AI BUYER'S GUIDE

AI for the Independent Insurance Agency Back-Office: A Henderson NV Buyer's Guide

What I'd actually buy, defer, or ignore for the back-office at a 3-15 employee P&C / commercial agency in 2026.

By Piers Rollinson, DomeWorks

Henderson, Nevada

The 30% Problem

I keep having the same conversation with Henderson P&C agency owners. Someone walked them through a demo of an AI intake tool, quoted a 40% reduction in data entry time, and they want to know if it's real. The number itself isn't fabricated. The problem is that it usually comes from a vendor's own customer survey, applied to a large book of personal-lines business, and then dropped into a conversation about a seven-person commercial agency where half the renewal submissions require supplemental forms that the tool doesn't touch. The gap between the demo number and your actual situation is where the money goes.

Here's the underlying reality: independent insurance agency staff report spending 30 to 50% of their day on duplicate data entry across forms, email, AMS, and carrier portals (Big "I" Future One 2024 Agency Universe Study; Vertafore 2024 Insurance Agency Workforce Report). ACORD forms average 8 to 12 minutes each with a roughly 20% manual error rate before rekeying corrections are factored in (Infrrd/Regure vendor benchmarks). McKinsey's "Insurance productivity 2030" puts 30 to 40% of underwriter time nationally on administrative rekeying and manual analysis. Those aren't fringe numbers. They're the reason your service staff is still at their desks at 5:45 PM on a Thursday when they should have been done at five. At the same time, Reagan Consulting's Q1 2024 GPS put organic growth at 8.4%, the lowest in 11 quarters. The next dollar of margin for a 3-15 person Henderson agency doesn't come from writing more new business. It comes from doing the existing business faster.

There's a second number underneath the first one. The labor cost is one thing. The role that's missing is the bigger thing. Every team running AI well has a person whose actual job is to figure out where AI agents create real leverage, map the workflow around them, connect the systems, set up the checks that catch errors before they reach clients, and own the KPIs week after week. Most Henderson agencies don't have that person. They have a tool, sometimes two, and an owner-operator hoping the team will adopt it on top of an already-full week. That's why Year 1 ROI usually doesn't show up. The tool is fine. The role isn't there. I'll come back to this in the math section, because it's the line item the cost stack everyone's modeling has been quietly leaving out.

What this guide is not trying to do is sell you on AI. I don't resell any of the tools named here. What I'm trying to do is give you the buyer's view that vendors won't: an honest read on what each workflow costs, what tools actually target it, where the evidence is thin, and what you'll run into when it stops matching the demo. Every vendor-sourced productivity claim in this guide is labeled accordingly. Where independent evidence exists, I use it. Where it doesn't, I say so.

TLDR

If you only read one section, here's what I'd want you to take away.

1. **The bottleneck is back-office labor.** Independent agency staff spend 30 to 50% of the day on duplicate entry, ACORD form rework, and renewal correspondence. That's where the time and money are being lost.
2. **The throughput frame matters more than the savings frame.** AI doesn't mostly save you hours. It lets a 5-person agency requote, summarize, and follow up at the volume of a 10-person agency. The agencies that win the next 24 months won't be the ones with the cheapest stack. They'll be the ones who actually run their stack.
3. **AI compresses repetitive work. It does not replace licensed judgment.** Renewal letters, intake summaries, and COI tracking are well inside AI's range. Coverage advice, suitability determinations, and ACORD forms submitted without human review are not.

4. **Year 1 only works if someone owns it.** A minimal 5-person stack runs $35,000 to $65,000 in Year 1; a full buildout reaches $110,000. Recovered capacity is worth $14,000 to $46,000 in Year 1 only if adoption lands at 70% or better, and adoption only lands at 70% if a named person inside the agency is responsible for making it land. Most Henderson agencies don't have that person yet. Without that role, the tools become shelfware and the math doesn't pencil.
5. **Nevada regulators and your E&O underwriter are already asking about AI.** The "What Not to Get Burned By" section covers the five failure modes with specific NRS citations and E&O market signals. Have answers before your next renewal conversation.

> *Independent agency staff spend 30 to 50% of their day on duplicate data entry across forms, email, AMS, and carrier portals.*
>
> BIG "I" FUTURE ONE 2024 AGENCY UNIVERSE STUDY

Three Buckets

Every AI tool pitched to independent agencies in 2026 falls into one of three buckets. The bucket that's right for you depends on what AMS you run, which workflows are actually your top time-sinks, and how much integration plumbing your operations team has capacity to manage. Getting this framework right before you book a demo saves you from buying the wrong class of tool for the problem you actually have.

Bucket A: Horizontal LLM wrappers. ChatGPT Business, Claude Team, Microsoft 365 Copilot, Gemini Business. These are general-purpose AI tools that agencies are adapting for renewal-letter drafting, email summarization, internal knowledge retrieval, and meeting notes. They're fastest to trial because staff already live in email and documents, and the per-seat cost is published and SMB-accessible: ChatGPT Business at about $20-25/seat/month, Microsoft 365 Copilot Business at about $18-21/seat/month (SMB tier through mid-2026), Claude Team at about $20/seat/month billed annually. The tradeoff is that none of these tools integrate natively with Applied Epic, AMS360, HawkSoft, or EZLynx. You're borrowing AI capability for workflows that live outside your AMS. Useful, but you're building the bridge yourself.

Bucket B: Vertical insurance-native tools. Quandri, Canopy Connect, Roots, Indio, Cara, Indemn, Sixfold, Agentech, and Semsee (now acquired by iBynd as of March 2026). These tools are built for insurance workflows. They touch ACORD forms, carrier appetite data, and AMS records in ways horizontal copilots don't. Quandri's April 2026 launch confirmed live integrations into Applied Epic, Vertafore AMS360, and HawkSoft [vendor claim, Quandri press release April 2026]. Canopy Connect Elevate publishes pricing at about $250/month for five users, which is rare transparency in this category. Most others are sales-led and require a demo conversation before you see a number. The tradeoff: closer to your actual bottleneck, but pricing opacity means you're often buying before you can fully model the ROI.

Bucket C: Embedded AMS AI. Applied Insurance AI (Book Builder, Epic Max, AutoFill via Cytora, acquired September 2025), Vertafore Client Communications plus AgencyZoom, NowCerts Momentum AMP. These sit inside your system of record. NowCerts is the clearest pricing example: Essentials at $99/month, Professional at $169/month, Business at $349/month. Applied and Vertafore do not publish separate pricing for AI add-on tiers. As a planning placeholder, budget 15 to 30% over your base AMS cost (uncited estimate, based on comparable SaaS tiering patterns rather than vendor disclosure; confirm actual pricing before signing). The advantage is that your policy data is already in the system, so integration friction is lowest. The downside is that AMS vendor AI innovation cycles are slower than dedicated insurtech startups, and you're adding another reason your contract renewal becomes difficult to walk away from.

How to choose before the demos start

Bucket	Best fit	Starting price signal	Integration depth
A: Horizontal copilots	Email, drafting, summarization, knowledge retrieval	$18-25/seat/mo (published)	None native to insurance AMS
B: Vertical insurance-native	Renewal, intake, quoting prep, COI tracking	$250/mo (Canopy); others opaque	Live AMS integrations vary by vendor
C: Embedded AMS AI	Workflows that live inside your AMS	NowCerts published; Applied/Vertafore opaque	Deepest; data already in AMS

The decision frame I'd suggest: buy Bucket C (embedded AMS AI) when your AMS vendor has published a feature that covers the workflow you're targeting. Add Bucket B (vertical) when your AMS doesn't cover it. Use Bucket A (horizontal copilots) for everything else, with PII guardrails in place before anyone uses them on client data.

One consolidation note worth knowing: Applied acquired Planck in July 2024 and Cytora in September 2025. Semsee was acquired by iBynd in March 2026. The embedded AMS suites are actively absorbing the vertical tier. If you're evaluating a standalone vertical tool, ask what happens to your data and integrations if that vendor gets acquired or shuts down before you're done with your contract.

A

Horizontal LLM wrappers

General-purpose, fastest to trial

- ChatGPT Business (~$20-25/seat/mo)
- Claude Team (~$20/seat/mo)
- Microsoft 365 Copilot Business
- Gemini Business

B

Vertical insurance-native

Built for insurance workflows

- Quandri (renewal + requoting)
- Canopy Connect (intake)
- Roots, Indio, Cara

- Indemn, Sixfold, Agentech

C

Embedded AMS AI

Deepest integration, recommended start

- Applied Insurance AI (Epic Max / Cytora)
- Vertafore Client Communications + AgencyZoom
- NowCerts Momentum AMP
- HawkSoft (API partners)

RECOMMENDED STARTING POINT

Three buckets for insurance agency AI tools. Start with Bucket C if your AMS vendor has a published AI module. Add B for specific workflow gaps. Use A for everything else.

> *Buy Bucket C when your AMS vendor has published a feature that covers the workflow. Add Bucket B when it doesn't. Use Bucket A for everything else, with PII guardrails in place.*

Workflow-by-Workflow Read

Five workflows that account for most of the back-office time drain in a 3-15 person Henderson P&C and commercial agency. Each section describes the current-state pain, names the tools with the best evidence base, gives one concrete data point, and flags where the honest-broker hedge belongs. If a claim comes from a vendor, it's labeled. If independent evidence exists, I use it. If it doesn't, I say so.

Workflow 1: Client intake and fact-find documentation

Your CSRs spend part of every new-business and renewal cycle chasing missing fields in ACORD forms, re-entering data from PDFs and emails into the AMS, and following up with clients when supplemental forms come back incomplete. The downstream cost is real: every quoting prep and compliance step that comes after intake inherits whatever errors and gaps intake produced. Getting intake right is where AI has its highest leverage, because the payoff compounds across every subsequent workflow.

The published benchmark is ACORD forms averaging 8 to 12 minutes each manually, with roughly a 20% manual error rate (Infrrd/Regure vendor benchmarks). That's the number to keep in your head. AI-assisted intake tools claim to cut that to 2 to 5% error rates and 5 to 10 minutes per form [vendor claim across the category]. **Canopy Connect Elevate** is the most pricing-transparent option I found: about $250/month for five users, with direct integrations into Applied Epic and EZLynx. Canopy's headline 5.6-second policy data transfer [vendor claim] is data movement (carrier-to-AMS), not form completion. It's useful for the prefill step, but it does not by itself shrink the 8-to-12-minute form-fill baseline. **Indio** (now owned by Applied Systems) carries a library of 14,000-plus forms, and one published customer example has a completed signed application returned in 12 minutes [vendor claim, Indio case study]. For document-heavy commercial books where the forms are less standardized, **Hyperscience** and **Docugami** are intelligent document processing options that go beyond simple OCR. For agencies already on Microsoft 365, **Copilot** can assist with email-to-structured-data drafting without requiring a new system.

The honest-broker note here: every published time-savings figure for intake AI is vendor-sourced. No IIABA Best Practices Study or independent peer-reviewed audit of AI intake performance at SMB agencies under 15 employees exists as of mid-2026. Use vendor-published case studies as directional evidence that a category of time can be recovered, not as a commitment that your specific book and carrier mix will see the same result.

Workflow 2: Policy comparison and quoting prep

Multi-carrier quoting prep is the most-cited time sink I hear about from Henderson agencies. One agency quoted in Quandri's own press release described spending "up to one hour, 15-plus screens per renewal" for commercial line reviews (Gibson Agency, [vendor case study]). The core problem is duplicate entry into raters and carrier portals, carrier appetite changes that require market-searching before you even start quoting, and commercial submissions that need manual field mapping across systems that don't talk to each other.

IVANS reported that 83% of agents would write more business if carriers provided real-time appetite and quoting in management systems, and 73% say they often lose opportunities because they can't find a market to quote [vendor claim, IVANS "State of Agency-Carrier Connectivity"]. For personal lines and standard commercial, **EZLynx Rating Engine** claims connections to

330-plus carrier markets [vendor claim]. **Applied Epic Quotes Commercial Lines / Tarmika** puts more than 35 markets inside Epic without leaving the AMS [vendor claim, Applied]. **Quandri Renewal Reviews and Requoting** targets the repetitive portion of the commercial renewal cycle specifically, with live integrations into Applied Epic, AMS360, and HawkSoft as of April 2026. **IVANS Ask Kodiak** handles appetite search, which is often the first bottleneck before quoting starts. Comparative raters typically show 40 to 60% time savings on routine renewals [vendor claim across rater category]; complex commercial renewals see smaller gains because carrier-specific appetite and supplemental requirements still require manual handling.

The honest-broker hedge: comparative rater coverage is real but incomplete. If you write specialty commercial, surplus lines, or niche programs, no rater covers 100% of your markets. The tools do the most for standard personal lines and commercial habitational, not for the hard-to-place accounts that already take your team the longest.

Workflow 3: Compliance and suitability paperwork

Compliance is where Nevada-specific risk concentrates for your agency. The Nevada DOI approved a best-interest annuity suitability regulation in December 2024, raising the documentation bar for any agency writing life and annuity alongside P&C. Nevada Bulletin 24-001 (February 2024) adopted the NAIC Model Bulletin verbatim, requiring documented AI governance, model testing, and audit-trail records at Section 4.2 if your tools touch regulated functions. Audit-trail gaps during a DOI examination translate directly to adverse market-conduct findings.

The current-state pain is carrier-specific forms arriving in different formats, compliance packets that come back incomplete, supervisory documentation scattered across email and a shared drive, and no reliable way to show an examiner which version of a suitability form was completed and when. **iPipeline iGO** is the established e-app tool for life and annuity, with adoption across dozens of carriers, hundreds of distributors, and tens of thousands of agents [vendor claim, iPipeline]. **RegEd** handles annuities training compliance, serving roughly 200,000 producers and more than 80% of annuity carriers [vendor claim, RegEd]. **Indemn** markets NAIC-aligned audit trails for compliance documentation [vendor claim]. For the AI prefill side, horizontal copilots can assist with form completeness checks, but they require careful prompt design to avoid generating coverage language that your E&O underwriter would object to.

The honest-broker note is not optional here: no AI tool eliminates the licensed-review requirement on suitability decisions. AI can prefill, flag missing fields, and check form completeness. It cannot determine whether a product is suitable for a given client. Treat any AI-generated suitability or coverage output as draft-only, requiring licensed sign-off before client delivery or carrier submission. This is not a conservative posture. It's what Nevada Bulletin 24-001 and NRS 686A.030/.040 require regardless of what a vendor demo implies.

Workflow 4: Renewal correspondence and remarketing

Renewal automation is the single workflow where the math works most reliably for SMBs, and it's where I'd start if I were running a 5-to-10 person Henderson agency building an AI stack for the first time. The work is repetitive, the failure mode of doing it poorly is client churn rather than a compliance event, and the tooling is more mature than anywhere else in this category.

Applied Systems' 2024 Voice of Customer survey found that managing renewals was the single most inefficient identified task across their agency base [vendor claim, Applied VoC 2024]. Quandri's published Vienneau Insurance case study claims 6,000 hours per year saved at a mid-size agency [vendor claim, Quandri/Vienneau]. I want to be specific about that number: it's one published agency, not a randomized study, and it involves a book size and team composition that may or may not match yours. Even 20% of that range scales cleanly to a payback period for a 5-to-10 person agency, which is why I use it as directional evidence rather than a commitment. **Quandri Renewal Reviews** targets the personalized outreach and review cycle. **AgencyZoom** (now owned by Vertafore) handles renewal automation inside a CRM-style workflow and reports 5,000 agencies and 35,000 users [vendor claim]. **Indio** handles renewal application prefill with its 14,000-plus form library. **EZLynx EVA** and **Applied Epic Book Builder** are the embedded-AMS renewal tools for agencies already on those platforms. For agencies that primarily need renewal letter drafting, **Microsoft Copilot** or **Claude Team** can handle that task with no insurance-specific integration required, as long as a licensed team member reviews before sending.

The honest-broker hedge: "6,000 hours/year saved" is a vendor-published case study from one agency, not an audited average. Your mileage depends on book size, how many manual touchpoints your current renewal process has, and whether your team actually changes the handoff process or just adds a new tool on top of the old one.

Workflow 5: Claims and underwriting follow-up

This is the smallest direct-AI opportunity for a retail P&C agency, and I want to be honest about why. Most of the work in claims handling lives at the carrier: adjuster assignment, coverage analysis, reserve decisions, settlement authority. What retail agencies actually do is track status, draft status emails, request loss runs, and manage certificate-of-insurance issuance for commercial accounts. That's a meaningful time drain, but the right tools here are often simpler than the ones vendors want to show you.

The clearest published data point I found is the Certificial-Zurich North America integration: Zurich saw more than 50% faster complex renewal certificate packages after implementing Certificial-Momentum (BusinessWire, January 2025, third-party reported). That's the cleanest data point in this section because it's reported by a third-party news outlet citing an actual enterprise implementation, not a vendor case study on their own website. Certificial estimates that manual COI tracking eats 3 to 4 hours per person per day in compliance-heavy commercial books [vendor claim, Certificial]. **Appulate LossRunner** is free for agents and publishes a 40-minutes-saved-per-submission figure [vendor claim, Appulate]. **Roots** reported 2,700 hours saved at Eastern Alliance Insurance Group (third-party reported via Insurance Innovation Reporter, which is a step above a vendor white paper). For agency-facing RPA on repetitive portal-check and email-drafting tasks, a basic Zapier or n8n workflow connecting your carrier portal status emails to your AMS task queue will often outperform a purpose-built AI tool on first-year ROI, because rule-based automation is faster to configure and easier to audit than a fine-tuned LLM.

The honest-broker note: carrier-side AI (Sixfold, Federato) is where the big underwriting productivity numbers come from. As a retail agency, you benefit indirectly through faster carrier turnaround, not through tools you deploy yourself. Don't buy a carrier-side underwriting AI tool as an agency. It's not designed for you and the ROI won't materialize.

> *AI compresses repetitive work. It does not replace licensed judgment. Renewal letters and intake summaries are well inside AI's range. Coverage advice and ACORD forms submitted without human review are not.*

Which Bucket, Which Workflow First

Sequence matters more than tool selection. I've seen agencies spend Year 1 budget on a vertical intake tool their CSRs stopped using by month three, because nobody had mapped the actual intake workflow before buying the tool. The decision tree below is designed to get you to a single starting point, not a shopping list.

Work through the questions in order. Stop at the first "yes."

Q1: Does your AMS vendor (Applied, Vertafore, NowCerts, HawkSoft) have a published AI module or feature that covers the workflow you want to improve?

Yes: Start with Bucket C. The data is already in your AMS and integration friction is lowest. Get a working demo with your actual AMS version before signing anything.

No: Continue to Q2.

Q2: Is your top back-office bottleneck renewal correspondence, client intake, or quoting prep specifically?

Yes: Start with Bucket B. Match the vertical tool to the specific workflow (Quandri for renewal and requoting, Canopy Connect or Indio for intake, EZLynx / Tarmika / Quandri for quoting prep). One tool, one workflow, 90 days.

No: Continue to Q3.

Q3: Is the bottleneck primarily email drafting, document summarization, meeting notes, or internal knowledge retrieval?

Yes: Start with Bucket A. A single ChatGPT Business or Claude Team seat, used under an Enterprise tier with a signed BAA if any client data is involved, is faster to trial and lower-risk than a vertical deployment. Map the prompts your team actually needs before buying seats for everyone.

No: Continue to Q4.

Q4: Are you trying to fix renewal correspondence, intake, quoting prep, and email drafting all at once?

Stop. Pick one bottleneck. Run a 90-day pilot. Review adoption honestly at week six. If adoption is below 60% of the team actively using the tool weekly, kill it before month three and diagnose why before adding a second tool. Buying four tools simultaneously and getting 20% adoption on each is a common way to spend $30,000 and end up no better off than before.

The deeper principle here is that the sequence of implementation compounds. Fixing intake quality first improves quoting prep. Improving quoting prep makes renewal reviews faster. Trying to run all five workflows in parallel means none of them get the workflow-redesign attention they require to actually stick.

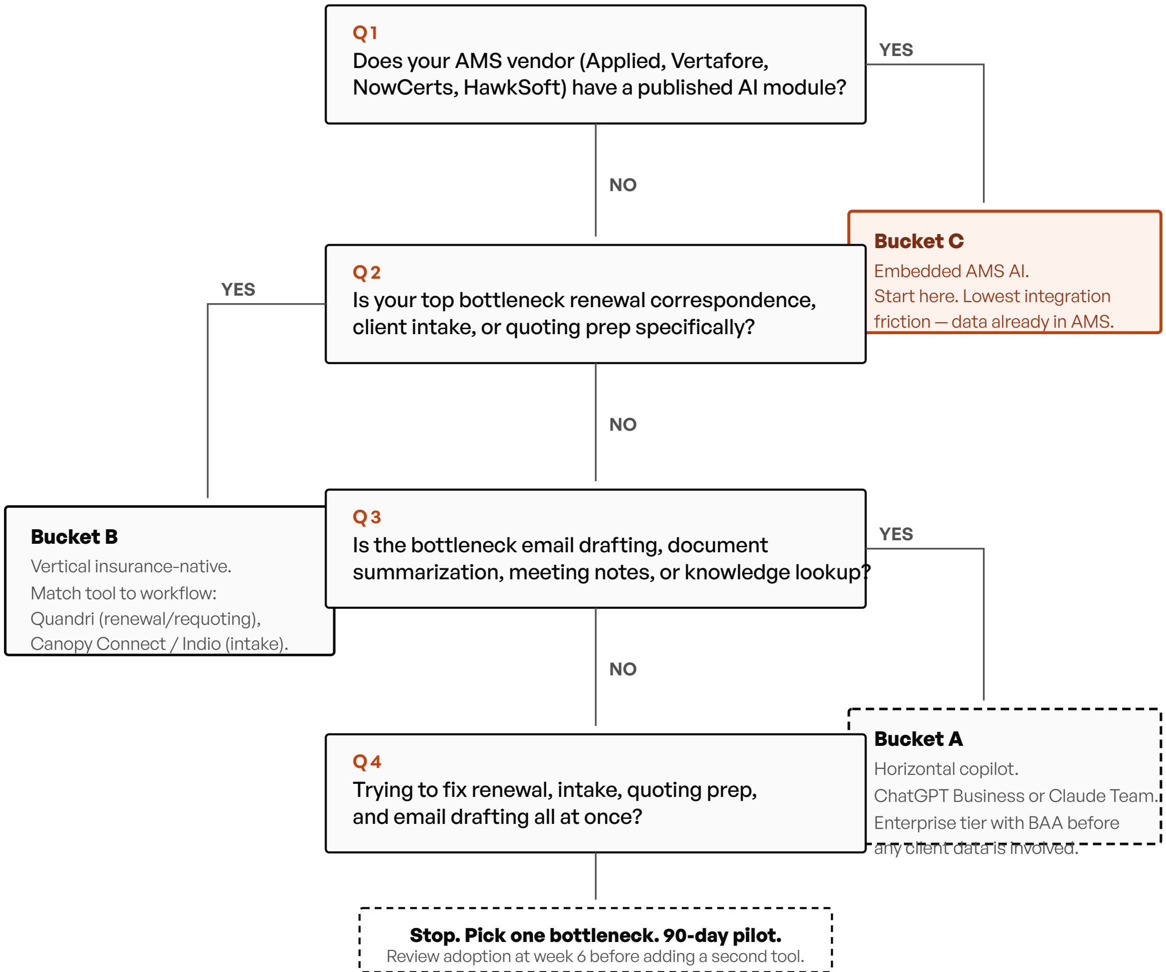

Walk any back-office workflow through these four questions in order. Stop at the first yes. The prose decision tree in the guide provides vendor names for each path.

The Math

What an AI back-office stack actually costs and what it actually saves for a 3-15 person Henderson P&C and commercial agency in Year 1 and Year 2.

Year 1 cost stack

The ranges below reflect the planning uncertainty in this market. Most vendors don't publish pricing. Where they do, I use it. Where they don't, I've estimated from comparable SaaS tiers and implementation benchmarks. Every uncited estimate is labeled.

Line item	Range	Confidence	Source
Horizontal copilot (1-2 seats, annual)	$1,440-$1,800	High	Published pricing: ChatGPT Business, Claude Team
Vertical AI tool (Quandri-class)	$18,000-$60,000	Low	No public price; estimated from SaaS norms [uncited estimate]
Embedded AMS AI add-on premium	$0-$6,000	Very low	Vendors don't disclose; budget 15-30% over base AMS cost [uncited estimate]
Setup and integration (consulting + training)	$5,000-$25,000	Medium	SMB SaaS implementation benchmarks [uncited estimate]
Implementation labor (agency staff, 40-160 hrs at $40/hr opportunity cost)	$1,600-$6,400	Medium	Internal time estimate [uncited estimate]

Line item	Range	Confidence	Source
Ongoing supervision (5-10 hrs/week for output review and QA)	$5,616-$11,284/yr	Medium	Henderson CSR at about $21.67/hr (Indeed.com April 2026)
AI ops lead (the line item most stacks leave out, see callout below)	$0-$180,000/yr	High	Existing repositioned role / fractional / full-time hire
Total Year 1, tool stack only	**$31,656-$110,484**	Medium	Synthesized from all three research tools

The wide range is real, not a hedge. A minimal stack (one horizontal copilot seat plus one vertical tool with no AMS upgrade) compresses toward the lower end: roughly $35,000 to $65,000 for a 5-person agency if setup is tight and implementation is managed internally. That's the planning scenario I'd use. The top of the range reflects a full embedded AMS AI upgrade plus a vertical tool plus outside consulting help, which is what happens when agencies underestimate implementation time and hire someone to rescue the rollout.

The line item the cost stack is missing: the person who runs this

Every team that's actually getting value from AI right now has a person whose job is to run it. The work is concrete: figuring out which workflows the AI is going to take real time off, mapping how data has to move (from carrier portal to AMS to email), designing where the licensed team member checks the AI's output before it reaches a client, writing the checks that catch when an AI-drafted renewal letter gets a coverage term wrong, and tracking whether the tool is actually moving the KPI it was supposed to. This person is technical enough to connect your AMS to the rest of your systems and operational enough to live inside the service team instead of bouncing between IT and ops.

Three realistic ways a 3-15 person Henderson agency staffs this:

Path	Cost	When it works
Repositioned existing operations lead (gives them training time, a clear mandate, and budget for tools)	Roughly 30-50% of one existing salary's time, redirected	You already have someone curious enough to be running ChatGPT in their evenings. Give them the role officially before someone else hires them away.
Fractional AI ops lead (outside specialist, 5-15 hrs/week)	$40,000-$80,000/yr	You don't have the internal candidate. You don't yet have enough scale to hire full-time. You need someone who's done this for similar agencies and will own the KPIs.
Full-time AI ops hire	$120,000-$180,000/yr (loaded)	You have 25+ employees, multiple service teams, and the orchestration problem is hand-offs between teams, not tool selection.

If none of these line items appears anywhere in your Year 1 plan, the stack is going to underperform the math above by 30 to 50%. Adoption rarely lands at 70% on its own. The shelfware path is what happens when you buy the seats without naming the person who owns making them work.

Recovered capacity

What does a well-implemented stack actually free up? The numbers below apply per FTE in a CSR or operations role whose work is heavy on the five workflows covered in this guide.

Workflow	Recovered time per week
Data entry and intake rekeying	5-10 hrs/week

Workflow	Recovered time per week
Quoting prep (comparative rater + appetite search)	2-5 hrs/week
Renewal correspondence	1-3 hrs/week
Claims and underwriting follow-up	1-2 hrs/week
Total per FTE	**9-20 hrs/week**

At $30 to $35 per hour fully loaded (consistent with Nevada BLS wage data for insurance operations roles), that's $270 to $700 per week per FTE, or $14,040 to $36,400 per year per FTE. For a 5-person agency with two to three CSR-equivalent roles, the ceiling is $72,800 to $109,200 in annual recovered capacity at full implementation.

I want to be direct about why those numbers don't just land in your bank account. Capacity is usually reinvested into book growth, not cash savings. Unless you simultaneously fix the workflow and build in adoption discipline, buying seats without changing handoffs does not create capacity. The CSRs are faster; they absorb more accounts instead of leaving at 4:45. That's good, but it doesn't show up in Year 1 as a P&L line unless you've connected recovered hours to either reduced overtime or increased revenue from accounts that now get the attention they weren't getting before.

What Year 1 actually looks like on a 5-person Henderson agency

If you're a 5-person Henderson P&C and commercial agency, plan on $35,000 to $65,000 in Year 1 total cost (tooling, setup, internal implementation labor, ongoing supervision). Against that, plan on $14,000 to $46,000 in recovered capacity if you have two CSR-equivalent roles and adoption lands at 70% or better. Year 1 is marginal to negative. Year 2 is plausibly positive as fixed setup costs are gone and prompt and workflow assets compound.

The confidence note matters here: no independent ROI study for agencies under 15 employees exists as of mid-2026. Every ROI figure in this category comes from vendor case studies with selection bias (the agencies that succeeded are the ones vendors put in press releases) or large-enterprise extrapolations from Bain, McKinsey, and Celent, which sampled top-100 brokers and carriers, not 7-person Henderson P&C shops. Use the numbers in this section as planning placeholders, not commitments.

One more thing: the single most important missing data point in this entire analysis is the E&O premium impact of AI use at your next renewal. I can't give you that number. Nobody publishes it. The only way to get it is the conversation with your E&O broker described in section 11. That conversation belongs in your Year 1 math before you sign any vendor contract.

3-15 PERSON HENDERSON P&C / COMMERCIAL AGENCY, YEAR 1

$31,656 - $110,484 tool stack only

Horizontal copilot (1-2 seats) ChatGPT Business / Claude Team	**$1,440 - $1,800/yr**
Vertical AI tool (Quandri-class) Insurance-native workflow tool	**$18,000 - $60,000/yr**
Embedded AMS AI add-on Applied / Vertafore / NowCerts premium	**$0 - $6,000/yr**
Setup + integration Consulting + training	**$5,000 - $25,000**
Implementation labor 40-160 hrs at $40/hr opportunity cost	**$1,600 - $6,400**
Ongoing supervision 5-10 hrs/wk QA at Henderson CSR rate	**$5,616 - $11,284/yr**

AI ops lead (the missing line it...
Fractional to full-time; most stacks leave this out
$40,000 - $180,000/yr

AI ops lead bar reflects fractional-to-full-time range; most agency cost models omit this entirely. Without it, adoption rarely reaches 70% and recovered-capacity math doesn't pencil. Minimal 5-person stack (one copilot + one vertical tool, no AMS upgrade): approximately $35,000-$65,000 in Year 1.

Year 1 cost stack for a Henderson independent insurance agency. Bars are proportional to midpoint. AI ops lead highlighted because it is the line item most stacks leave out.

> *Year 1 only works if someone owns it. Without that role, the tools become shelfware and the math doesn't pencil.*

ONE A MONTH, NO DRIP

Want the next monthly note from the field?

I send one email a month with what's actually working in small-firm AI rollouts and what isn't. Same standard as this guide. Drop your email if you'd like the next one.

To subscribe from the PDF: email piers@domeworks.tech with "Monthly notes" in the subject.

What Not to Get Burned By

Five concrete burn modes that are already actionable, not theoretical. Three connect to Nevada regulation. Two connect to what your E&O underwriter will ask about at your next renewal. Holland and Knight confirmed in May 2025 that no Nevada DOI enforcement action against an agency specifically for AI use had surfaced through Q1 2026. That's the order to plan around: E&O market signals first, regulators second.

Burn 1: Hallucinated coverage advice

Your CSR asks an AI tool to draft an email to a client explaining why their claim was declined, or to summarize coverage terms on a commercial general liability renewal. The tool produces a confident, fluent paragraph. Your CSR reviews it quickly and sends it. It contains a material misstatement about the policy terms.

This hits NRS 686A.030 and NRS 686A.040 (false and misleading statements in connection with the sale, solicitation, or negotiation of insurance) and NRS 686A.310 (unfair claims practices). Administrative fines run up to $500 per act for licensed agents under NRS 686A.183, and license suspension is on the table for repeated violations.

The hedge is not complicated: treat all AI-generated client communications as draft-only and require a licensed team member to review and approve before sending. This is not a conservative position. It's what the statute requires regardless of whether AI was involved. The only change AI introduces is the volume and speed at which incorrect drafts can reach your outbox if you don't put the review step in the workflow.

Burn 2: PII leakage via unvetted AI tools

A team member pastes a client's name, policy number, social security number, and loss history into the consumer version of ChatGPT to draft a summary for a renewal presentation. The session data may be used in model training. You have no data processing agreement, no Business Associate Agreement, and no audit trail.

NRS 603A.040, .210, and .215 require encryption and appropriate security measures for personal information including Social Security numbers, medical and health insurance IDs, and other regulated identifiers. A violation is a deceptive trade practice under NRS 603A.260. If the client data includes any health insurance information, HIPAA penalties can reach $2.13 million per year at the highest tier (HHS-OCR 2024).

The hedge: use only Enterprise tiers of AI tools, with signed Business Associate Agreements where PHI is involved and signed Data Processing Agreements for all PII. ChatGPT Business, Claude Team, and Microsoft 365 Copilot all offer BAAs on Enterprise tiers. Consumer tiers do not. Create a written policy stating that no client-identifying information is entered into any consumer AI tool. Document that the policy exists and that staff received it. This documentation belongs in your Bulletin 24-001 compliance file.

Burn 3: The Nevada Bulletin 24-002 filed-rate trap

Nevada Bulletin 24-002 (August 27, 2024) is the one that most agencies haven't heard about. It states that any mathematical model used in underwriting or rating, including surcharging and non-renewal decisions, must be filed with the Nevada Division of Insurance for prior approval under NRS 686B.110. The bulletin explicitly says calling something an underwriting model rather than a rating model doesn't change the filing requirement.

Most agency AI deployments don't touch underwriting or rating decisions. But some tools, particularly AI-assisted quoting tools or appetite-screening tools, surface scoring outputs that influence which accounts get priority or which markets get quoted. If your tool affects eligibility determinations, surcharging, or non-renewal recommendations, you're in filed-rate territory and your vendor probably hasn't told you.

The hedge: before deploying any tool that touches underwriting screening, appetite scoring, or risk prioritization, get confirmation in writing from the vendor describing exactly what the model does and does not determine. Then confirm with your E&O broker or the Nevada DOI directly before deployment. This is a one-conversation diligence step, not a legal engagement.

Burn 4: E&O exclusions and renewal questioning

ISO/Verisk Form CG 40 47 01 26, effective January 2026, introduced a generative AI exclusion in General Liability coverage. W.R. Berkley rolled out what their underwriters are calling an "absolute" AI exclusion in D&O, E&O, and Fiduciary lines. Swiss Re, which backs the Big "I" Westport E&O program, is heavily questioning AI usage at agency E&O renewals. This is the live market signal that matters more than any bulletin, because it affects your coverage before any regulator files an enforcement action.

Your next E&O renewal questionnaire will likely ask whether your agency uses AI tools, how they're governed, and what audit trail documentation exists. If you've disclosed nothing, budgeted nothing, and have no written AI governance policy, you're walking into a conversation without answers your underwriter already expects.

The hedge: disclose every AI tool currently in use to your E&O broker now, before your renewal cycle starts. Ask specifically about premium impact and what documentation the underwriter wants to see. Budget for some premium movement. And require audit-rights clauses and SOC 2 Type 2 attestation in every AI vendor contract you sign, so you have something to show the underwriter when they ask.

Burn 5: Vendor lock-in with no audit rights

You deploy a vertical AI tool that integrates with your AMS. The vendor gets acquired. Or shuts down. Or raises prices 40% at renewal because you've already migrated your workflows onto their platform. You go to export your data and discover your contract has no audit-rights clause and no specified data-return process. The vendor's terms say they can delete your data 30 days after contract termination.

NYDFS Circular Letter No. 7 (July 2024) explicitly requires audit-rights clauses in AI vendor contracts. Nevada follows the same governance framework under Bulletin 24-001. The real risk isn't regulatory, though. It's operational: you've become dependent on a vendor whose continuity you can't independently verify.

The hedge: before signing any AI vendor contract, require three things in writing. First, a data portability clause specifying the export format and timeline if you cancel. Second, audit-rights language giving you the right to review or have reviewed the model's decision logic. Third, a SOC 2 Type 2 attestation or equivalent showing that the vendor's security controls have been independently tested. If a vendor won't include these terms, that's useful information before you integrate their system into your AMS workflows.

---

01 HIGH SEVERITY

Hallucinated coverage advice

AI-drafted client communication contains a material misstatement about policy terms. Your CSR reviews it quickly and sends it.

REGULATION NRS 686A.030, NRS 686A.040, NRS 686A.310; fines up to $500 per act (NRS 686A.183)

02 HIGH SEVERITY

PII leakage via unvetted AI tools

Client SSN, policy number, and loss history pasted into a consumer AI tool with no BAA. Data may be used in model training.

REGULATION NRS 603A.040, .210, .215; HIPAA penalties up to $2.13M/yr if PHI involved (HHS-OCR 2024)

03 MED SEVERITY

Nevada Bulletin 24-002 filed-rate trap

AI-assisted quoting or appetite-screening tool surfaces scoring outputs that affect eligibility or prioritization without prior DOI filing.

REGULATION Nevada Bulletin 24-002 (Aug 27 2024); NRS 686B.110 filed-rate authority

04 HIGH SEVERITY

E&O exclusions + renewal questioning

ISO/Verisk CG 40 47 01 26 (Jan 2026) and W.R. Berkley introduced AI exclusions. Swiss Re (Big "I" Westport program) questioning AI use at E&O renewals.

REGULATION ISO/Verisk Form CG 40 47 01 26; W.R. Berkley absolute AI exclusion in D&O/E&O/Fiduciary lines

05 MED SEVERITY

Vendor lock-in with no audit rights

Vertical tool integrates with your AMS. Vendor gets acquired or raises prices 40%. Contract has no data-portability clause and no audit-rights language.

REGULATION NYDFS Circular Letter No. 7 (Jul 2024); Nevada Bulletin 24-001 governance framework

Five burns. Three connect to Nevada regulation. Two connect to what your E&O underwriter will ask about at your next renewal.

Henderson Sidebar: Why the Local Math Is Different

If you've been sitting on AI deployment because the national numbers don't pencil in your margin model, Henderson's labor market is reason enough to reconsider the math. The national AI-in-insurance conversation assumes replacing a CSR is annoying but manageable. Locally, that assumption breaks.

The Las Vegas-Henderson MSA experienced 24.9% employment growth from 2020 to 2025 (nv.gov employment data). As of February 2026, local unemployment sat at 5.8%, above the national average, which sounds contradictory until you factor in that a fast-growing market draws workers who then compete across sectors rather than concentrating in any one industry. Nevada health insurance premiums are projected at 30 to 84% increases in the next budget cycle (DHHS/state budget documents), which compounds the benefits cost burden for small agencies trying to stay competitive as employers. A fully loaded Henderson CSR runs roughly $60,000 to $70,000 per year in wage plus benefits (BLS OEWS Las Vegas-Henderson MSA, May 2024, plus 30 to 35% benefits load per BLS ECEC norms). That number has been moving up, not staying flat.

The implication for the buy-versus-wait question: capacity recovery is worth more in Henderson than national averages suggest, because replacing the CSR you just lost to a larger competitor is harder and more expensive than in most markets. The math that looks marginal in a market with cheap and plentiful back-office labor looks more favorable when your next hire is competing against four other industries that all have "hiring now" signs on their front door. I'm not saying this tips the calculation to a

> clear yes. I'm saying if you've been using a flat "the ROI doesn't pencil" conclusion as a reason to wait, the local cost context is worth factoring in more seriously.
>
> Two local resources actually worth your time. Nevada DOI bulletins (doi.nv.gov) are the source of record for 24-001 and 24-002 cited in the burn list; check it before deploying any tool that touches rating or underwriting. The Nevada chapter of the Independent Insurance Agents and Brokers (Big "I") is where peer conversations about what other agencies are actually deploying happen, which is more useful than vendor demos for calibrating expectations.

Three Things This Quarter

Three 90-day actions for a 3-15 person Henderson P&C and commercial agency. None of these require a capital commitment before week seven. All three will tell you something useful about where you actually stand.

Action 1 (Week 1-2): Map your top three back-office workflows, and have the E&O broker conversation while you're at it.

The workflow piece first: not an estimate, a real count. Have each CSR and operations team member track their time for one week, split by the five workflow categories in this guide: intake, quoting prep, compliance paperwork, renewal correspondence, claims and underwriting follow-up. Be specific about where the handoff fails in each one. One workflow will account for more time than the others, and that's your starting point. Trying to automate all five simultaneously is how agencies end up with four tools running at 20% adoption each.

While the count is running, call your E&O broker. Three questions: how is your underwriter currently scoring AI use at renewal for independent agencies? What does the renewal questionnaire ask about AI tools and governance, and what documentation do they want to see? What audit-trail format would satisfy them if they asked you to show how AI outputs are reviewed before reaching clients? Document the answers. Whatever you hear changes your Year 1 planning, your vendor contract requirements, and your disclosure posture before the next renewal cycle starts.

Action 2 (Week 3-6): Run one horizontal copilot trial for renewal-letter drafting and certificate-of-insurance summaries.

Pick one. ChatGPT Business or Claude Team, Enterprise tier with a signed BAA before any client data is involved. Budget about $20 to $25 per seat per month. Two seats maximum for the trial. The specific task: renewal follow-up letters and COI summary drafts that a licensed team member reviews and approves before sending. Track actual hours saved per CSR weekly, not impressions. Kill the trial if adoption is below 60% of the target team members using it actively by week six. "We bought it" is not adoption. The kill-it-early rule matters because it forces a diagnosis: if adoption is low, figure out whether the tool is wrong, the workflow wasn't mapped correctly, or the team needs a different onboarding approach. That diagnosis is worth more than six more months of hoping adoption improves.

Action 3 (Week 7-12): Decide who runs your AI ops, and write the role down.

This is the action that determines whether the first two actions compound or evaporate. By the end of the trial in Action 2, you'll have evidence about adoption, and adoption will tell you something specific: does this team have someone who can own this work and run it, or does that person need to come from outside? The math section of this guide makes the case. The tools alone don't move the KPI. The person who runs them does.

You have three realistic options for a 3-15 person agency. Option one: reposition an existing operations or service-team lead. They take 30 to 50% of their time on this, with a clear mandate and budget for tools. This is the right choice if you already have someone curious enough that they're running ChatGPT in the evenings without being asked. Give them the title before someone else hires them away. Option two: bring in a fractional AI ops lead at 5-15 hours per week, $40,000 to $80,000 per year. This is the right choice if you don't have the internal candidate and aren't yet at the scale that justifies a full-time hire. The fractional path is the most common starting point for a 5-10 person agency that's serious about this. Option three: full-time hire. This makes sense at 25-plus employees with multiple service teams and a hand-off problem between them, not before.

Whichever you pick, write the role down. One page. Reports to whom, owns which workflows, owns which KPIs, has authority to choose and cancel tools, has a budget. Without that page, the role doesn't exist; you have a tool and a hope.

These three actions together cost less than $2,000 out of pocket and about 20 hours of management time. The workflow mapping tells you where to start and what your E&O exposure looks like. The trial tells you whether your team can actually adopt a tool. The role decision tells you whether any of this compounds into a 24-month advantage or fades back into shelfware. None of them commit you to a full deployment before you know what you're buying.

If your agency has more than 25 people

If your agency has more than 25 employees and three or more service team leads, Action 3 above lands differently. At that scale, the orchestration problem is mostly hand-offs between teams, not tool selection, and the AI ops role usually needs to be full-time or fractional with senior weight behind it. I do separate Tier-2 work for agencies in that band: a fractional AI ops lead on retainer with named workflows, named KPIs, and weekly delivery. For a 25-person commercial P&C agency that typically means owning the renewal-review workflow end to end, setting a KPI like mean time from renewal notice to requote completed, and moving that number weekly instead of quarterly. See the About section for how that engagement works.

About & Ways to work together

About Piers

I'm Piers Rollinson. I run DomeWorks, a Henderson-based independent advisory practice helping small firms across professional services figure out where AI actually fits and where it doesn't.

Before DomeWorks, I spent 15 years building products and leading engineering teams. Most recently I was Director of Engineering at Mudflap (Series B fintech serving 500K+ truckers) where I led the AI Platform and Voice AI teams and redesigned the engineering organization around AI and agentic coding. Before that I was an Engineering Manager at DoorDash and Square, and Head of Engineering at Zesty (acquired by Square in 2018). I've been building software since 2009 and managing engineers since 2017.

I'm independent. I don't resell AI tools. I'm not partnered with the vendors named in any of my guides. I get paid by the firms I work with, not by tool referrals. That's the only way the honest-broker positioning works.

Ways to work together

In order from lowest commitment to highest.

01 **Free**

Free 45-minute walkthrough

I'll walk through your top three highest-volume workflows against the framework you just read, show you which one I'd start with first, and answer questions about specific tools. No slides, no pitch. Henderson or Las Vegas in person, or video.

Book the walkthrough →

02

$999

Tier-1 audit

Four-phase deliverable for firm owners who want a written diagnosis. 45-minute discovery call, AI analysis of your specific workflows, custom report identifying 5-7 opportunities ranked by impact and effort, 30-minute walkthrough call. 5+ hours-per-week recovered-capacity guarantee or full refund.

Book a discovery call →

03

2 weeks

Tier-1.5 measurement engagement

If your bigger problem is that you're not measuring AI ROI at all, you're not alone. 96% of small firms aren't. Two-week engagement to set up the measurement system that turns recovered hours into billable hours instead of earlier evenings. Pre-baseline, per-workflow attribution, conversion-to-billable tracking.

Ask on the walkthrough →

MID-MARKET

If your firm is bigger than 25 people

The audit is calibrated to firms with 5-25 back-office staff. If you're larger and looking at AI as an org-wide capability, the fractional CAIO engagement is a different conversation. Strategy work, governance, ROI measurement at the program level.

Read about the fractional CAIO offer →

Stay in touch

If you want monthly notes from me on what's actually working in the field, drop your email below. No drip campaign. One email a month, written to the same standard as this guide.

To subscribe from the PDF: email piers@domeworks.tech with "Monthly notes" in the subject.

Sources

Sources marked [vendor claim] are vendor-published. Where independent evidence exists, it is identified by publication name. Where it does not, the guide says so.

Problem Shape

- BLS OEWS May 2024 -- Insurance Claims and Policy Processing Clerks (SOC 43-9041), national mean and Nevada MSA data. bls.gov/regions/west/news-release/occupationalemploymentandwages_lasvegas.htm and bls.gov/oes/2023/may/oes439041.htm
- BLS OEWS Las Vegas-Henderson-North Las Vegas MSA -- office and administrative occupations. bls.gov/oes/2023/may/oes_29820.htm
- 2024 Big "I" Future One Agency Universe Study (n=1,269). insurancejournal.com/news/national/2024/09/27/794611.htm
- McKinsey "Insurance productivity 2030." mckinsey.com/industries/financial-services/our-insights/insurance-productivity-2030-reimagining-the-insurer-for-the-future
- Vertafore 2024 Insurance Agency Workforce Report. (URL pending verification.)
- Reagan Consulting Q1 2024 GPS (8.4% organic growth). theinsurer.com/ti/news/reagan-consulting-brokers-ended-2024-with-9-organic-growth-2025-02-20/
- IA Magazine "Half of Independent Insurance Agency Employees Feel Burned Out." iamagazine.com

- J.D. Power 2024 U.S. Independent Agent Satisfaction Study (n=4,918). (URL pending verification.)
- Bain 2025 P&C claims productivity study (20-25% loss-adjusting expense reduction; up to 50% on pilot tasks; up to one-third of underwriter time reclaimable). bain.com/insights/100-billion-dollar-opportunity-for-generative-ai-in-p-and-c-claims-handling/
- Nevada 2025 Insurance Market Report, Nevada Division of Insurance. (URL pending verification.)

Workflow Inventory

- Infrrd/Regure ACORD form benchmarks (8-12 min, 20% error rate). infrrd.ai/blog/acord-140-form-automation-guide; getregure.com/use-cases/reduce-manual-data-entry/ [vendor claim]
- Canopy Connect Elevate pricing and integration data. canopyconnect.com [vendor claim]
- Applied Epic Quotes Commercial Lines / Tarmika (35+ markets). reinsurancene.ws/applied-systems-unveils-ai-advancements-aimed-at-revolutionising-agency-workflows/ [vendor claim]
- EZLynx Rating Engine (330+ carriers). EZLynx product pages [vendor claim]
- IVANS "State of Agency-Carrier Connectivity" (73% lose opportunities, 83% would write more). (URL pending verification) [vendor claim]
- Gibson Agency quote (up to one hour, 15+ screens per renewal). Quandri press release; quandri.io/us/newsroom/quandri-unveils-ai-powered-requoting-in-us
- Indio (14,000+ form library, 12-minute signed application). Indio brochure [vendor claim]
- iPipeline iGO documentation. [vendor claim]

- RegEd annuities training (200,000 producers, 80%+ carriers). RegEd materials [vendor claim]
- Indemn compliance audit trail claim. indemn.ai [vendor claim]
- Nevada DOI annuity suitability press release (December 2024). doi.nv.gov
- Applied 2024 Voice of Customer survey (renewals as #1 inefficient task). [vendor claim]
- Quandri Vienneau case study (6,000 hours/year saved). quandri.io/us/customer-story/vienneau-insurance-ai-powered-personal-lines-renewals [vendor claim]
- AgencyZoom (5,000 agencies, 35,000 users). vertafore.com/resources/press-releases/vertafore-acquires-agencyzoom [vendor claim]
- Appulate LossRunner (40+ minutes saved per submission, free for agents). Appulate agent pages [vendor claim]
- Roots / Eastern Alliance Insurance Group (2,700 hours saved). iireporter.com/eastern-alliance-insurance-group-adopts-roots-automations-digital-coworker/ (third-party reported)
- Certificial COI tracking (3-4 hours/person/day). certificial.com [vendor claim]
- Certificial-Zurich integration (50%+ faster complex renewal cert packages). businesswire.com/news/home/20250128270839/en/ (third-party reported, January 2025)
- HawkSoft pricing (about $94/user/month). selecthub.com/p/insurance-agency-management-systems/hawksoft/

Vendor Landscape

- OpenAI ChatGPT Business pricing (as of April 2, 2026). chatgpt.com/pricing

- Microsoft 365 Copilot Business pricing. microsoft.com/en-us/microsoft-365-copilot/pricing/enterprise
- Anthropic Claude Team pricing. claude.com/pricing
- Google Workspace Gemini Business ($14-22/user/month). Google Workspace pricing pages
- NowCerts Momentum AMP pricing (Essentials $99, Professional $169, Business $349). nowcerts.com/pricing
- Quandri Series A+ ($12M, July 2025). businesswire.com/news/home/20250716833650/
- Canopy Connect Series A ($6.5M, 2023). (URL pending verification)
- Roots Series B ($22.2M, September 2024). crunchbase.com/organization/roots-automation
- Sixfold Series B ($30M, January 2026). sixfold.ai/content/post/series-b-ai-underwriter
- Cara Seed ($8M, April 2026). fintech.global/2026/04/01/cara-raises-8m-seed-to-build-ai-platform-for-brokers
- Agentech Seed ($3M, October 2024). businesswire.com/news/home/20241022350833/
- Indemn pre-seed ($1.9M, February 2024). crunchbase.com/organization/indemn
- Applied Systems acquisition of Planck (July 2024). BusinessWire announcement. (URL pending verification)
- Applied Systems acquisition of Cytora (September 2025). globenewswire.com/news-release/2025/10/07/3162837
- Vertafore acquisition of AgencyZoom (November 2021). vertafore.com/resources/press-releases/vertafore-acquires-agencyzoom
- Vertafore acquisition by Roper Technologies ($5.35B, 2020). businesswire.com/news/home/20200813005244/en/Vertafore-to-be-Acquired-by-Roper-Technologies

- Semsee Series B (2022) + iBynd acquisition (March 2026). businesswire.com/news/home/20260407446762/en/ (URL pending verification)
- Instabase Series D ($100M, January 2025, down round). techcrunch.com/2025/01/17/instabase-raises-100m
- Hyperscience Series E ($100M, late 2021). hyperscience.ai/newsroom
- Gallagher Re Q2 2024 insurtech funding (AI firms = 33% of total). (URL pending verification)
- HawkSoft API-partner model (~68 partners as of April 2026). hawksoft.com/about/partners

---

What Actually Works (Evidence-Based)

- NBER "Generative AI at Work" (14% productivity gain, 5,179 customer support agents). (URL pending verification)
- Science/MIT "Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence" (40% faster, 18% higher quality). (URL pending verification)
- INFORMS/MIT "Navigating the Jagged Technological Frontier" (758 consultants, 25% faster, 40% higher quality; "jagged frontier" finding). (URL pending verification)
- McKinsey "State of AI" 8th Annual (2025). mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai
- Bain P&C claims (20-25% loss-adjusting expense reduction; up to 50% on pilot tasks; ~4% of total claims expense scaled). bain.com/insights/100-billion-dollar-opportunity-for-generative-ai-in-p-and-c-claims-handling/
- Deloitte 2024 survey (200 US insurance executives, 76% deployed gen AI). deloitte.com/us/en/insights/industry/financial-services/scaling-gen-ai-insurance.html

- EY Responsible AI Pulse (August-September 2025, 975 C-suite leaders, 99% reported AI-related losses, 64% over $1M). ey.com/en_us/insights/insurance/gen-ai-in-insurance-key-survey-findings
- Celent March/May 2025 (44% active deployments, LLM hallucination "actively hazardous" for underwriting). celent.com/en/insights/what-the-evolution-of-gen-ai-in-underwriting-says-about-the-future-of-gen-ai-in-insurance
- Big "I" Agents Council for Technology (ACT) Research 2020 ("Accelerating the Quoting Process"). (URL pending verification.)
- Robotics and Automation News March 2022 (RPA COI: 15 min to 2 min per COI). (URL pending verification.)
- Journal of Financial Services Research 2023 (chatbot client satisfaction study). (URL pending verification.)
- IVANS download benchmark (76% of agents save 2+ hours/day via download). IVANS connectivity report [vendor claim]. (URL pending verification.)

What Gets People Burned

- NAIC Model Bulletin on Use of AI Systems by Insurers (adopted December 4, 2023). content.naic.org/sites/default/files/inline-files/2023-12-4%20Model%20Bulletin_Adopted_0.pdf
- Nevada Bulletin 24-001 (February 23, 2024). doi.nv.gov/News_Notices/Bulletins/Bulletin_24-001/
- Nevada Bulletin 24-002 (August 27, 2024). doi.nv.gov/uploadedFiles/doinvgov/Content/News_Notices/Bulletins/24-002.pdf
- NRS 603A (data security for insurance licensees). leg.state.nv.us/nrs/nrs-603a.html

- NRS 686A (unfair practices, false/misleading statements, unfair claims practices). leg.state.nv.us/nrs/NRS-686A.html
- NRS 686B.110 (filed-rate authority). Nevada Revised Statutes
- NYDFS Circular Letter No. 7 (July 11, 2024) -- audit-rights requirement. dfs.ny.gov/industry-guidance/circular-letters/cl2024-07
- ISO/Verisk Form CG 40 47 01 26 (effective January 2026). independentagent.com/vu_resource/verisk-to-roll-out-new-general-liability-exclusions-for-generative-ai-exposures/
- W.R. Berkley AI exclusion. zellelaw.com/AI_Update_The_Growing_Trend_of_AI-Related_Insurance_Policy_Exclusions
- Holland and Knight analysis confirming no Nevada DOI enforcement through Q1 2026 (May 2025). hklaw.com/en/insights/publications/2025/05/the-implications-and-scope-of-the-naic-model-bulletin
- HHS-OCR HIPAA tier penalties (up to $2.13M/year). HHS regulatory guidance
- Huskey v. State Farm (N.D. Ill., 2022). sanfordheisler.com/case/discrimination-harassment/state-farm-algorithm-bias-lawsuit/
- Kelly v. State Farm (M.D. Ala., October 2025). insurancebusinessmag.com/us/news/claims

---

The Math

- BLS Nevada wage data (Las Vegas-Henderson MSA, office and administrative occupations). bls.gov/regions/west/news-release/occupationalemploymentandwages_lasvegas.htm

- Henderson CSR salary (about $21.67/hour, April 2026). indeed.com Insurance CSR listings
- All major LLM pricing pages as of April 2, 2026: chatgpt.com/pricing, claude.com/pricing, microsoft.com/en-us/microsoft-365-copilot/pricing/enterprise
- SaaS implementation cost estimates: uncited, based on comparable SMB SaaS deployments. No independent benchmark exists for insurance AI implementation costs at agency scale.

What's Changing

- NAIC AI Systems Evaluation Tool (12-state pilot through September 2026). content.naic.org/committees/h/big-data-artificial-intelligence-wg
- Hawaii adoption of NAIC Model Bulletin (December 2025). (URL pending verification.)
- Federal Executive Order on state AI authority (December 16, 2025). content.naic.org/article/statement-national-association-insurance-commissioners-naic-ai-executive-order
- EU AI Act enforcement timeline (August 2, 2026). artificialintelligenceact.eu/implementation-timeline/
- Applied Systems Hellman and Friedman ownership (since 2014, IPO speculation). hf.com/portfolio/applied-systems-inc/
- Vertafore Roper Technologies 2025 10-K filing. sec.gov Roper Technologies filings
- Zywave consolidation activity (Clearlake/Aurora). (URL pending verification.)

DomeWorks. Henderson, NV. domeworks.tech