REGESTA

FOR INDEPENDENT INSURANCE AGENCIES

The Work Between the Systems

How workflow architecture turns fragmented agency work into an AI-native insurance operation.

RECORD	RENEWAL PREPARATION	OPEN
MOVES THROUGH	INBOX / AMS / PORTAL / SHEET	6 HANDOFFS
CONTROL	LICENSED REVIEW REQUIRED	HUMAN

DECISION OWNER IDENTIFIED

BY **ANDREW CHERRY**

The Work Between the Systems

How workflow architecture turns fragmented agency work into an AI-native insurance operation.

Andrew Cherry · Regesta

> The premise Most agencies do not have a software problem. They have critical work spread across inboxes, the AMS, carrier portals, spreadsheets, and what a few experienced people carry in their heads.

A NOTE FROM ANDREW

This is not an AI book.

It is a book about the work AI is being asked to enter.

I keep seeing the same mistake. An agency feels the strain of growth, sees a new technology that promises capacity, and starts with the tool. The vendor gets a demo. A few licenses get handed out. One or two curious people find useful tricks. The rest of the team keeps working the same way.

The experiment looks like an AI problem. It is usually a workflow problem that was already there.

A renewal still moves through email, the AMS, carrier portals, documents, spreadsheets, and the memory of the account manager who knows where everything lives. A submission still depends on someone noticing what is missing. A service inbox still relies on judgment nobody has written down. AI lands on top of that and inherits the ambiguity.

> *If the work is unclear, automation does not create clarity. It creates speed without control.*

This field guide is for owners, presidents, and operators who want the upside of AI without pretending the agency is a clean system waiting for a model. It translates a broad body of agent and broker research into a practical position for independent commercial agencies: make the workflow visible, protect the decisions that must stay human, and only then decide what technology belongs in the future state.

That is the work I do through Regesta. One workflow at a time.

Andrew Cherry Workflow Architecture for Independent Insurance Agencies

CONTENTS

Start with the work.

I The workflow underneath the agency Why growth exposes work that was never truly designed. 06

II What the research is actually saying Four signals from the front line, and why the agency must keep learning after launch. 11

III Five rules for useful AI A practical architecture for integration, evidence, authority, continuity, and knowledge. 20

IV Three workflows, redesigned Renewals, submissions, and the service inbox before and after. 26

V The first move A diagnostic, a stress test, and a build order you can defend. 31

> How to use this book Read it straight through if you are deciding how AI fits into the agency. If one workflow is already causing pain, go to Part Four, find the closest example, then use the stress test in Part Five with one recent real case.

The agency is already automated.

It is automated through habits, workarounds, inbox rules, spreadsheets, and the memory of experienced people.

That system can run for years. It can even look efficient from the outside. Then the agency grows, a key person leaves, a workload spikes, or leadership adds a new tool. The hidden system becomes visible because it stops bending.

> *Your best people are not using the workflow. In many agencies, they are the workflow.*

The first job is not to replace them. It is to make what they know visible enough to improve, teach, and protect.

Once the workflow is visible, AI becomes a design choice. Before that, it is a bet.

I

The workflow underneath the agency

Growth does not create operational fragility. It reveals where the agency has been borrowing capacity from its best people.

Most workflows were never designed.

They accumulated around exceptions, carrier requirements, staffing changes, and whatever worked last time.

Ask leadership how a renewal works and you will hear a sequence. Ask the account manager and you will hear a set of decisions. Ask a newer employee and you will hear where they get stuck.

All three descriptions can be true. None is the workflow by itself.

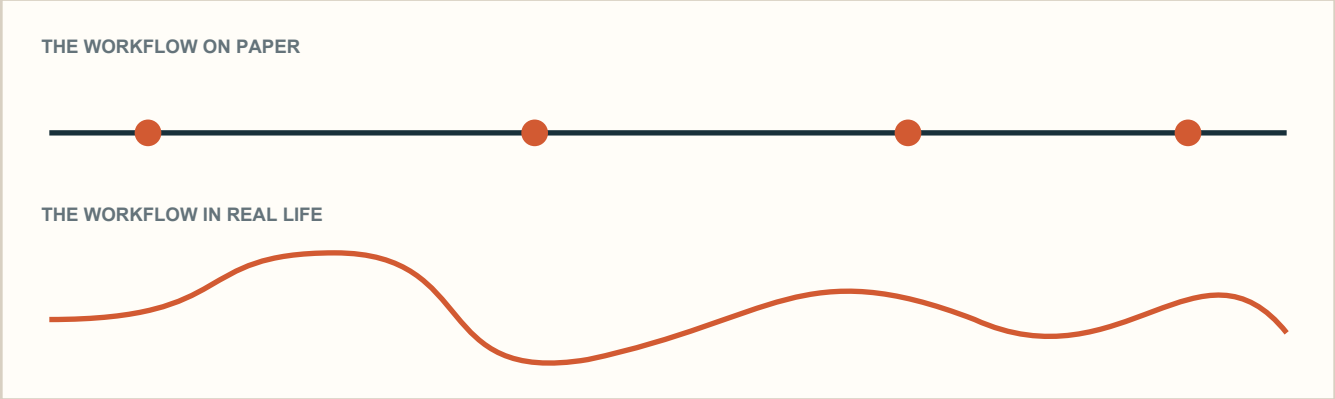

The gap between those two lines is where service capacity disappears. It is also where risk hides. If an exception is resolved only because one person recognizes it, that person is carrying a control the agency does not formally own.

Key person dependency is a workflow signal.

The most valuable person in a process is often the one who knows what the documented process leaves out.

They know which carrier portal behaves differently. They remember that a client always needs a certificate issued a certain way. They can look at a submission and sense that one fact will create a market problem before anyone else sees it.

If your best account manager disappeared for two weeks, which workflow would be hardest for the rest of the team to reproduce correctly?

What the answer reveals

- Unwritten routing logic. The right next step depends on context that lives in memory.
- Invisible quality control. Someone catches errors because they know what wrong feels like.
- Exception ownership. Difficult cases drift toward the same person regardless of formal role.
- Training debt. New people learn by interruption because the reasoning is not available when they need it.

AI can help capture, surface, and reuse parts of this knowledge. It cannot be allowed to convert hidden judgment into hidden automation. The first design task is to separate what can be made repeatable from what must remain an explicit human decision.

One task. Many systems. One human bridge.

Agents rarely complain that they lack software. They complain that the software does not behave like one system. The account manager becomes the integration layer, carrying information from one place to another and reconstructing context every time the work moves.

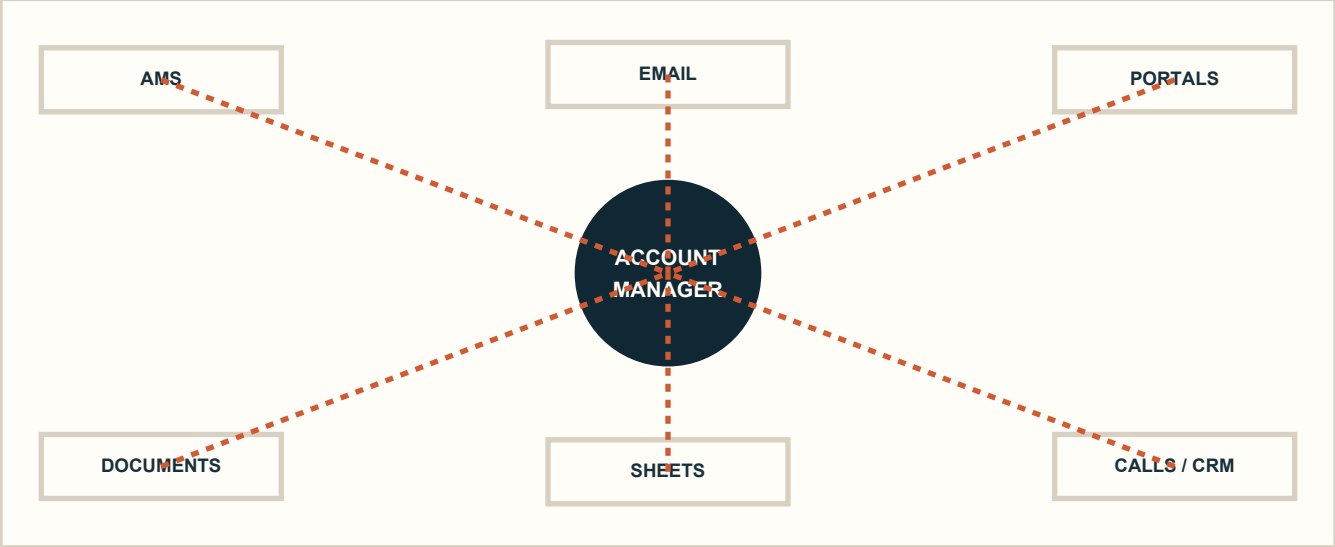

> *Every manual hop is a chance to wait, repeat, forget, or reinterpret.*

The objective is not a fantasy platform that replaces everything. It is a designed flow of context across the systems the agency already uses, with fewer transfers and clearer ownership.

Five questions before you buy anything.

1. What event starts the workflow? Not the department. Not the job title. The observable trigger.
2. What counts as finished? A draft is not complete if someone still has to chase approval or enter the same data elsewhere.
3. Where does the work wait or double back? These points usually matter more than the time spent typing.
4. Which decisions require licensed, fiduciary, or relationship judgment? Make authority explicit before adding assistance.
5. What evidence would prove the future state is better? Cycle time, touches, corrections, interruptions, backlog, or a risk control that now exists.

> The purchasing rule If leadership cannot answer these questions using one recent real case, it is too early to choose the tool.

That is not caution for its own sake. It is how the agency avoids paying a vendor to automate an assumption.

II

What the research is actually saying

The strongest signal is not that agents resist AI. It is that curiosity has outrun structure.

Useful research, translated honestly.

The agent and broker study behind the original report interviewed 16 participants across 13 states, multiple lines of business, and a wide range of experience.

That is qualitative research. It is valuable for finding patterns, language, and design constraints. It is not a census of the industry, and this book does not present it as Regesta's own field study.

What the study can show
Recurring friction, common behaviors, unmet needs, and the conditions under which agents trust or reject AI.

What it cannot prove alone
How frequently every behavior occurs across the entire independent agency market.

The four signals that matter

- AI is spreading without a consistent roadmap.
- People are finding real value, but most use remains shallow and personal.
- Agents want integration more than another layer of automation.
- Domain expertise determines whether the output can be trusted.

Regesta's contribution is the operational translation: what those signals mean when an independent agency is deciding what to redesign, what to protect, and what to build first.

Investment is ahead of operating readiness.

8%

of independent agents were reported to use AI daily.
Liberty Mutual, 2025

27%

of executives reported having a comprehensive AI strategy.
Gartner, 2026, as cited in the source report

66%

of organizations had not adopted financial guardrails around AI spend.
Gartner, 2026, as cited in the source report

40%

of enterprises were predicted to demote or decommission AI agents by 2027 after governance failures.
Gartner, 2026, as cited in the source report

Read the pattern, not just the numbers Low use, incomplete strategy, weak spending controls, and expected production failures all point to the same missing layer: operational design.

Curiosity is doing the work of strategy.

AI often enters the agency through the most curious employee, not through a defined operating decision.

That person experiments, finds a useful prompt, shares it with a colleague, and creates local value. Meanwhile, someone else is unsure whether they are allowed to use the same tool. Leadership may believe AI use is low because sanctioned adoption is low. The team may already be using consumer tools for business work.

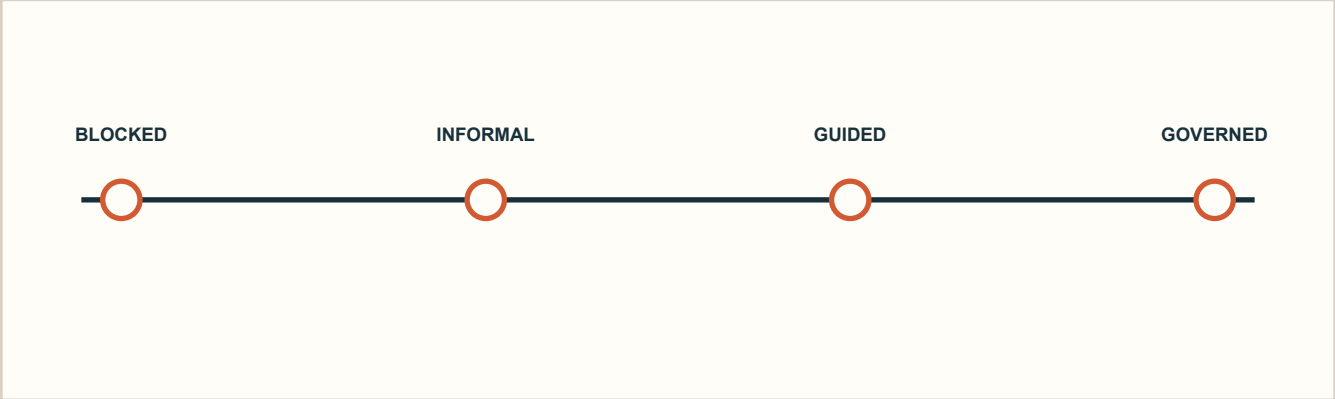

What an agency needs instead

A short, usable boundary: approved tools, prohibited data, allowed use cases, required review, and a named owner for exceptions. Governance should tell a capable employee how to proceed safely, not simply say yes or no.

Small wins can hide a large exposure.

Drafting emails, summarizing policy language, comparing documents, and preparing meeting notes can save real time. Those wins matter. They also create a dangerous illusion: because the output looks like a draft, the process around it can feel informal.

Faster writing, less blank-page work, quicker synthesis, lower cognitive load.

Client data in personal tools, uncertain retention, no source trail, inconsistent review.

> *Shadow AI is often a policy failure before it is an employee failure.*

The research cited in the source report found that 88 percent of employees with enterprise AI access also used personal AI tools for business tasks. The behavior is easy to condemn and important to understand. People route around friction when the unofficial tool works better than the official path.

The answer is not to reward the workaround. It is to design an approved path that is useful enough to compete with it.

Operating requirement Every AI-assisted task needs a data boundary, an accountable reviewer, and a clear definition of what may happen after the output is produced.

Agents do not want another destination.

They want the existing work to require fewer searches, transfers, and repeated entries.

Another chat window can be impressive in a demo and irrelevant in a workflow. If the agent must leave the AMS, locate the right documents, paste context into a tool, validate the response, then re-enter the result elsewhere, the agency has added a step while calling it automation.

ADDITION	INTEGRATION
• One more login	• Works where the task already lives
• Manual context gathering	• Pulls only the approved context
• Copy and paste between systems	• Returns structured work to the right record
• Output with no source trail	• Cites the underlying source
• Adoption depends on memory	• Approval is part of the flow

The practical test

Count the total touches before and after. If the new tool saves typing but adds searching, copying, checking, and transferring, the workflow may not be better.

> *The right AI should reduce the pile, not become the newest thing on it.*

Experience is the quality system.

The agents getting the most value from AI are often the people most capable of recognizing when it is wrong.

That creates a trap. The tool looks easiest to deploy for routine work and newer employees. Yet a newer employee may be least equipped to detect a plausible mistake, a missing exposure, or a recommendation that crosses an authority boundary.

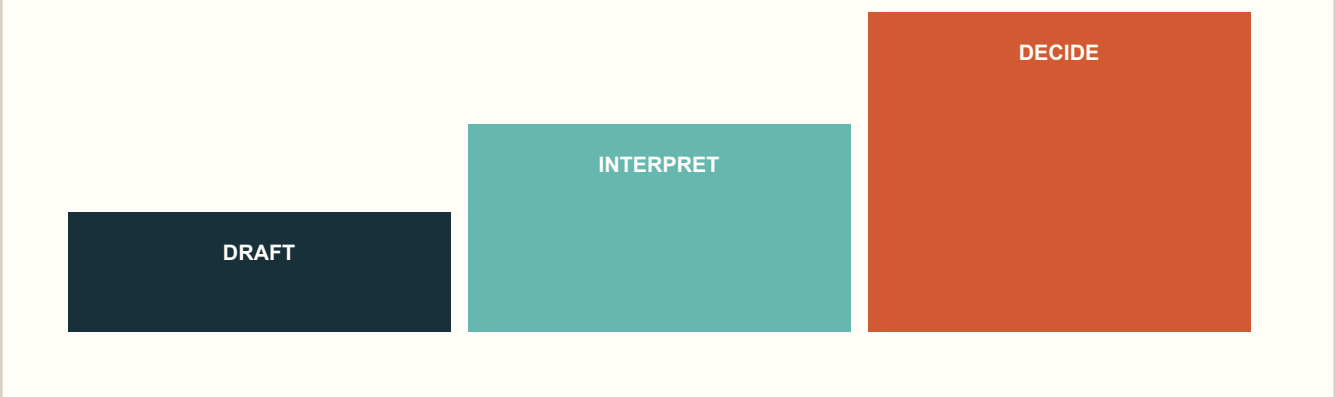

The design opportunity is knowledge transfer, not silent substitution. Cite the source. Show the reasoning trail. Capture corrections. Surface the questions an experienced person would ask. Keep the licensed decision visible and owned.

> The principle AI can help a newer employee reach better questions. It should not disguise the absence of experience as certainty.

AI has no go-live date.

An agency can finish an installation. It cannot finish becoming better at using intelligence inside the work.

A normal software launch answers familiar questions. Is the system configured? Does the integration work? Can people log in? Who handles support? Those questions matter. They simply do not finish the job.

AI changes through use. A producer finds a better way to prepare for a call. An account manager catches a confident mistake. A difficult renewal exposes an approval boundary nobody defined. A model changes. A carrier changes its appetite. A regulator clarifies an expectation. Each event changes what good use looks like.

If attention ends at deployment, the agency freezes its operating design while the technology, the risk, and the people keep moving. The tool remains available, but the way it is used becomes private, uneven, and harder to govern.

> *The software may launch once. The agency has to keep learning.*

An AI-native agency is therefore not defined by the number of tools it has installed. It is defined by whether it can absorb what people learn, revise the workflow, strengthen the control, and turn one useful experiment into a repeatable way of working.

Build the loops that survive the launch.

AI becomes durable when the agency has a way to keep the work, the judgment, the learning, and the value connected.

1. The work loop
A useful method moves from private experiment into the actual workflow, with an owner, a trigger, a finish line, and fewer manual transfers.

2. The judgment loop
Sources, uncertainty, approval, and escalation remain visible so assistance never quietly becomes authority.

3. The learning loop
Corrections and exceptions are captured, reviewed, and used to improve the process instead of disappearing inside individual chat histories.

4. The value loop
The agency measures whether the change reduces waiting, touches, corrections, backlog, or decision time in the whole workflow.

Ownership follows the outcome

The business leader owns the result. Operations stewards the workflow. Technology maintains the foundation. Risk defines the boundary. The people doing the work supply the evidence that tells everyone what must change next.

> The test Six months after launch, can the agency explain what it learned, what it changed, who approved it, and what improved? If not, it installed a tool. It did not build a new way of operating.

Five rules for useful AI

Useful AI is built inside the workflow and improved through use. Good design makes the system quieter, the evidence clearer, and human authority harder to miss.

Do not add a destination.

Put assistance inside the path the work already follows, or remove enough existing steps to justify a new path.

Integration does not mean connecting everything to everything. It means the right context appears at the right moment, in the system where the accountable person is already working.

Bring context forward
Assemble the client, policy, task, and communication history needed for this step.

Return work cleanly
Write approved outputs back to the proper record without a second round of entry.

Count every touch
Measure searches, transfers, checks, and interruptions, not just keystrokes.

> *The goal is not more capability. It is less operational surface area.*

Design question

If this feature disappeared tomorrow, would the team miss the outcome or only the novelty?

Make the output inspectable.

Trust grows when a person can see where a statement came from and correct it without leaving the workflow.

An AI-generated renewal summary should point to the policy language or document section it used. A comparison should show the source values. A flagged gap should identify the evidence and the uncertainty.

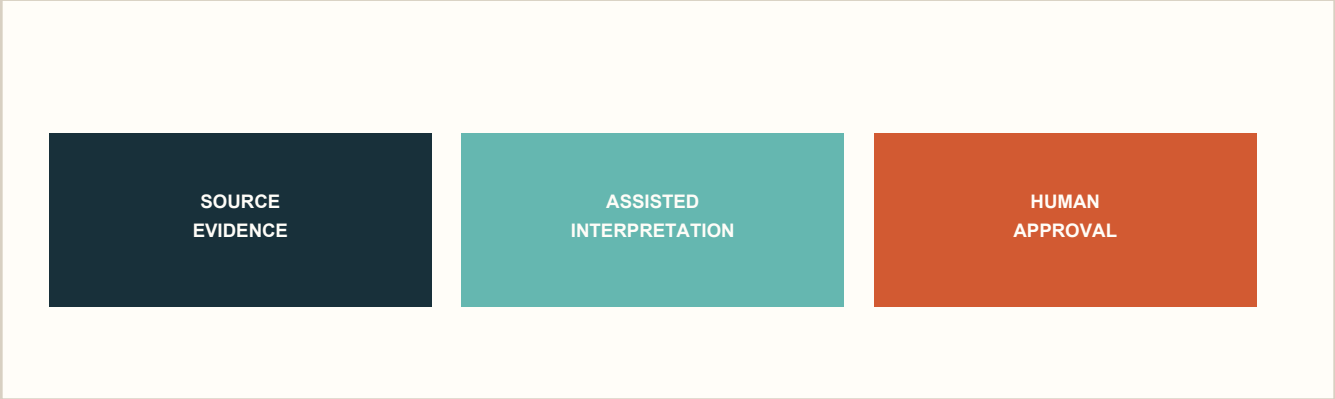

- Cite the source beside the claim.
- Show uncertainty where the source is incomplete.
- Capture the correction as part of the record.
- Never let fluent language masquerade as verified fact.

Assistance is not permission.

A system may prepare, organize, compare, and draft. That does not mean it owns the decision.

Every future state should distinguish three kinds of work.

Machine work
Retrieval, formatting, routing, structured comparison, and other bounded operations.

Assisted work
Drafting, synthesis, and recommendations that require validation against evidence.

Human authority
Coverage advice, exceptions, fiduciary choices, relationship decisions, and final approval.

A checkpoint is more than a button The reviewer needs the source, the proposed action, the relevant uncertainty, and enough time to make a real decision. An approval click without context is theater.

Human review should be designed as a meaningful operating step, not a legal label added after the architecture is finished.

> *Keep the person in control of the decision, not trapped in cleanup after it.*

Do not automate the moment and ignore the workflow.

Many tools help at one attractive point: write the email, summarize the call, compare the documents. The surrounding work still begins too early and ends too late.

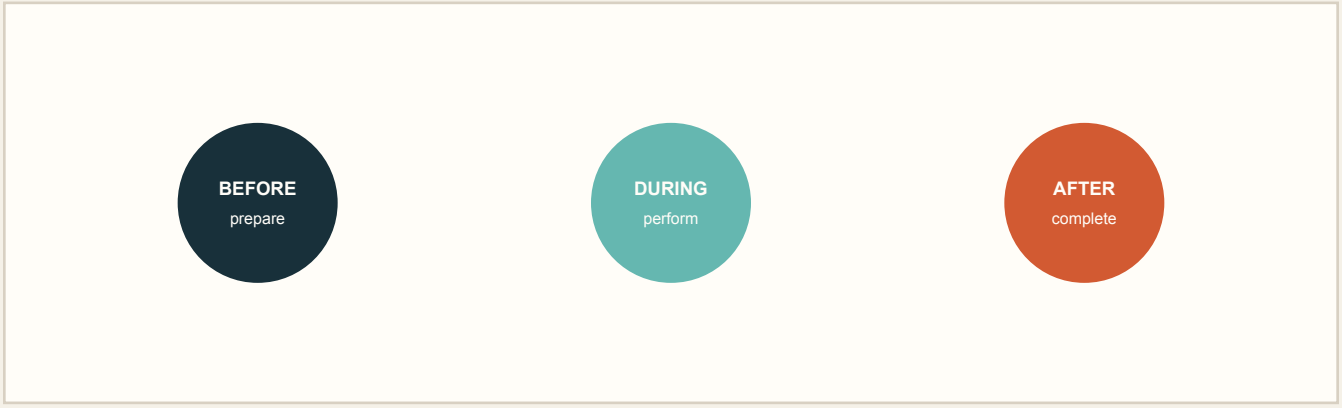

A renewal email is not the workflow. The workflow begins when the renewal enters the horizon and ends when the client decision, carrier action, documentation, and follow-up are complete.

Designing the loop prevents a common failure: saving five minutes at the visible task while leaving hours of searching, waiting, and reconciliation untouched.

> Measure the whole loop Trigger to completion. Not prompt to draft.

Use efficiency to fund knowledge transfer.

The long-term prize is not a faster email. It is an agency that can retain and teach the reasoning its best people use.

When an experienced account manager corrects a summary, rejects a market, or asks a question nobody else considered, that action contains operating knowledge. A well-designed system can make the reasoning easier to capture and reuse without pretending it can replace the person who supplied it.

Capture decisions
Record what was decided, the evidence considered, and why an exception was made.

Surface precedent
Bring similar prior cases forward when they are relevant, with their sources intact.

Teach the questions
Give newer staff the prompts an experienced person uses to investigate the work.

Protect the boundary
Do not turn historical judgment into automatic authority without review.

> *Efficiency saves the hour. Transferable judgment protects the agency.*

IV

Three workflows, redesigned

The value appears when context moves cleanly and the right person still owns the decision.

From document chase to decision preparation.

A renewal should arrive at the account team as a prepared decision, not a scavenger hunt.

AS IT OFTEN RUNS

- Dates and tasks live in several places.
- Documents are gathered manually.
- Changes are compared line by line.
- Open items are reconstructed from email.
- The experienced account manager catches what the process misses.

A DESIGNED FUTURE STATE

- The renewal horizon triggers the workflow.
- Approved context is assembled automatically.
- Changes and missing information are presented with sources.
- The account manager resolves exceptions and approves client communication.
- The decision and follow-up return to the record.

What AI may assist

Document extraction, change comparison, plain-language drafting, activity synthesis, and checklist preparation.

What stays human

Coverage interpretation, advice, market strategy, exceptions, client framing, and final approval.

> Evidence of improvement Fewer searches, fewer late surprises, lower preparation time, fewer corrections, and a visible approval trail.

From repeated entry to a controlled handoff.

The submission should be built once, validated once, and routed with the gaps made visible.

AS IT OFTEN RUNS

- Information is captured in notes, email, and forms.
- The same facts are re-entered.
- Missing details appear late.
- Carrier appetite lives in personal memory.
- Producers and service staff interrupt each other for status.

A DESIGNED FUTURE STATE

- One structured intake becomes the working record.
- Existing approved data is reused.
- Missing or conflicting information is flagged early.
- Market options are supported by evidence and reviewed by an experienced person.
- Status and ownership are visible throughout.

What AI may assist

Call capture, structured extraction, form preparation, document classification, completeness checks, and draft narratives.

What stays human

Risk interpretation, market selection, representations to carriers, exception handling, and final submission approval.

> Evidence of improvement Lower re-entry, fewer missing fields, fewer producer interruptions, faster ready-to-market time, and fewer carrier follow-ups.

From queue anxiety to visible ownership.

A shared inbox should not require constant vigilance to remain safe.

AS IT OFTEN RUNS

- People scan for urgency by instinct.
- Ownership is implied, not explicit.
- Requests are copied into other systems.
- Similar work is handled differently.
- Backlog becomes visible only when something is late.

A DESIGNED FUTURE STATE

- Requests are classified within approved boundaries.
- Ownership and service level are explicit.
- Relevant client context follows the request.
- Routine responses are prepared for review.
- Exceptions escalate to a named person with the evidence attached.

What AI may assist

Classification, summarization, routing suggestions, duplicate detection, draft replies, and structured task creation.

What stays human

Urgency overrides, coverage or fiduciary interpretation, sensitive client responses, exceptions, and final communication where authority is required.

> Evidence of improvement Lower time to ownership, fewer unassigned items, reduced backlog age, fewer duplicate touches, and a clearer escalation trail.

Make the path visible. Protect the decision.

A useful future state is not a software diagram. It is an operating agreement.

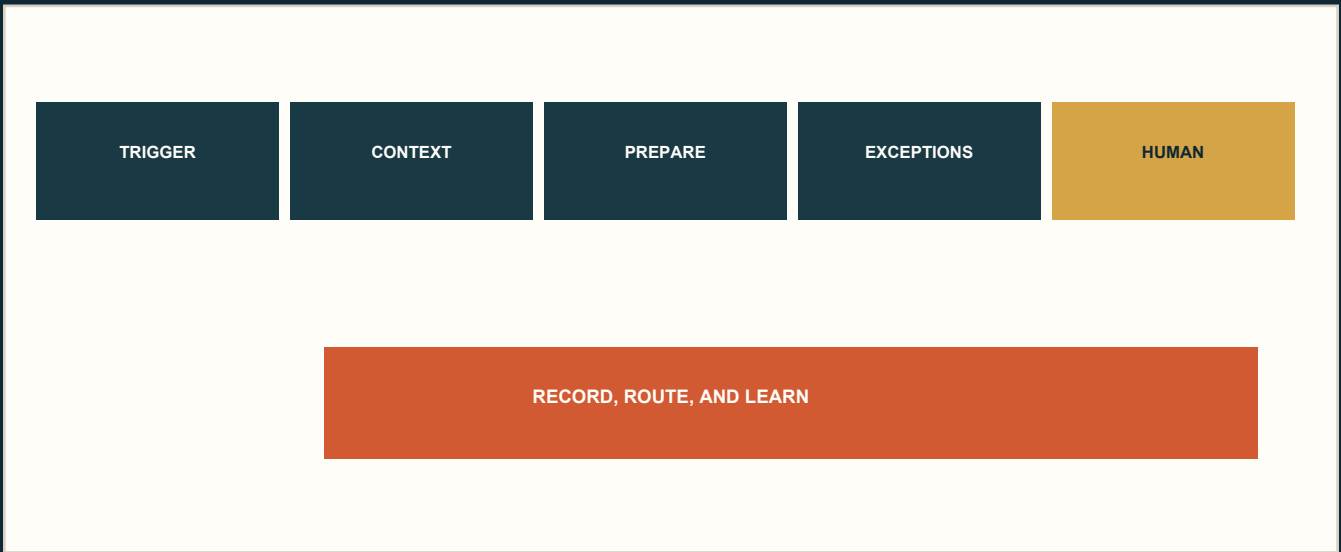

The loop matters. Corrections, exceptions, and approved decisions should improve the operating record. They should not disappear into a private chat history.

V

The first move

One workflow examined properly is more useful than an agency-wide transformation plan built on assumptions.

Diagnose one workflow before you build.

Regesta maps the work from trigger to completion using recent real cases, not the clean version people remember in a conference room.

The Independent Agency Workflow Bottleneck Diagnostic is deliberately narrow. It gives leadership enough evidence to decide what is actually broken, where human judgment must remain, and what to improve first.

Current state
Steps, owners, systems, handoffs, waiting, rework, exceptions, and failure points.

Protected judgment
Licensed, fiduciary, relationship, and exception decisions that stay human.

Future state
A practical workflow architecture based on the evidence, not on a preferred vendor.

Build order
Prioritized improvements, a proposed scope, a budget range, and a recommended first move.

What you own at the end A complete decision instrument that remains useful whether or not Regesta implements anything afterward.

Which workflow should you examine first?

Choose one workflow and check every statement that is true using a recent real case.

- [ ] The same information was entered or checked in more than one place.
- [ ] The work waited because ownership or the next step was unclear.
- [ ] Someone had to search across email, the AMS, portals, or documents to rebuild context.
- [ ] A senior employee was interrupted because the reasoning was not available.
- [ ] An exception was handled from memory rather than from a visible rule.
- [ ] Leadership could not see the backlog, age, or true status without asking someone.
- [ ] A correction or delay reached the client, carrier, or producer.
- [ ] The team believes another person or another tool is the only way to create capacity.

0 to 2: monitor it. 3 to 5: map it. 6 to 8: it is already an operating constraint.

> Do not score the ideal process Use what happened last time. The exception is often where the real workflow becomes visible.

A SAFE BUILD ORDER

Fix clarity before complexity.

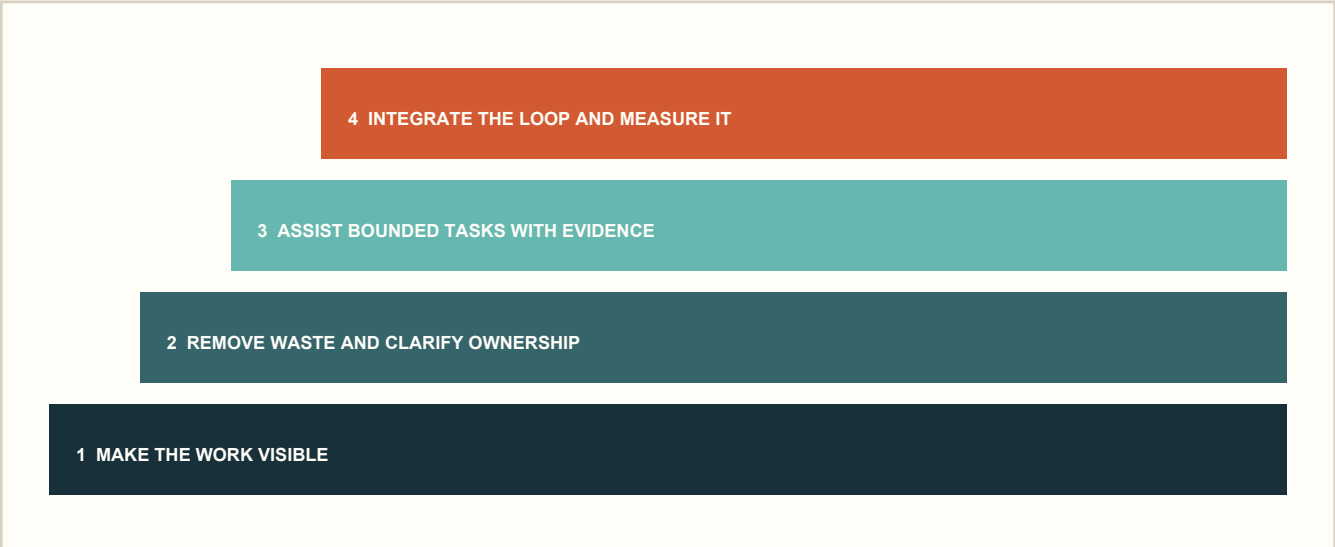

Agencies often want to begin at step four because integration looks like the solution. It becomes the solution only after the workflow, ownership, evidence, and authority are defined.

> *The first build is usually not software. It is shared clarity about how the work should run.*

THE CONCLUSION

The workflow comes first because the people come first.

The point is not to preserve every habit. It is to preserve judgment while removing the work that keeps good people from using it.

Agents are not waiting to be convinced that useful technology exists. Many are already experimenting. What they need is an operating environment that makes the safe path clear, connects the work, and lets them verify what the system produces.

That environment will not emerge from another license. The work begins at launch and becomes part of the agency's operating rhythm.

> *Make the work visible. Protect the decision. Then choose the technology.*

That is how an independent agency turns AI from a private workaround into durable capacity.

I map the work your agency runs on.

Regesta is workflow architecture for independent insurance agencies.

I am Andrew Cherry. I work with independent commercial P&C; agencies on the operational workflows that quietly consume service-team capacity: renewals, COIs, submissions, and shared-inbox work.

I am not here to sell software or a generic AI audit. I map one workflow exactly as it runs today, show where capacity, rework, delay, and risk accumulate, protect the decisions that must remain with licensed people, and give leadership a practical future state and build order.

> The first conversation Bring one recent example of a workflow that went wrong, ran late, or required more senior attention than it should have. We will use the real case to decide whether the problem is worth mapping properly.

Which workflow gives your team the most trouble today?

That is the one worth examining first.

BOOK A WORKFLOW CONVERSATION

Regesta · Workflow Architecture for Independent Insurance Agencies regesta.ai

The research behind the argument.

The four research signals and selected statistics in this field guide were derived from The Connective Thread: From agent and broker research to a new design vision for AI-enabled insurance work , supplied by the user. Regesta did not conduct the 16 participant study and does not claim ownership of its interviews or participant language. The operating thesis, independent-agency translation, workflow examples, diagnostic method, and architecture in this edition are Regesta's own.

- Agent for the Future / Liberty Mutual. Artificial Intelligence for Insurance Agents: 2025 Research. 2025.
- Boston Consulting Group. AI Adoption in 2024: 74% of Companies Struggle to Achieve and Scale Value. 2024.
- Gartner. Research on AI strategy, governance, productive users, financial guardrails, and agent decommissioning, 2026, as cited in the source report.
- McKinsey & Company. Building the Foundations for Agentic AI at Scale.
- MIT Sloan Management Review. Cracking the Code on Gen AI Value. 2025.
- U.S. Chamber of Commerce. The America Works Report: Industry Perspectives.

Editorial note Statistics are presented as reported in the supplied source document. Before public distribution, links and exact publication titles should be checked against the original publisher pages if Regesta wants a fully linked bibliography.

AI-NATIVE INSURANCE OPERATIONS

The agency runs between the systems.

That is where the work must be made visible.

METHOD	MAP THE REAL WORK	ONE WORKFLOW
BOUNDARY	PROTECT LICENSED JUDGMENT	EXPLICIT
OUTPUT	FUTURE STATE AND BUILD ORDER	OWNED