hours lead response	~$25K/year recovered revenue	Build month 2
3	Renewal & recall outreach	~$25K/year revenue protection	Build month 3
4	Internal carrier knowledge RAG	~12 hrs/week recovered · ~$29K/year	Build months 4 – 5
5	Voice receptionist	Real pain, insufficient ROI at current volume	**Not yet**

Build order, and one to skip

Start with the carrier document parser in month one. The pain is universal across the CSR team, the ROI dwarfs every other finding, and getting it right early demonstrates the managed-agent model before scope expands. Opportunities 2 through 4 phase across months two through five at a pace your team can absorb without disruption.

Skip the voice receptionist for now. Call volume is currently around 80 inbound calls per week, handled comfortably by your operations manager. The install cost for a quality voice agent ($10,000 to $16,000 all-in) is real, and the recoverable value at this volume doesn't justify going ahead of higher-leverage work. Revisit if call volume crosses ~150 per week or your operations manager's role changes.

Several other things you might have expected to see also don't make the cut: a migration off HawkSoft, predictive analytics on your book of business, or a custom client portal. Each is covered in **Section 6, What I'd not do, yet**, with the reasoning.

The full pricing for the four-agent install is detailed in **Section 5**. Briefly, the install lands in the range of **$16,500 – $22,000** before the $2,500 you paid for this Audit is credited ($14,000 – $19,500 after), and the monthly retainer lands in the range of **$2,500 – $3,200**.

---

2. What I observed

2.1 The agency

Acalanes Insurance Partners was founded in 1987 by Dale Acalanes, a former CSAA underwriter who left the carrier side to build his own book. Karen Acalanes-Reyes, Dale's daughter and the agency's principal since 2019, took over operations after a decade of producing under Dale, and has spent the last six years modernizing the back office in deliberate, conservative steps. The agency runs on HawkSoft, with EZLynx Rating Engine for comparative quoting, Microsoft 365 for email, RingCentral for phones, and SharePoint for the document overflow that doesn't live in HawkSoft attachments. The office occupies a 3,200-square-foot storefront on Mt. Diablo Boulevard in Lafayette, open 8:30am to 5:30pm Monday through Friday.

The book is approximately 1,400 active policies, weighted ~75% personal lines (homeowners, auto, umbrella) and ~25% small commercial (BOPs for local retail, light contractors, and professional liability for the East Bay's small-firm legal and dental clientele). Total annual revenue is approximately $3.2 million. Twelve people run the operation:

- **Karen Acalanes-Reyes**, owner and lead producer
- **Jim**, producer, personal-lines lead
- **Lina**, producer, commercial book (hired Fall 2025)
- **Maria, Sarah, Brian, and Olivia**, CSRs, the backbone of the daily operation
- **Carl**, operations manager, who also fields most of the inbound calls
- **Tasha**, new-business intake admin
- **3 additional administrative staff**, split across billing, claims liaison, and reception

The carrier roster is the usual East Bay independent mix: Travelers, Mercury, Safeco, CSAA, Stillwater, Foremost, Kemper, and a handful of E&S markets. The California homeowners market has consumed an outsized share of Karen's attention for the last three years, though. With Allstate, State Farm, and others pulling back from California wildfire exposure, the agency has rebuilt large parts of its homeowners book around Mercury, CSAA, Stillwater, and the harder-to-place specialty markets. The shift added a real per-policy paperwork load, since carriers with appetite for California risk tend to have stricter documentation requirements, more frequent re-underwriting, and faster non-renewal cycles.

2.2 The week on-site

I was at the agency from Monday May 18 through Friday May 22. Most days I worked from a desk Karen set up for me in the back corner of the open floor, near the CSR pod. I rotated between observation and interviews.

The interviews:

- **Karen**, two long sessions: Monday morning to set context, Friday afternoon to walk through preliminary findings.
- **Carl** (operations manager), one structured sit-down on Tuesday, plus questions throughout the week as I observed call handling and renewal queue management.
- **All four CSRs**, 30 to 45 minutes each, scheduled around their workload. The conversations were specifically structured around "walk me through your morning" rather than abstract questions about what's hard.
- **Both producers** (Jim and Lina), 30 minutes each.
- **One admin** (Tasha), who handles new-business intake, including most after-hours forms and weekend voicemail recovery.
- **Karen's HawkSoft account rep** by phone, to confirm what integrations are currently active and what's available without a contract upgrade.

The observation:

- **Tuesday and Thursday mornings, 8:30am to 11:00am**, sitting at Maria's desk during the morning carrier-email triage. This is where the bulk of the document parser recommendation comes from: direct observation of the work, timed with Maria's permission.
- **Wednesday afternoon**, sitting in on a renewal call Jim conducted with a long-term homeowner client whose Mercury policy was up for renewal and whose premium had increased 22% year-over-year. Useful for understanding the renewal conversation pattern, the supporting documentation pull, and what an avoidable customer-loss event looks like in real time.
- **Throughout the week**, monitoring the agency's general inbox and after-hours form submissions to count and classify the volume.

The findings preview on Friday afternoon was Karen, Carl, and Lina (by Karen's invitation). I shared the ranking that became this report's executive summary, with the express purpose of pressure-testing the headline numbers before writing them down. They corrected one estimate (call volume, since my Tuesday count was unusually low) and confirmed the rest.

2.3 The shape of the work

Four work patterns dominate the agency's day. Each one is a category of recurring, mostly rules-bound work currently done by a person, that could mostly be done by a managed AI agent. The four become the four "build" recommendations in Section 3. The fifth opportunity in this report (voice reception) sits outside these four patterns and is addressed separately.

The morning carrier-email triage. Each CSR begins the day with between 40 and 80 carrier emails: policy declarations pages, endorsements, renewal offers, non-renewal notices, billing changes, premium adjustments, and routine system notifications. The CSR opens each one, identifies the document type, locates the correct policy in HawkSoft, attaches the document to the policy record, updates any fields that have changed (premium, expiration date, coverage limits), and where the change is material to the client, drafts an email to the client explaining what happened. Across four CSRs this consumes roughly 50 hours per week in aggregate. The work is rules-bound. It also has to be done before any client-facing work can start, which is why mornings at the agency are inbox mornings: the CSR pod is silent until 10:30 or 11:00.

The renewal pipeline cadence. Personal-lines renewals are managed in HawkSoft's renewal report, which surfaces policies expiring in the next 60 days. A CSR or producer pulls the list weekly, prioritizes by premium-change magnitude and client value, and works the high-touch renewals down: pulling supporting documents, drafting the renewal-review email, scheduling the producer call. The work itself is good work, in that it preserves revenue, deepens client relationships, and catches coverage gaps. The problem is throughput. At 1,400 policies on roughly annual cadence, the agency processes around 117 renewals per month. Karen estimates the team works the top 40% intensively; the remaining 60% receive a templated email and a renewal certificate. Lapse and slippage in that bottom 60% is the agency's largest invisible revenue leak.

The lead intake reality. New leads arrive through three channels: the agency's contact form on its website (about 4-6 per week), the main phone line during business hours (handled by Carl or a producer if available), and after-hours voicemail or form submissions (about 8-12 per week, of which a meaningful fraction never receive a callback). Tasha, the new-business intake admin, works the after-hours queue first thing each morning, but her capacity is finite and older leads at the bottom of the queue often go cold. Jim, the personal-lines producer, surfaced this concern unprompted during his interview: he can name three specific leads in the last six months he is fairly certain went to another agency because the response window stretched past 24 hours.

The knowledge-finding pain. Independent agencies live on carrier underwriting guidelines: what each appointed carrier will and won't write, at what rate, with what exclusions, under what conditions. The guidelines change frequently. They live in PDFs, in carrier-portal pages, in internal email threads, and in the experiential knowledge of the most senior staff. When a producer is on the phone with a prospect describing an older home in a flood zone, they need the Travelers underwriting guidance in seconds. They currently get it by interrupting Karen or Carl, or by digging through the SharePoint folder where guidelines are filed. Lina, hired last fall, is the most affected: her first three months were dominated by the "is this writeable, where, at what rate" question, and her ramp time would have been weeks shorter with a working internal carrier knowledge system.

3. The five opportunities

The four "build" recommendations are presented in priority order. The fifth (voice receptionist) follows, with the reasoning for the "not yet" recommendation.

Each entry follows the same structure: headline value, the observed problem, the agent pattern that solves it, the math behind the value claim, install effort, ongoing cost contribution to the monthly retainer, and the risks and dependencies you should weigh.

3.1 Carrier document parser

> **Headline value: ~40 hrs/week recovered · ~$73,000/year Recommendation: build first.**

The problem. Tuesday morning, May 19. I sat at Maria's desk from 8:32am to 10:50am with her permission and a stopwatch. Her HawkSoft-routed carrier inbox held 67 emails when she opened it at 8:35am, a slightly heavy Tuesday, since Monday is the catch-up day for weekend carrier system jobs. She processed 47 before our 11am stop, averaging 2 minutes 14 seconds of handling time per email (the elapsed morning ran longer than the handling time alone, since phones and walk-up questions interrupted her roughly every fourth email). The remaining 20 deferred to the afternoon, joining the next morning's arrivals.

The distribution across those 47 emails was representative of every CSR's morning at the agency:

DOCUMENT TYPE	COUNT	WHAT THE CSR DOES
Declarations pages and endorsements	18	File to policy record
Renewal offers	12	File, flag if premium change > 10%
Non-renewal notices	6	Escalate to producer, file, draft client notification
Billing changes	5	File, occasionally notify client
Routine system notifications	4	File
Coverage change confirmations	2	File, notify client

Across four CSRs, this triage consumes roughly **50 hours per week in aggregate**, which is a full FTE doing rules-bound work that doesn't require a human's judgment for the majority of cases. Mornings at the CSR pod are silent until 10:30 or 11:00. Client-facing work waits for the inbox to clear.

The pattern. The carrier document parser AI agent reads every incoming carrier email, classifies the document type, locates the matching policy in HawkSoft using the policy number and carrier appointment, attaches the document to the correct record, and updates the fields where the carrier-provided data differs from the HawkSoft record (premium, expiration date, coverage limits). For material changes (premium increases over 10%, non-renewals, coverage reductions) it drafts a client notification email and surfaces it for CSR review and send. Edge cases (multi-policy documents, scanned faxes, handwritten amendments) route to CSR review by design rather than guess.

The CSR's role shifts from triage to review. The hour their morning currently begins moves to the hour they currently see their first client.

The math.

INPUT	VALUE	SOURCE
Current state	50 hrs/week aggregate	4 CSRs × ~2.5 hrs/morning × 5 days, confirmed by stopwatch sample
Post-build state	~10 hrs/week aggregate	CSR review pass at ~30 min/morning each
Recovered time	40 hrs/week	
Annualized	2,080 hrs/year	40 × 52
Loaded CSR cost	$35/hr	Base ~$25/hr + 40% load (benefits, taxes, overhead); see Appendix A
Annual value	**~$72,800/year**	2,080 × $35

Install effort: $5,500 – $7,500. Includes HawkSoft API integration (confirmed available at your current contract tier), email ingestion from the carrier-routed inbox, classification model tuning on a sample of your historical carrier documents, the notification-draft template library, and the CSR review interface.

Ongoing retainer contribution: ~$700 – $900/month. Covers monitoring, on-call coverage, and the few hours per month of tuning as carriers change their document formats (typically 2-4 format changes per year across your appointed carriers).

Dependencies & risk factors.

- **HawkSoft API access:** confirmed with your HawkSoft account rep; included in your current contract tier.
- **Carrier format drift:** carriers occasionally change document formats. Tuning is covered by the retainer, and classification accuracy is monitored continuously.
- **Edge-case handling:** multi-policy documents, scanned faxes, and handwritten corrections escalate to CSR review by design. The agent is not autonomous.
- **Onboarding cadence:** 4-6 week shadow period during which the agent runs in observe-only mode, output reviewed by CSRs before any action. Go-live only after sustained >95% classification accuracy on your live mail. This is a build-trust phase, not a soft launch.

3.2 After-hours lead response

> **Headline value: ~$25,000/year in recovered revenue Recommendation: build in month two.**

The problem. New leads arrive through three channels: the agency's website contact form (about 4-6 per week), the main phone line during business hours (handled by Carl or whichever producer is available), and after-hours voicemail or form submissions (about 8-12 per week, counted directly during the week on-site). Tasha, the new-business intake admin, works the after-hours queue first thing each morning, but her capacity is finite and the older items at the bottom of the queue often go cold. Jim, the personal-lines producer, surfaced this in his interview without prompting:

> "I can name three leads in the last six months I'm pretty sure went to another agency because we didn't get back to them in time. There are probably more I don't know about."

The industry data backs the intuition. Studies of inbound lead response across financial-services and insurance consistently find that contact rate drops sharply after the first 5 minutes and continues to fall through the first 24 hours. The 24-hour delay typical for an after-hours form submission to your agency results in roughly half the contact rate of a same-hour response.

The pattern. The after-hours lead response agent receives every after-hours form submission and voicemail. Within 60 seconds, it sends a personalized acknowledgment, asks 3-4 qualifying questions appropriate to the line of business indicated (homeowners, auto, commercial), captures the answers via reply or a brief web form, and books the appropriate producer's calendar based on availability and line specialty. The producer arrives at the discovery call already briefed. The morning queue, instead of holding a backlog of cold contacts, holds pre-qualified, calendar-booked leads.

The math (conservative).

INPUT	VALUE	SOURCE
After-hours leads currently lost to delay	~2/week	Conservative; Jim's gut estimate was higher
Annualized	104 leads/year	2 × 52
Close rate on inbound personal-lines leads	28%	Industry-standard range 25-30%; midpoint used
Incremental policies sold	~29/year	104 × 28%
Avg first-year commission (blended auto + home)	$850	See Appendix A
Year-1 revenue	**~$24,650/year**	29 × $850

The year-1 figure is the conservative headline. Personal-lines policies that stick renew at ~85% annually, so the lifetime value of each recovered policy is roughly 3-4× the year-1 commission. The agent's actual financial contribution compounds over the relationship.

Install effort: $6,500 – $9,000. Includes website contact-form integration, RingCentral voicemail transcription pipeline, the qualification conversation flow, integration with the producers' Microsoft 365 calendars, and the producer-briefing email template.

Ongoing retainer contribution: ~$500 – $650/month.

Dependencies & risk factors.

- **Voicemail transcription quality:** RingCentral's transcription accuracy varies (~85-92%). Cases where transcription is unintelligible get a "could not transcribe your message, please call back at..." response and escalate to the morning queue.
- **Producer calendar discipline:** calendar booking requires producer M365 calendars kept current. Already the case at the agency; flagging because it's a soft prerequisite.
- **Qualification flow tuning:** too many questions feels like a robot; too few leaves the producer without useful context. The 4-week shadow period is when this gets tuned to your prospect base.
- **Voice tone calibration:** Karen will want to read and approve the standard reply templates before go-live. The agent speaks for the agency.

3.3 Renewal & recall outreach

> **Headline value: ~$25,000/year in revenue protection plus winback Recommendation: build in month three.**

The problem. Personal-lines renewals are managed through HawkSoft's renewal report. A CSR or producer pulls the list weekly, prioritizes by premium change and client value, and works the high-touch renewals down. At 1,400 policies on roughly annual cadence, the agency processes around 117 renewals per month. Karen estimates the team works the top 40% intensively; the remaining 60% receive a templated email and a renewal certificate.

The Wednesday renewal call I observed Jim conduct is the high-touch case: a long-term homeowner client whose Mercury policy renewed with a 22% premium increase. Jim's prep took 15 minutes, the call ran 20 minutes, and the client renewed after Jim walked through coverage adjustments that brought the increase to 14%. Without that call, the client would have shopped, and likely would have left, since 22% is comfortably past the threshold where Lafayette homeowners get serious about alternatives.

That kind of call is exactly the work the agency should be doing. The problem is the bottom 60%, the roughly 70 renewals per month receiving templated outreach without the human-touch follow-through. Slippage in that segment is the agency's largest invisible revenue leak. Karen's gut estimate, corroborated by a rough count of policies that lapsed in the last 12 months without contact, is roughly 1.5% annual slippage in the templated segment.

Separately, there is a recall opportunity. Roughly 80 policies lapsed in the last 18 months without systematic re-engagement. Most are not coming back on their own. Some would, with the right outreach at the right interval.

The pattern. The renewal & recall agent pulls policies expiring in 60 days from HawkSoft, classifies each by complexity (premium change, coverage change, claims history, client tenure), and drafts a personalized renewal-review email for CSR or producer review. For high-touch cases, it schedules the producer call. For all cases, it tracks renewal status through close. Separately, the agent works the lapsed-policy queue: every policy that lapsed in the last 18-24 months gets a re-engagement email drafted, with cadence and tone tuned to time-since-lapse.

The agent does not auto-send. Every email is drafted for review. The agent's job is to remove the prep time, not the human judgment.

The math.

INPUT	VALUE	SOURCE
Bottom-60% renewals annually	~840 policies	117/month × 60% × 12
Current slippage rate	~1.5%	Karen's estimate, spot-checked
Post-build slippage rate	~0.5%	Conservative assumption based on systematic outreach
Policies saved	~8/year	(1.5% – 0.5%) × 840
Avg renewal commission	$2,000	Mature policy; higher than new-business; see Appendix A
Slippage prevention	**~$16,000/year**	8 × $2,000
Lapsed-policy winback rate	~10%	Industry standard for 12-24 month lapse window
Policies recovered	~8/year	80 × 10%
Avg winback commission	$1,500	Between renewal and new-business commission
Recall winback	**~$12,000/year**	8 × $1,500
Total annual value	**~$28,000/year**	Conservatively rounded to **$25,000** in the headline

Secondary benefit: the templated email drafts shortcut high-touch renewal prep by ~10 minutes per renewal × 47 high-touch renewals/month × 12 = ~94 hrs/year of recovered producer and CSR time. That figure is not in the headline, but it is real.

Install effort: $4,500 – $6,500. Includes HawkSoft renewal-report ingestion, classification rules, the email template library for both renewal and recall flows, calendar booking integration, and the status-tracking dashboard.

Ongoing retainer contribution: ~$600 – $800/month.

Dependencies & risk factors.

- **HawkSoft renewal-report data quality:** currently good, per spot-check during the week.
- **Voice and tone tuning:** renewal emails are relationship-preserving, not transactional. The 4-week shadow period is critical for getting the voice right before drafts go to CSRs and producers for send.
- **Email deliverability on recall:** outreach to 80 lapsed clients in a short period could trigger spam filters. Recall rollout is staged across 6-8 weeks rather than batched.
- **The agent does not auto-send:** CSR or producer review is mandatory on every outbound. This is intentional, not a limitation.

3.4 Internal carrier knowledge RAG

> **Headline value: ~12 hrs/week recovered · ~$29,000/year Plus: faster new-hire ramp time**
> **Recommendation: build in months four and five.**

The problem. Independent agencies live on carrier underwriting guidelines: what each appointed carrier will and won't write, at what rate, with what exclusions, under what conditions. The guidelines change frequently. They live in PDFs filed in SharePoint, in carrier-portal pages, in internal email threads, and in the experiential knowledge of the senior staff.

When a producer is on the phone with a prospect describing an older home in a flood zone, they need the Travelers guidance in seconds. They currently get it by interrupting Karen or Carl, or by digging through the SharePoint guideline folder. Across the agency, roughly 15 of these questions arise per day. Average resolution time is around 10 minutes, ranging from 2 minutes for routine questions to 30+ minutes for the esoteric ones. Across the team, this consumes roughly 12.5 hours per week, most of which interrupts Karen and Carl, whose institutional knowledge is the bottleneck.

Lina, hired last fall, is the most affected. Her first three months were dominated by the "is this writeable, where, at what rate" question. Karen estimated in her interview that Lina's ramp time would have been weeks shorter with a working internal knowledge system. The next hire, currently being scoped for early 2027, would benefit even more, because the knowledge base will have grown.

The pattern. The internal carrier knowledge RAG (retrieval-augmented generation) agent ingests every carrier underwriting guideline, every internal policy doc, every carrier-portal page the agency has access to, and the institutional email threads worth keeping. Producers and CSRs ask plain-English questions, like "can Travelers write a 1920 home in Lafayette without a wildfire surcharge if it's been recently re-roofed?", and get answers grounded in the actual source documents, with citations. The agent is configured to say "I don't have a confident answer; ask Karen" rather than guess. New documents are auto-ingested as they arrive. When the answer requires judgment, the agent surfaces the relevant sources and routes the question to Karen or Carl.

The interface lives in the producers' Microsoft Teams chat, so questions can be asked from where they already work, not from another tab.

The math.

INPUT	VALUE	SOURCE
Questions per day across the agency	~15	Counted during the week on-site
Avg resolution time per question	~10 min	Observed range 2-30+ min, weighted average
Weekly time	12.5 hrs/week	15 × 10 min × 5 days
Annualized	650 hrs/year	12.5 × 52
Blended labor cost (producers, CSRs, Karen, Carl)	$45/hr loaded	See Appendix A
Annual value	**~$29,250/year**	650 × $45

Secondary benefit, not in the headline: **new-hire ramp acceleration.** Lina's first 90 days had roughly 30% of her billable capacity consumed by knowledge-finding. For the next hire, with the RAG agent in place, this could drop to ~10%. On a $75,000 loaded annual cost for a new producer, that's ~$15,000 of recovered new-hire capacity per hire, realized when hiring happens rather than continuously.

Install effort: $4,500 – $6,000. Includes the document ingestion pipeline (PDFs, web pages, SharePoint), the RAG retrieval system, the query interface (web app + Microsoft Teams integration), and the citation/source-link presentation.

Ongoing retainer contribution: ~$700 – $850/month. Slightly higher than the other agents because an interactive agent needs more ongoing tuning attention than a background one.

Dependencies & risk factors.

- **Source-document curation:** the agency must commit to keeping the source SharePoint folder current. The agent is only as good as its sources. This is a small operational discipline rather than a heavy lift, but it's a real prerequisite.
- **Carrier portal scraping:** some carrier portals have terms of service that prohibit automated access. The agent uses agency-credentialed scraping only where the carrier permits; otherwise the document is ingested via manual PDF export. Travelers, Mercury, Safeco, and CSAA all permit; a few others require manual ingest.
- **Hallucination risk:** managed by grounding answers in retrieved sources and showing citations. The agent is configured to refuse rather than fabricate.
- **Quality compounds over time:** this agent benefits from sustained tuning. The system improves over months as Karen and Carl correct edge cases. Don't expect month-one quality to be the ceiling.

3.5 Voice receptionist: recommended NOT YET

> **Real pain, and insufficient ROI at the agency's current call volume. Recommendation: revisit when call volume crosses ~150/week or the operations manager role changes.**

The problem. Inbound call volume is approximately 80 calls per week, counted from RingCentral logs Monday through Friday. Roughly 75% land during business hours; 25% after hours or during the lunch hour. Carl handles the majority of business-hours calls; producers and admin staff cover overflow. After-hours calls go to voicemail; some are urgent claim notifications that aren't well served by morning callback.

Carl's interview surfaced call handling as a real pain. He estimates 15-20% of his day is consumed by call routing and interruption. The cost is not only the time itself but the context-switching tax it imposes on ops manager work that requires concentration, like vendor escalations, accounting close, and hiring loops.

The pattern, if built. A voice receptionist agent would answer inbound calls 24/7, route business-hours calls to the appropriate party (Carl for general, the relevant producer for known clients, the claims process for claim calls), and after hours, would handle claims intake, lead intake (the 3.2 agent in voice form), and emergency escalation. The receptionist would know your tone, your team, and the routing rules.

Why not yet. The install cost for a quality voice agent is $10,000 to $16,000 all-in: speech model selection, RingCentral integration, custom routing logic, voice-tone tuning, the claim-intake conversation flow, and the lead-intake voice variant of 3.2. The recoverable value at the current 80 calls per week is modest relative to that cost and ranks behind the four recommended agents on leverage:

VALUE SOURCE	ANNUAL ESTIMATE
Carl's reclaimed time (~5 hrs/week × $50/hr loaded)	~$13,000
After-hours claim intake improvement	Real but hard to dollarize; modest
After-hours lead capture	Already captured by 3.2 in form/voicemail form
Total	**~$13,000 – $18,000/year**

The 3.2 agent (after-hours lead response) addresses the highest-value subset of what a voice receptionist would do, at a fraction of the cost. The remaining value (Carl's time and claim intake) would pay back a $10,000 – $16,000 build over roughly a year, but ranks behind the four recommended agents on leverage. The "not yet" call reflects sequencing at the current call volume, not a permanent veto.

When to revisit.

- **Call volume crosses ~150/week.** The threshold where Carl can no longer comfortably absorb routing on top of his ops manager role.
- **Carl's role changes.** Promotion, departure, or a shift in focus to higher-leverage work.
- **The agency adds a third producer who needs dedicated routing**, or extends business hours, or opens a second location.
- **After-hours claim volume increases.** Currently low because most claims funnel through carrier portals rather than the agency.

I am flagging this opportunity in this report rather than omitting it because Carl raised the pain himself and the team should know it has been considered. The recommendation may change in 12-18 months. I would rather you weigh the gap honestly than have me recommend a $16,000 build that doesn't pay for itself today.

---

4. What I'd build first, and in what order

The four "build" recommendations sequence across five months. The order is not arbitrary; each month below explains why that build comes when it does.

Month 1: Carrier document parser (3.1). This is the highest-ROI recommendation in the report and the broadest in impact across the CSR team. Building it first does three things at once: it delivers the largest visible win to the team most affected by current-state pain, it builds the HawkSoft API integration that three of the four agents will share, and it gives the agency a managed-agent track record under real conditions before scope expands. If the carrier document parser cannot meet its accuracy threshold during the shadow period, the right call is to halt scope expansion rather than press forward.

Month 2: After-hours lead response (3.2). Low integration complexity, low risk to ongoing operations, fast revenue impact. The configuration work is tailoring the pattern to your producer calendars, your line specialties, and Karen's voice for the standard reply templates. Most of month 2's effort is configuration and review rather than new construction.

Month 3: Renewal & recall outreach (3.3). This agent depends on the HawkSoft integration from month 1, and the email-template tuning benefits from the team having internalized the review-and-send rhythm established in months 1 and 2. The recall component (re-engagement of the 80 lapsed policies) phases across 6-8 weeks for deliverability reasons, so the revenue impact builds gradually through Q2 of the engagement.

Months 4 & 5: Internal carrier knowledge RAG (3.4). This is the most content-heavy build of the four. Document ingestion is straightforward; the work is in source-document curation (which Karen and Carl need to do alongside me) and in the iterative quality-tuning that takes the system from "useful in 60% of cases" to "trusted in 90%". Two months gives the curation work time to land and gives the team time to build the habit of asking the agent before interrupting Karen.

What your team will experience

During month 1, install week looks like this:

- **Week 0** (week before install): scope-confirmation call with Karen, HawkSoft credential setup, list of carriers to prioritize for the email ingestion sample, and a brief team kickoff so CSRs know what's coming and why.
- **Week 1** (on-site): I'm at the agency for 2-3 days. I provision the managed cloud environment the agents run in, connect it to the HawkSoft API, and ingest 90 days of historical carrier mail for the agent's classification training. Nothing new gets installed in your office and nothing changes on your team's machines.

- **Weeks 2 – 5** (shadow mode): the agent runs in observe-only. Every email triggers a proposed action that goes to CSR review rather than firing. Accuracy is measured and reported weekly. Maria, Sarah, Brian, and Olivia continue working as they do today, but they're also rating the agent's proposed actions, which is how I tune it.
- **Week 6** (go-live decision): if classification accuracy is sustained above 95% across the sample, the agent goes live, and CSRs shift to review-only on the agent's drafted notifications. From that point their mornings start with client work instead of the inbox.

Months 2 through 5 follow the same pattern with smaller week-1 footprints, because the HawkSoft integration and the shared cloud environment are already in place.

Throughout the engagement, the monthly business review with Karen and Carl covers what the agent did the previous month, what I changed, what's coming next, and whether anything in the prior month's behavior suggests a scope adjustment. The report is short; the conversation is the value.

---

5. Pricing for this scope

The pricing here is sized to the four-agent bundle as recommended in Sections 3 and 4. Per-agent standalone pricing is in each opportunity's entry in Section 3. Those numbers are higher than the bundle equivalents because shared infrastructure (HawkSoft integration, cloud environment setup, project setup, shadow-period monitoring) cannot amortize across a single-agent build.

Install (one-time)

COMPONENT	RANGE
Four-agent bundle (3.1 + 3.2 + 3.3 + 3.4)	$16,500 – $22,000
Discovery Audit credit	–$2,500
Install total after credit	**$14,000 – $19,500**

Bundle pricing is approximately 20-25% below the sum of the per-agent standalone ranges in Section 3 because the HawkSoft API integration is built once (rather than once per agent), the cloud environment is provisioned once, the project setup overhead is amortized, and a single shadow-period monitoring framework covers all four go-live decisions.

The variability inside the bundle range maps to scope decisions you and I will make during the design call: the depth of the renewal-recall historical sweep, the breadth of carrier-portal scraping for the RAG agent, and the degree of customization in the after-hours lead-response qualification flow.

The install covers the 4-6 week shadow period for each agent. Monthly retainer billing begins at that agent's go-live, not at engagement start.

Monthly retainer

COMPONENT	RANGE
Operations and on-call coverage	included
Monthly tuning as carriers change formats	included
Monthly business review with Karen and Carl	included
Monthly retainer total	**$2,500 – $3,200**

The retainer total is the sum of the four per-agent contributions listed in Section 3, and it covers everything required to keep the four agents working over time. There are no per-seat licenses and no surprise bills. Set against the roughly $12,700 per month of recovered time and revenue documented in Section 3, the retainer works out to a four-to-five-times monthly return on the fee. The variability inside the range maps to how much ongoing tuning the agents turn out to need in their first year; heavy interactive use of the carrier knowledge agent lands at the higher end.

What is NOT in the price

- **New hardware or software for your team.** The agents run in a managed cloud environment I operate on your behalf; nothing new gets installed in your office.
- **HawkSoft contract upgrades** if you choose to add modules outside the scope of this engagement. Your current contract tier supports the API access the agents require.
- **Carrier-portal credentials** that violate the carrier's terms of service. Where a carrier prohibits automated access, the RAG agent ingests via manual PDF export rather than scraping.
- **Discretionary scope additions during the engagement.** A new agent pattern, a major workflow change, or a large integration outside the four scoped agents would be quoted separately. The retainer maintains the four; new builds are new work.

How the Audit credit applies

The $2,500 you paid for this Audit credits 100% against the install total. If you choose to engage, the install invoice is the post-credit number above. If you choose not to engage, no credit applies; the Audit is the deliverable you paid for, and this report is yours regardless.

6. What I'd not do, yet

Some recommendations in an Opportunity Report are about what to build. Equally important are recommendations about what *not* to build, particularly when a buyer has good reason to expect a recommendation that doesn't hold up under scrutiny. Below are three things I would not recommend for your agency at this stage, along with the reasoning.

The voice receptionist (Section 3.5) belongs in this section too. The "not yet" reasoning is documented there and not repeated here.

6.1 A migration off HawkSoft

You did not raise this as a question, but other consultants might. The category of AMS-migration recommendations targeting independent agencies of your size, typically toward Applied Epic, AMS360, or Vertafore, is well-funded and well-marketed.

The math does not work for your agency. HawkSoft is fit-for-purpose at 1,400 policies and 12 staff. The migration cost (software, data conversion, training, and productivity loss during transition) runs $50,000 to $150,000 for an agency at your scale. The functional benefits over your current state are marginal. The disruption is real and immediate.

Revisit if the agency grows past ~25 staff or the commercial book grows past ~25% of total revenue, which is the point where the multi-user collaboration and commercial-line workflow gaps in HawkSoft become real.

6.2 Predictive analytics on your book

Tooling vendors selling agency-level predictive modeling (cross-sell scoring, churn prediction, lifetime-value segmentation) are active in the independent agency market. The pitch is compelling, but the math doesn't work at your book's size.

Modern carriers run their own segmentation at orders-of-magnitude greater data scale. Agency-level predictive modeling at 1,400 policies does not have enough data to outperform the heuristics your producers already carry in their heads. The vendor's model is statistically thin; your producers' intuition is statistically richer.

Revisit at ~5,000 policies, when the book is large enough that agency-level modeling produces signal the carrier's own modeling doesn't already capture.

6.3 A custom client portal

Your clients have access to the carrier portals already. A custom agency-side portal would duplicate functionality, add support burden, and isn't being requested by your clients in the inbox.

The legitimate version of this recommendation (improving the carrier-portal handoff so clients can self-serve common requests without contacting the agency) is real work, but it's a workflow refinement rather than a portal build. The RAG agent (3.4) does some of this work indirectly by reducing the questions producers and CSRs can't answer quickly.

Revisit if client-portal feature requests start appearing repeatedly in the agency's general inbox.

7. Next step

If the recommendations in this report look like work you want done, the next step is a 30-minute design call. The agenda:

- Confirm the four-agent bundle scope or adjust based on your reading of this report.
- Walk through any sections where the math, the assumptions, or the recommendations don't match your read of the agency.
- Schedule the install. There is no hardware to procure, so the start date is calendar coordination.

I will follow the call with a statement of work documenting the bundle, the install timeline, the final pricing within the ranges in Section 5, and the retainer terms. Sign and return; I'll set a kickoff date two to three weeks out.

If you want to read this report once more first

The offer in this section stands regardless of whether you engage:

> **A 30-minute review of this report, at no additional charge.** Bring questions, push back on numbers you disagree with, or walk me through anything I missed. Karen, you'd benefit most from this if you're forwarding the report to Carl or your accountant for a second read.

The Audit is the deliverable. This report is yours, and the review call is a way to make sure you got what you paid for, independent of whether the engagement continues.

Appendix A. Methodology and inputs

This appendix documents the inputs and benchmarks behind the dollar values in Section 3, so you can substitute your own numbers and re-run the math if you disagree with any of mine.

Inputs used in the report

INPUT	VALUE	NOTES
CSR loaded hourly cost	$35	Base ~$25/hr + 40% load (employer-paid taxes, benefits, paid time off, overhead allocation). Adjust if your actual numbers differ.
Producer loaded hourly cost	$60	Producers earn commission as well; this figure is loaded time-cost only, not opportunity-cost of production.
Blended labor cost (producers + CSRs + ops + admin)	$45/hr	Weighted average across the team for the carrier knowledge RAG calculation.
Avg first-year personal-lines commission	$850	Blended across auto and home for inbound leads. East Bay homeowners with current carrier-market conditions skew this higher than national averages; conservative use here.
Avg renewal commission, mature personal-lines policy	$2,000	Higher than first-year due to growth in premium base and bundling over a multi-year client relationship.
Inbound personal-lines close rate	28%	Industry range 25-30%; midpoint used. Your actual close rate may differ.
P&C personal-lines annual retention	85%	Industry standard. The agency's actual retention is unmeasured but anecdotally consistent with this range.
Lapsed-policy winback rate (12-24 month window)	10%	Industry standard for systematic re-engagement campaigns.
Slippage rate, templated-segment renewals	1.5% currently → 0.5% post-build	Current rate is Karen's estimate corroborated by spot-check; post-build rate assumes systematic outreach without coverage of every single low-touch renewal.

Where the benchmarks come from

The industry benchmarks above are sourced from publicly available data and standard industry references, including:

- Independent Insurance Agents & Brokers of America (Big "I") agency operations research.
- Public carrier rate filings for the California personal lines market (CDI public filings).
- Independently published research on inbound-lead response time and contact rate (notably Oldroyd, McElheran, and Elkington's response-time research, originally published in Harvard Business Review and widely replicated).

Where a number is Karen's estimate or my direct observation during the week on-site, the source is named inline in Section 3.

How to substitute your own numbers

Each math table in Section 3 names the inputs separately from the result. If you believe your CSR loaded cost is $42/hr rather than $35, multiply 3.1's recovered-hours figure (2,080) by the new rate ($42) to get a revised annual value of $87,360 rather than $72,800. The methodology is the same; only the inputs change.

The same applies to close rate (3.2), slippage rate (3.3), and team labor blend (3.4). Every input is named and adjustable.

If you'd like me to re-run any of these calculations with your actual numbers, I'll do it on the review call described in Section 7.

End of report.

Sample: Acalanes Insurance Partners is a fictional composite. All names, figures, and findings are illustrative. Prepared by Lee McDole, McDole Group, May 25, 2026.