

A Faster, Smarter Onboarding System for Dave Creek Media

Plus one new AI product your insurance clients will pay for every month

Prepared by Marjohn Robillo

GHL Systems Architect | AI Automation Specialist

www.marjohnrobillo.me

Executive Summary - the whole proposal in one minute

Your onboarding system is already built on the right idea: AI generates, automation orchestrates, humans refine. Roughly 80% of it should stay exactly as it is. This proposal does two things: it gives you one new product you can sell across your insurance clients (Exercise 1), and it identifies the three highest-leverage upgrades to onboarding, ranked by impact against effort (Exercise 2).

A note on method: I treated Exercise 1 as a live demonstration of the research standard I propose in Exercise 2. Every claim I make about American Safeguard Insurance came from actually researching them - so you can judge the system by its output.

Exercise 1 in one line

The Instant Policy Review Engine: prospects upload their declarations page, AI produces a branded coverage review in minutes, and - the part that matters - the system captures their renewal date, the single most valuable data point in insurance marketing. Built entirely on your existing stack. Productized across every insurance client in your portfolio.

Exercise 2 - keep 80%, upgrade these three

Upgrade	What it fixes	Impact	Effort
1. Structured intake + client confirmation	Garbage-in is the root cause of "AI filler" output	High	Low
2. Analysis-first generation order	Three deliverables that contradict each other	High	Low
3. Two early human gates + learning loop	Expensive late-stage rewrites; a system that never improves	High	Medium

Exercise 1 - The Instant Policy Review Engine

Where the idea comes from

American Safeguard's own homepage says it: "Have a current policy? Send your declarations page and we'll review it with you." That is a manual, slow, agent-hours-heavy process today - and it is also their best offer. A dec page review is genuinely useful to the prospect, requires no commitment, and positions ASI as the helpful advisor their brand promises ("we listen first, explain your options clearly... without added pressure"). The product below takes the thing they already do best and makes it instant, scalable, and measurable - without changing their conservative, relationship-first brand voice.

What it is




Figure 1 - The Instant Policy Review Engine, end to end on the existing DCM stack

A GHL funnel where a prospect uploads their current declarations page. Claude parses the PDF and extracts carrier, coverage lines, limits, deductibles, premium - and the renewal date. It cross-references against a coverage rules library (built once with ASI's agents, per line of business) and generates a branded, plain-English Coverage Review report. GHL automation then books the prospect with a licensed agent and drops them into a renewal-timed nurture sequence.

Why the renewal date changes everything

In insurance, prospects are only truly buyable in the 30-45 day window before their policy renews. Most agencies have no idea when that window is. Every dec page uploaded into this system tells you exactly when this prospect becomes winnable - so the funnel doesn't just generate a lead today, it schedules a perfectly timed sales conversation months from now, automatically. The database compounds in value every month it runs.

How AI powers it - on the stack you already pay for

Step	Tool	What it does
Landing page + upload	GoHighLevel	Branded funnel per client, form with file upload, CRM record created
Document parsing	Claude API (Sonnet)	Extracts structured policy data from the dec page PDF

Step	Tool	What it does
Gap analysis	Claude + rules library	Compares coverage against per-line rules the agents define once
Report generation	Claude Design	On-brand Coverage Review PDF from the client's brand system
Orchestration	Make.com	Upload triggers parse, parse triggers report, report triggers review task
Agent approval	ClickUp	Report lands as a review task - nothing is sent until an agent approves
Booking + nurture	GoHighLevel	Calendar booking, then renewal-window email/SMS sequence

How it scales across your insurance portfolio

Nothing in the engine is ASI-specific except three swappable parts: the brand kit (already per-client in Canva/Claude Design), the rules library (a one-workshop setup per client), and the funnel skin (a GHL snapshot - deploy in under a day). That makes this a productized offering: a setup fee plus a monthly license per insurance client.

Insurance agencies are ideal buyers - renewal-driven revenue, high customer lifetime value, and almost none of them have anything like this.

Exercise 2 - Upgrading the Onboarding System

Answered in the order of your seven questions. The redesigned pipeline below is the map; the sections that follow are the reasoning.




Figure 2 - Redesigned pipeline. Red boxes are the only two places a human touches the process before delivery.

1. Process improvements - fix the intake first

The current flow treats intake as a formality and research as the heavy lifting. It's backwards. Every downstream prompt inherits the quality of the intake, and today that intake is a client-completed form (inconsistent), sales notes (unstructured), and assets (variable). Three changes:

- **Record the sales call, not just the notes.** The transcript contains the client's own words - pain points, promises made, how they describe their customers. AI structures it; nothing new for the sales team to write.
- **One canonical Client Brain record.** AI merges form + transcript + contract + assets into a single structured record in ClickUp. Every prompt downstream pulls from this one source - not from four documents that disagree.
- **A five-minute "What We Heard" confirmation.** Before any research runs, the client confirms a one-page summary of their record. This catches wrong assumptions on day zero, and it's a trust-building touchpoint inside the exact window where you say trust is built or lost.

2. AI architecture - a staged pipeline, not a mega-prompt

One giant prompt produces confident mush. A staged pipeline with structured outputs at each stage produces checkable work. Model routing, with reasons:

Stage	Model	Why this model
Research collection	Gemini Pro (Flash-Lite for lookups)	Search-grounded, cheap at volume - collection is a breadth job
Strategy synthesis	Claude Opus - one call	The single judgment-heavy step; pays for itself in draft quality
Deliverable drafting	Claude Sonnet, batch mode	Structured drafting from an approved analysis; batch = 50% savings
Visual production	Claude Design	Reads the client brand system; outputs the polished blueprint

Consistency across industries comes from an industry module library: per-vertical prompt inserts (benchmarks, channel norms, vocabulary, a gold-standard past example as a few-shot reference). Every stage outputs structured JSON against a fixed schema, so stage 4 can never receive freeform text that drifts.

3. Research depth - evidence-linked, or it doesn't ship

- **Review mining is the sleeper source.** The client's and competitors' Google reviews are voice-of-customer gold - the exact words real buyers use, what they praise, what they complain about. No generic AI summary produces this; it feeds directly into messaging.
- **Facebook Ads Library for what competitors actually spend on.** Not what their website says - what they're running right now, how long ads have been live (long-running = working).
- **Search Atlas for the keyword gap** between the client and each named competitor - real volumes, not AI guesses.
- **The rule that prevents surface-level output:** every claim in the market analysis carries a source link or data point. If AI can't cite it, it gets flagged as an assumption - visibly. This one rule is most of the hallucination answer.

4. Deliverable design - one source of truth, three shapes

The market analysis generates first and becomes the single source of truth. The blueprint and creative brief are derived from it - which is why the three documents can never contradict each other.

Deliverable	Built for	Skeleton
Internal market analysis	The team's reference	Competitive landscape (evidence-linked) - positioning opportunities - audience and voice-of-customer findings - keyword gaps - ad benchmarks - strategic rationale - flagged assumptions
Client blueprint	The client's first impression	"What we learned about you" (mirrors their own words) - where you stand today - the strategy and why - channel plan - 90-day roadmap with milestones - what happens next week
Creative brief	Open it, start producing	Brand voice do/don'ts with examples - visual notes - content pillars with ratios - messaging hierarchy - campaign themes - first two weeks of content, listed by piece

5. Automation design - Make.com wiring and the two handoffs

Make.com stays the orchestrator. Contract signed in GHL fires the scenario chain: intake structuring, client confirmation step, then parallel research routes (competitors, ads, reviews, keywords, knowledge base) that converge into the synthesis call. Two - and only two - human handoffs:

- **Gate 1, after research (10 minutes).** A human skims the research summary before any generation. Catching a wrong competitor or a bad assumption here costs minutes; catching it after three polished deliverables exist costs hours. Cheap gate early beats expensive gate late.
- **Gate 2, after drafting (approve/annotate).** Reviewers approve or annotate - they don't rewrite. Annotations route back through a revision call. This keeps review fast and, critically, produces the edit log that powers the learning loop below.

Error handling that scaling will demand: every scenario writes status to the ClickUp record, failed routes retry then alert, and no silent failures - a research route that dies quietly today produces a hollow deliverable tomorrow.

6. Quality assurance - guardrails that don't slow you down

- **Source-or-flag rule** (section 3): claims carry evidence or are marked as assumptions. Hallucination becomes visible instead of hidden.
- **Schema validation between stages:** structured JSON contracts mean a malformed output stops the line at the cheap stage, not in a client document.
- **Fact table check:** names, locations, URLs, and numbers in deliverables are validated against the Client Brain record - the most embarrassing errors are the most mechanical to catch.
- **Brand lint:** a fast Sonnet pass that checks drafts against the client's voice rules before human review - reviewers spend their attention on strategy, not typos.

7. Scaling from 5 to 20 clients a month - design for it now

Compute scales trivially. Two things actually break: human review hours and the Make ops budget. The review answer is the approve/annotate workflow above plus the industry module library - reviewers confirm rather than create. The ops answer is measuring cost per onboarding from day one (ops consumed, tokens spent, review

minutes) so you know your real unit economics before you commit to volume. And one thing that turns scale from a threat into an asset:

THE LEARNING LOOP - why client #40's draft is better than client #4's




Figure 3 - Reviewer edits are logged and folded back into the prompt modules monthly.

Most AI systems produce the same quality at client 100 as at client 1. This one doesn't: every reviewer annotation is data about where the prompts fall short, and it feeds a monthly refinement cycle per industry module. Client #40's first draft is measurably better than client #4's - which means review time per client falls exactly as volume rises. That crossover is what makes 20 clients a month profitable instead of painful.

Honest Gaps - what I'd test before promising

- **Claude Design automation limits.** I'd verify how far its API can be pushed inside a Make scenario for fully unattended blueprint generation - the fallback is a Canva template populated by structured output, which is proven but less elegant.
- **The 24-hour SLA needs a measured baseline.** I'd run three pilot onboardings and time every stage before committing to the number publicly. My estimate is that the pipeline lands well under 24 hours - but an estimate isn't a baseline.
- **Higgsfield has no role in onboarding - and I'm saying so.** Video production belongs in fulfillment, not in the signature-to-strategy window. Forcing every tool into every workflow is how stacks get slow.
- **The insurance product's revenue ceiling** depends on how many insurance and adjacent clients (financial services, legal) sit in your portfolio - I'd want that count before projecting revenue. The engine itself also adapts cleanly to any "upload a document, get a review" vertical: leases, contracts, utility bills.

Why I built the proposal this way

You said the window between signing and first results is where trust is built or lost. Every choice above serves that window: the intake confirmation builds trust on day zero, the analysis-first order makes the deliverables coherent, the early gate protects speed, and the learning loop means the system you buy gets better every month you run it. I'd be glad to walk through any section - or build the pilot.

Marjohn Robillo