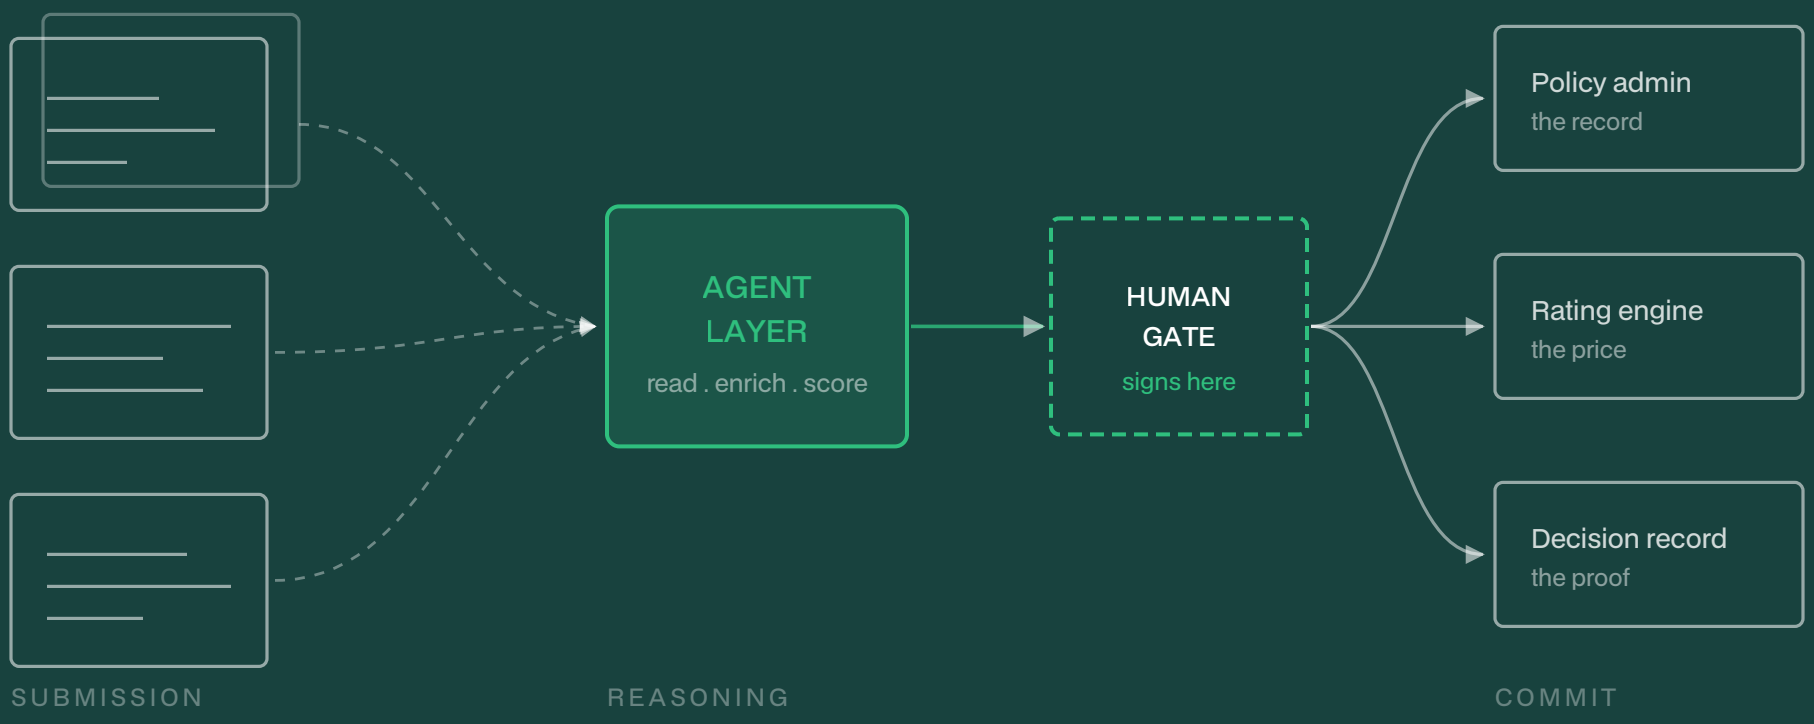


Six seconds. Then thirty one hours.

A machine read an insurance file in six seconds and got it right. The decision still took another thirty one hours. This is about the gap, and about why it is not going away.

48 slides
One file, nine stops

Scope
Insurance, and everyone next

October 2026
Global AI Forum

Global AI *forum*

A food company, and a spreadsheet nobody has opened.

Everything that follows happens to this one email. The details are illustrative, built from the pattern rather than taken from one account.

IT ARRIVES	**09:14, Tuesday** An email with four attachments, from a broker they have dealt with before
THE COMPANY	**A mid-sized food manufacturer** Eleven states, refrigerated warehouses, one recent product recall
THE SPREADSHEET	**2,800 locations** Excel, merged headers, no consistent layout, building types typed in by hand
THE HISTORY	**Four previous insurers** Four formats, four ways of counting a claim, two different valuation dates
WHAT IT IS WORTH	**About $140,000 a year** Roughly the average new piece of business in this part of the market
TIME AVAILABLE	**48 hours** After that the broker takes it somewhere else

The machine finished in six seconds.

6

seconds for the machine to read all of it

31

hours before anyone could act on it

0

insurers who have published whether it worked

> Everyone is arguing about whether machines can do the work. That argument is over. They can. The interesting question is why the work being finished did not finish the job.

HOW TO READ THIS

One file, nine stops.

Rather than survey an industry, this follows a single piece of business through it. The clock runs across the top. Each stop looks at four things.

01 **Where the file is**

What just happened to it, and what is about to.

02 **What we observe**

The pattern across the deployments we could verify.

03 **The shape of it**

One picture per stop, showing the mechanism underneath.

04 **What can be built**

Four places a company could stand once that is understood.

© www.gaiforum.com | 2026

01

A machine learned to read the mess.



The spreadsheet has 2,800 rows, headers that merge, and building types typed in by hand. **For twenty years no software could read this.** Now it can. Six seconds, and 98 percent of it right.

This is still how business arrives.

Half of all commercial insurance requests turn up as an email with attachments. 49 percent of agencies send them that way. Seven percent use software built for the job.

Seventy percent then retype the numbers into another screen by hand. And 90 percent say they have **taken business to a different insurer** because one was too slow to answer.

ONE THING YOU NEED TO KNOW

What is in the email

A request to be insured. It says who the company is, what it owns, where, and what has gone wrong before. Everything that follows is a judgement about those four things.

WHAT FOLLOWS

Reading was never the hard part. Watch what the machine does next.

NEXT . SIX SECONDS LATER, IT HAS AN ANSWER AND STOPS

What the machine had to read.

Four attachments, in order of how hard each one actually is.

HOW A COMMERCIAL SUBMISSION ARRIVES

49%
arrive by email

7%
use capture software

Ivans / Applied Systems, 702 agents, 2026

WHAT IS IN THE ATTACHMENT, BY DIFFICULTY

Solved	**ACORD PDFs**	Commodity optical character recognition
Hard	**Loss runs**	One per prior carrier, each a different layout and taxonomy
Hardest	**Statements of values**	Arbitrary Excel, merged headers, thousands of locations, schema inferred per file
Unmeasured	**Semantic normalisation**	"CONST TYPE: JM" to an ISO construction class. No public accuracy benchmark exists.

70% of commercial agencies re-key the data by hand.
90% have moved business away from a carrier over submission friction. This is a distribution problem, not a cost one.

Figure 01 . Channel figures are from the Ivans and Applied Systems connectivity survey, 702 independent agents, fielded April and May 2026. The difficulty ranking is ours.

The money is in the meaning.

Reading the words is now cheap. Knowing what they mean is not.

01

Translate, do not just transcribe

A building listed as JM has to become a known construction class with a known fire behaviour. That needs an insurance dictionary nobody wants to maintain.

02

Guess the shape of every file

These spreadsheets have no standard layout. The structure has to be worked out fresh, every single time.

03

Ask the broker, then wait

Half of what you need is missing. Asking for it, waiting, and picking up where you left off is the hard engineering, not the reading.

04

Publish an accuracy score

Nobody has published an independent measure of how well any of these tools read a spreadsheet. The first company to do it sets the standard for everyone.

09:14:06 TUE

THE ELEVEN

02

Reading is not deciding.



Here is what people outside the industry do not know. Saying yes to a risk is not one decision. **It is eleven.** The machine does the first four alone. Then it slows down. Then it stops.

Automation stops where the signature starts.

Sorting the post, checking for duplicates, deciding whether this is even the kind of business we write, pulling in flood and fire data: all of that runs without a person now.

Working out what the risk really is, what to charge, and what to promise: that is **helped by a machine, not made by one**. And the moment of actually saying yes is human, everywhere it matters.

ONE THING YOU NEED TO KNOW

What saying yes means

It is called binding. The company is now legally on the hook for whatever happens to that business for the next year. It cannot be taken back. That is why it is the one step nobody automates.

WHAT FOLLOWS

Four of eleven. That is what has actually been automated.

NEXT . 09:21, WHERE DOES THE REASONING GO

Eleven decisions, one email.

Where the machine has reached, where it helps, and where a person still holds the pen.

THE DECISION	AUTOMATED	ASSISTED	HUMAN
Submission intake Email, portal, broker feed, parsed to a record	●	○	○
Clearance Duplicate, conflict, broker of record	●	○	○
Appetite triage In or out, winnability, what to work first	●	●	○
Enrichment Hazard, financials, imagery, prior claims	●	○	○
Exposure analysis COPE, schedule of values, cat accumulation	○	●	●
Risk selection Accept, decline, refer	○	●	●
Pricing Technical price, then the market price	○	●	●
Terms and wording Limits, attachment, exclusions, endorsements	○	◐	●
Quote, bind, issue The signing act, governed by authority	○	◐	●
Portfolio steering Line size, aggregation, appetite setting	○	●	●
Audit and peer review Was it done right, and can we prove it	◐	●	○

Figure 02 . Commercial property and casualty insurance. Based on the deployments documented in this deck, not on a survey. Home and car insurance automated most of this twenty years ago.

Automate the chain, not the job.

Four places to build, all of them before anyone signs anything.

01 Turn the rulebook into code

Which business we will and will not look at is already written down. Most of it is unambiguous. Very little of it is in a machine.

02 Decide what to look at first

Choosing which email to open is a decision about staff time, not about a customer. Getting it wrong costs revenue, not harm. That makes it the safest thing to automate.

03 Clean the history

Four previous insurers, four different formats, four different ways of counting a claim. This is pure data work and it is where accuracy is actually lost.

04 Hand over a finished file

When it does reach a person, give them the evidence assembled, the similar cases pulled, and the question stated.

The answer and the reason end up in different places.

The policy	the legal document. Lives in the old system.
The workbench	where the thinking happened. A newer system.
The models	what scored it and priced it.
The reason	why this answer. Nobody built a home for this.

The machine has an answer, and it has a reason for the answer. The answer is written into the system that holds the policy. **The reason is not.** In eighteen months somebody will ask about the reason.

The old systems are not being replaced.

Everyone expected the new AI tools to take over from the decades-old systems that hold the policies. The opposite happened. The old systems made themselves the place everything has to pass through, and the AI companies got bought.

One supplier states in writing that its software **takes no action without a person approving it**. None of them claims the machine writes the final answer on its own.

WHAT FOLLOWS

The record of why is not a feature to add later. It is the product.

NEXT . 09:22, WHAT HAD TO EXIST FOR ANY OF THIS

Three ways an answer reaches the record.

The difference is not how clever the machine is. It is who presses the button, and what gets written down next to it.

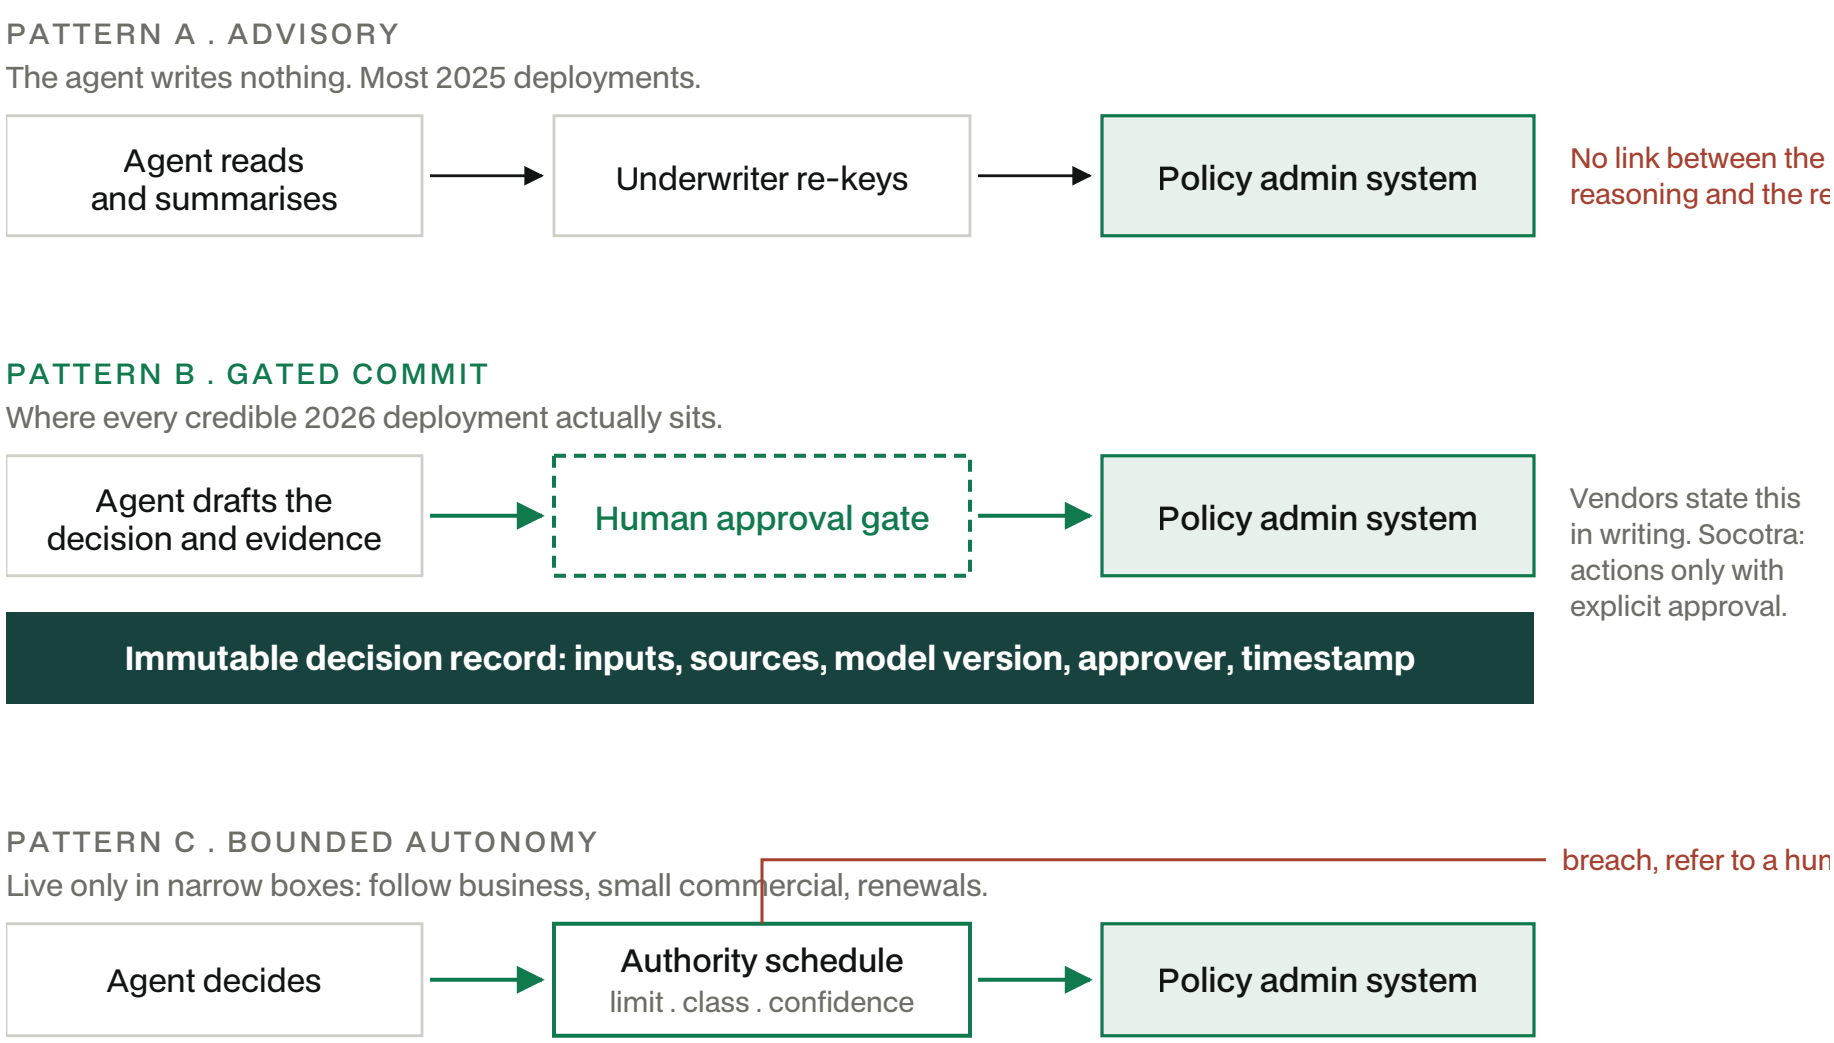


Figure 03 . The middle pattern is where every deployment we could verify currently sits.
Source: vendor product documentation and release notes, read October 2026.

Build for the moment of commitment.

Four things that are cheap now and impossible to add afterwards.

01

Write down why, as it happens

What it read, where that came from, which version of which model, who approved it, and when. Not reconstructed later. Captured live.

02

One door into the old system

Named, logged, reversible, and scoped. Not a shared password in a config file.

03

Make the record do three jobs

The same file answers the regulator, trains the next version, and settles the argument. Build it once.

04

Let it be rewound

A run you can take back to any point and start again from there.

04

Eight minutes of work sits on seven layers.

Reading, context, data	where everyone is building
The agents	the part people talk about
Permission	where every system stops
The record	what everyone is protecting

None of this is one clever model. It is seven layers, and three of them get built last. Those three are the first ones anybody asks about when something goes wrong.

Everyone builds the safety rails last.

Who the machine is, what it did, and whether it is still behaving: these run alongside everything else and are consistently the final thing anyone gets to.

European rules now treat every outside AI service in this chain as critical infrastructure. US insurance regulators expect any decision to be **traceable back to a specific version and a named person**.

WHAT FOLLOWS

Layers one to four are where the work is. Layer six is where it stops.

NEXT . 09:40, SOMETHING DISAGREES WITH THE ANSWER

© www.gaiforum.com | 2026

What the eight minutes sits on.

Read down for how the email moved. Read across for what a regulator will ask about.

THE STACK	WHAT BINDS IT
1 **CHANNEL** — Email, portal, agency feed, bordereaux, slip	49% email, 7% capture
2 **INTAKE** — OCR, schema inference, semantic mapping, gap query	No public benchmark
3 **CONTEXT** — Risk ontology, guidelines, appetite, treaty wordings	AIG: 4m+ data points
4 **ENRICHMENT** — Hazard, cat, firmographics, imagery, prior claims	Metered. Licence bound.
5 **AGENTS** — Supervisor, appetite, exposure, pricing, wording, critic	Deterministic state
6 **AUTHORITY** — Confidence, line size, class, referral, approval gate	Where deployments stop
7 **RECORD** — Policy admin, rating engine, cession, billing	Agent drafts, human commits

THREE CONTROL RAILS, SPANNING EVERY LAYER, USUALLY BUILT LAST

IDENTITY	Whose authority is this agent acting under
LINEAGE	Inputs, model version, approver, timestamp
EVALUATION	Shadow mode, drift, silent appetite creep

Figure 04 . Assembled from what Guidewire, Duck Creek, Socotra, Cytora and AIG have each described separately. No single company has published the whole of this.

Ship the rails first.

Four decisions that are cheap at the start and very expensive later.

01

Give the machine its own name

Not a borrowed login. Its own identity, so that what it did is separable from what a person did. The standards for this arrive in 2027. Build toward them now.

02

Put a meter on the data

Every flood score and credit check costs money and comes with rules about reuse. A machine left to loop runs up a bill and breaks a contract at the same time.

03

Keep the steps fixed

The order of a regulated process should not be decided by a model in the moment. The conductor routes. It does not improvise.

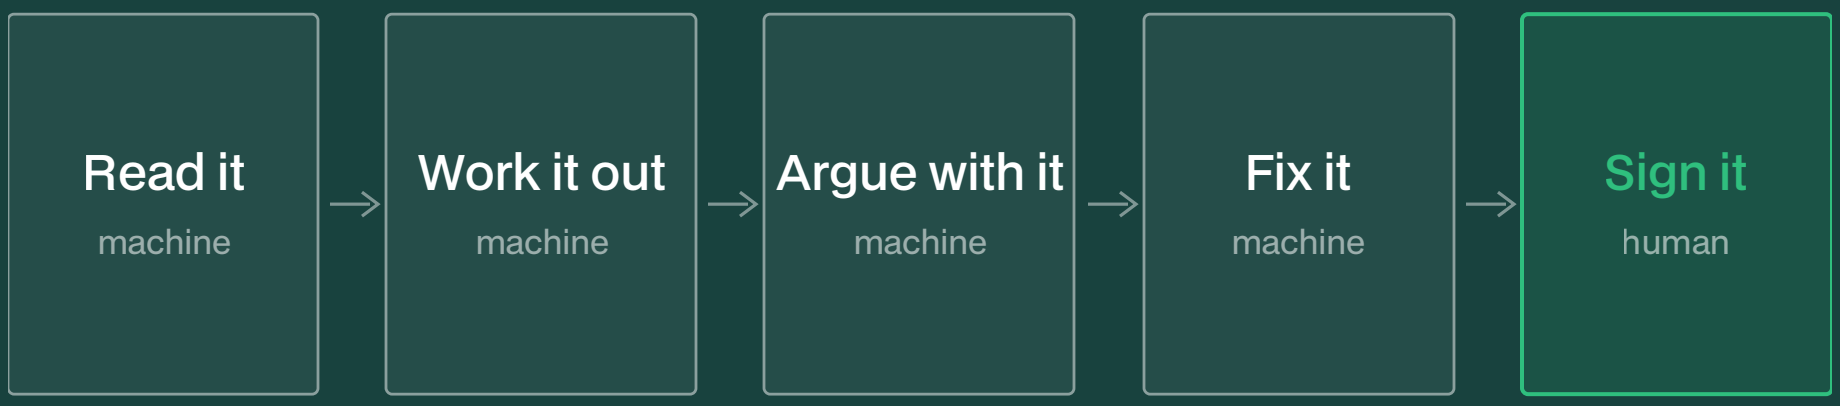
04

Be able to leave

European rules now expect a tested exit from any supplier, not a promised one. Building deeply into one provider is now a finding, not a strategy.

05

A second machine argues with the first.



A second agent reads the first one's work and disagrees about flooding at nine of the 2,800 sites. The answer changes before any person sees it. This is peer review. It just moved inside the computer.

Giving it more documents made it worse.

Tested on 635 real applications, a plain model got 77.6 percent right. The same model with a document search bolted on got 76.9 percent. Slightly worse.

The version that could tell when it was missing something, go and find it, and **give up and ask a human when it could not**, got 86.5 percent. The gain is not in the reading. It is in knowing what you do not know.

WHAT FOLLOWS

The argument is not the expensive part. The false alarms are.

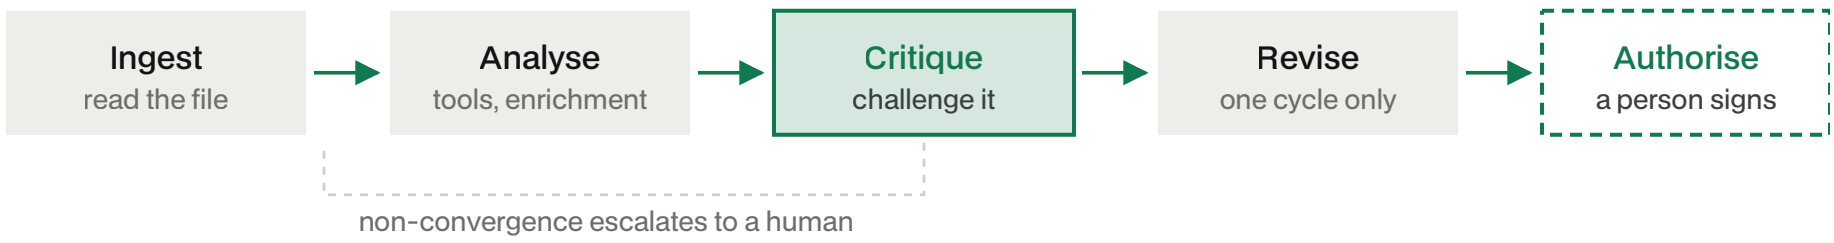
NEXT . 10:00, WHO BUILT THE THING THAT JUST DID THAT

© www.gaiforum.com | 2026

What the second opinion buys.

One round of challenge, one round of revision, and a handover to a person when it cannot settle.

THE RUN LOOP



WHAT THE CRITIC BUYS, MEASURED ACROSS 500 EXPERT-VALIDATED CASES

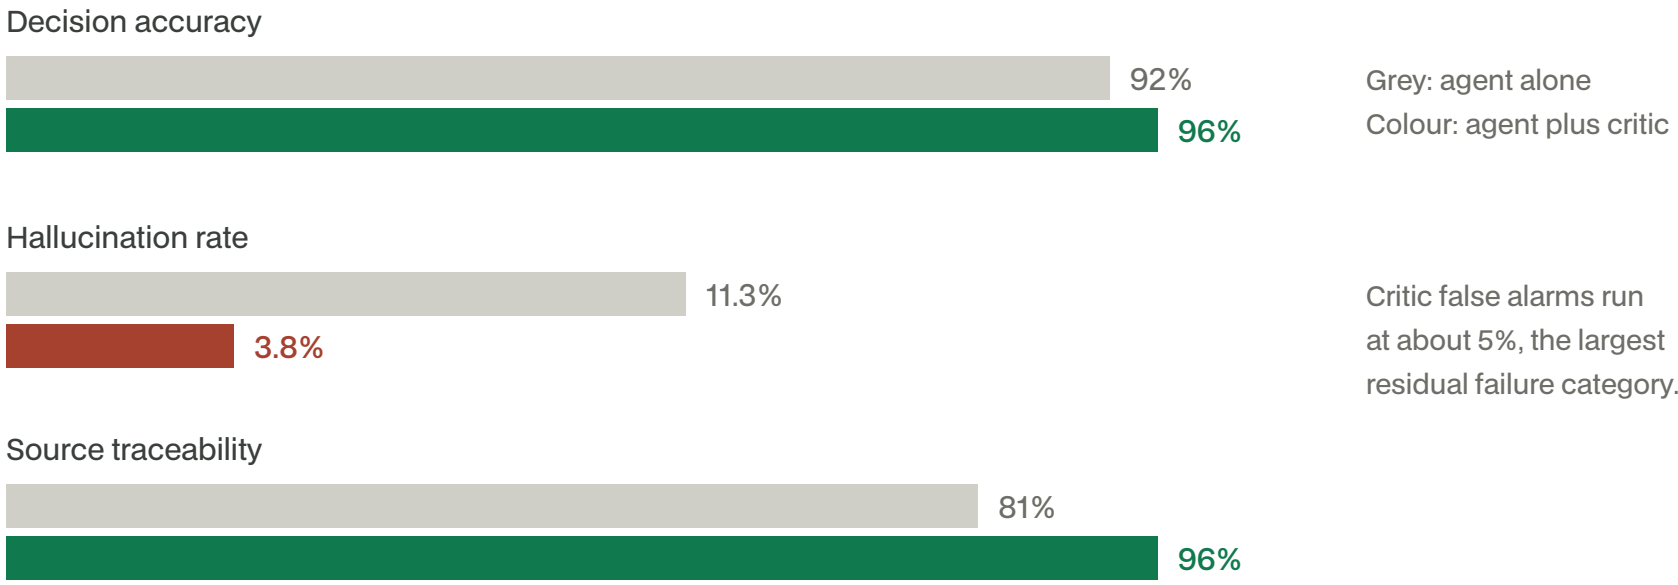


Figure 05 . Measured across 500 cases checked by experienced underwriters, in published research. Two limits carry: the challenger uses the same underlying model, so it shares the same blind spots, and the human comparison is the answer key by definition.

Put the second opinion inside the machine.

Four choices that make a system defensible rather than impressive.

01

Take away the pen, not just the permission

The cleverest control in the research: the machine's answer form has no box for saying yes. It cannot commit because there is nowhere to write it, not because it was told not to.

02

Argue once, then stop

One round. Letting a machine critique itself repeatedly amplifies its mistakes rather than correcting them.

03

Score the false alarms

Roughly one flag in twenty is wrong, and that is now the biggest remaining problem. Nobody has published a working fix. That is an opening.

04

Search for words, calculate with code

Asking a model to do the arithmetic in its head is a known way to get a wrong number. Separate the two jobs.

Who built the thing that just did that.

Cytora	bought by Applied Systems, September 2025
Send	bought by Duck Creek, July 2026
CAPE Analytics	bought by Moody's, January 2025
Planck . Groundspeed	both bought, 2023 and 2024
Still independent	Federato, hyperexponential, Sixfold, Akur8

Three years ago this was a startup's product. Today most of those startups belong to the companies that sell the old systems. **The new layer was not competed with. It was bought.**

One word, four different things.

Agentic is being used for workflow software with a score attached, for systems that genuinely run on their own under supervision, for assistants that wait to be asked, and for compliance checkers.

Only the second one is agentic in the sense the word implies. And exactly one supplier of the core systems has a product that is **finished, available and documented**. Everything else is early access.

WHAT FOLLOWS

The suppliers consolidated. Four places to stand were left behind.

NEXT . 11:30, FOUR OTHER INSURERS, SAME MORNING

Where everyone actually sits.

How far along the decision each supplier reaches, against how much it is allowed to do alone.

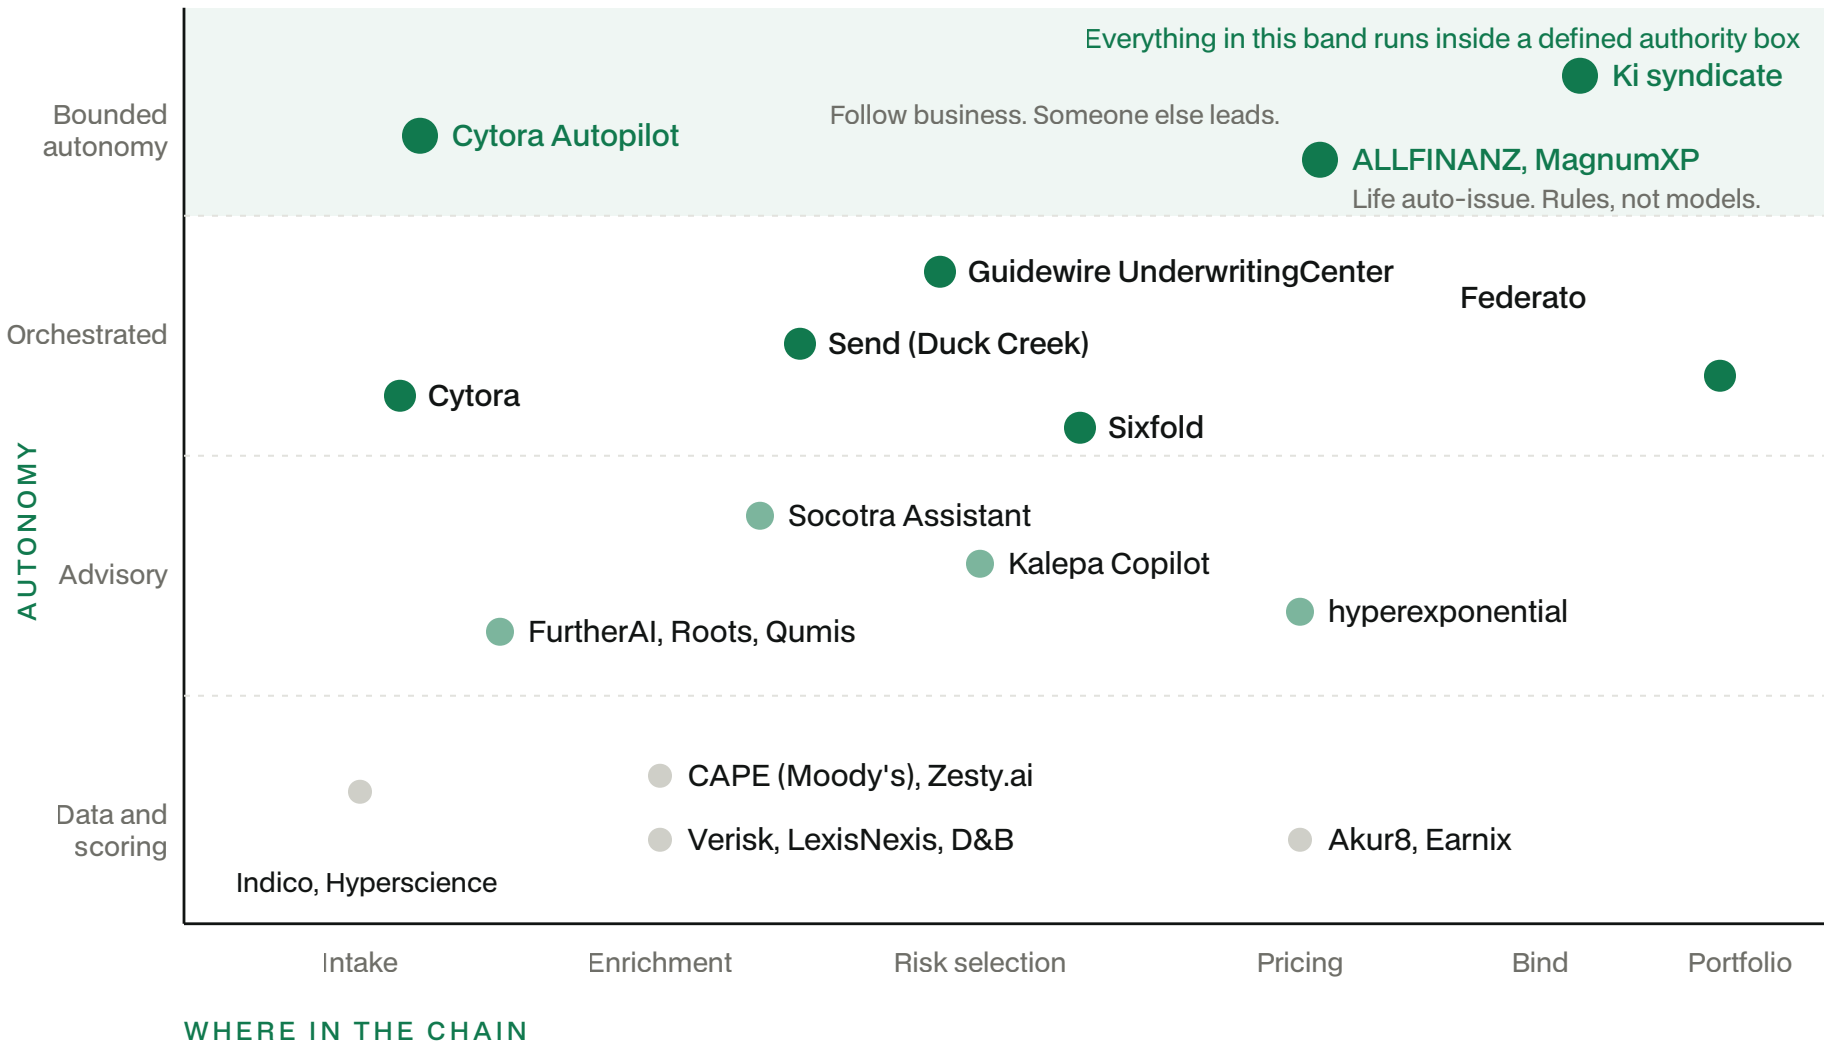


Figure 06 . The two most independent systems in this picture are the oldest technology in it. Both work by removing the judgement from the problem rather than by adding intelligence to it.

Four places still open.

What the buying spree left behind.

01 ## The record of why

Nobody owns it. It sits between the new systems and the old ones, and it is the thing every regulator actually asks to see.

02 ## Proof that it works

No insurer has published a test of its machine against its own people. The tools to run that test do not exist as a product.

03 ## The checking machine

Reviewing work already done is lower risk than doing the work, and it is the thing that makes everything else defensible.

04 ## Everyone outside the top tier

The largest insurers are served. A regional company running software older than the web is not, and cannot buy a two-year programme.

07

Four other insurers, the same morning.

AIG	three to four weeks, down to under a day
The Hartford	75 percent of small quotes, no human involved
Travelers	reading the post, hours down to minutes
W.R. Berkley	more than 20 percent efficiency gain
Zurich	five countries in ninety days

The same email went to four other insurers this morning. Five companies in the world have said publicly what happens next. Everyone else has done it and said nothing.

The confidence went down, not up.

In February 2026 one insurer described its programme with numbers and said it was running ahead of plan. By August a new chief executive described the same programme with adjectives and no figures.

Another chief executive said in October that the gains are **hard to measure**, and that savings will go into lower prices rather than profit. In a competitive market, getting faster becomes a discount, not a margin.

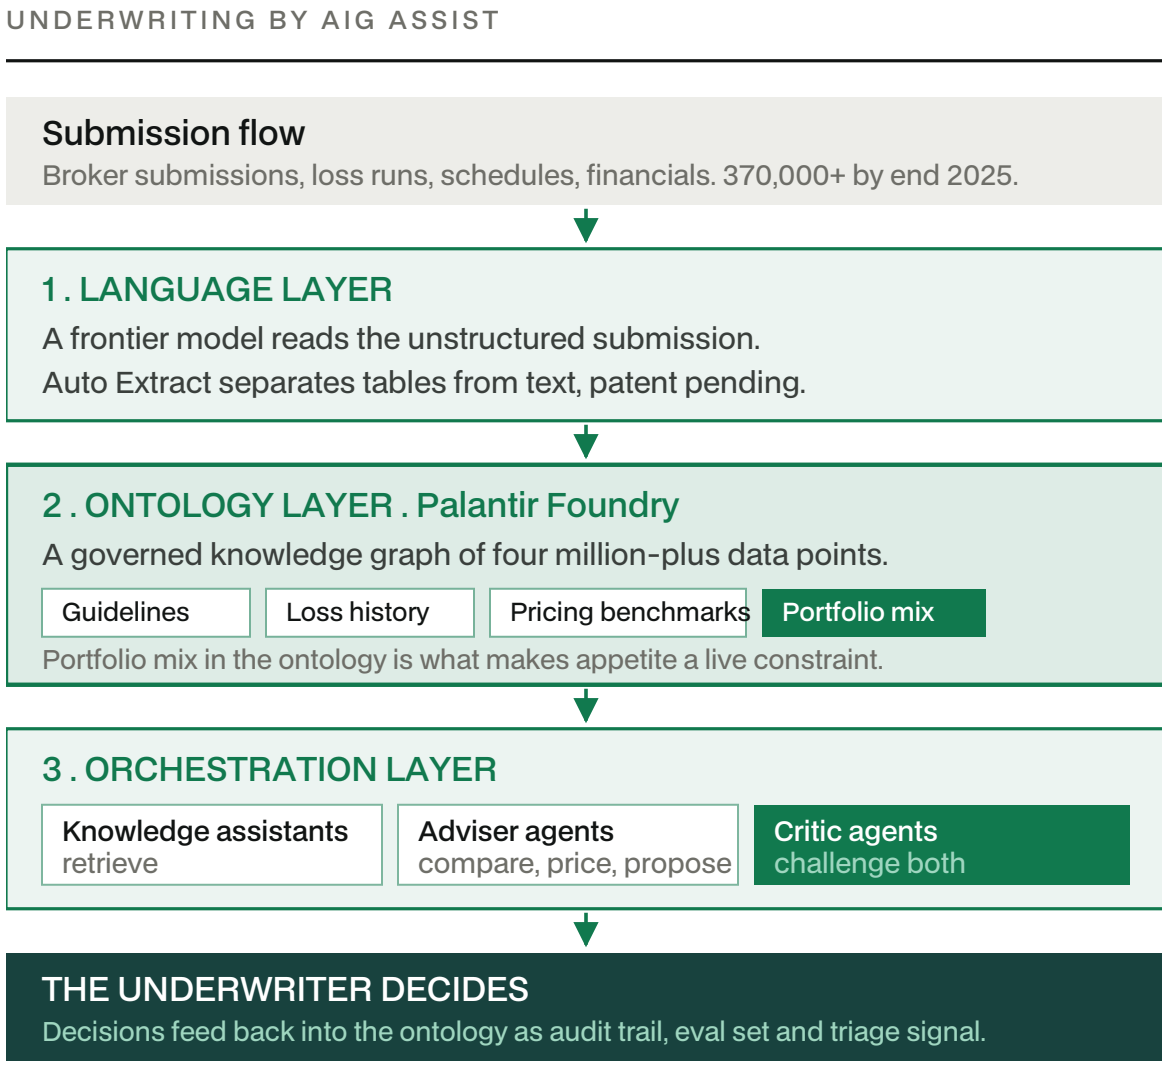
WHAT FOLLOWS

Described with numbers in February. Described in adjectives by August.

NEXT . 14:00, THE ONLY QUESTION THAT MATTERS

The most documented programme, rebuilt.

Three layers. The one worth copying is the middle one, and it is the one nobody talks about.



DISCLOSED

3–4 wks
to under one day, per submission
NA financial lines, Investor Day 2025

15 to 20%
bind-to-submit ratio

100%
of submissions reviewed, no underwriters added

+26%
Lexington submission count

~40%
increase in binding, but
one line only. No baseline given.

Every figure is company-disclosed and unaudited.

Figure 07 . Assembled from earnings call transcripts, granted patents and company announcements between November 2024 and August 2026. Every figure here is company-disclosed and unaudited.

What a real programme looks like.

Four things the documented ones have in common.

01

A map before a model

The leading programme gave its machines a structured map of four million facts about its own business, rather than a pile of documents. The map is the asset. It was the map they patented.

02

One product line, then seven

Eighteen months from first line to many, at the best resourced company in the market. Nobody went wide first.

03

More work, same people

Every company on the record says it handled more without hiring. None of them says it handled the same with fewer. Those are different claims.

04

A growth reason, not a cost one

One insurer looks at a hundred requests and takes two. The prize is not the cost of the two. It is the ninety eight it never properly read.

14:00 TUE

THE QUESTION

Did the company make better bets?

How fast it went	published by everyone
How accurately it read	published by everyone
How much ran untouched	published by everyone
How many it won	published by one company
Whether the bets were good	published by nobody

Here is the only question that matters, and nobody can answer it. Reading faster is not the same as choosing better. **Not one insurer anywhere has published whether this improved the thing the business exists to do.**

Four reasons it is missing. One of them is awkward.

It takes two to three years for a year of insurance to show whether it was profitable. And the AI arrived at the same time as price rises and a changing market, so no one can separate the effects.

And any company that has worked out a real advantage has **every reason to keep quiet about it**. So hold both ideas at once. The claim is unproven, and the silence is partly explained by self interest.

ONE THING YOU NEED TO KNOW

The number everything hangs on

Of every hundred pounds an insurer collects, how much it pays out again. Below a hundred it made money on the business. Above it, the business lost money and only the investments saved it. Faster reading does not

WHAT FOLLOWS

A machine that copies your judgement perfectly copies your results exactly.

NEXT . 16:40 WEDNESDAY, A PERSON SIGNS

Where the claims actually sit.

Four levels of proof, and one empty shelf. Note which claim travels furthest.

THE EVIDENCE LADDER

AUDITED

Ki syndicate FY2025: adjusted combined ratio 91.3%, from 95.2%.
Profit after tax $110.4m, up 50%. Company financial statements.

1 entity

And it writes follow business, where someone else sets terms.

CARRIER-ATTRIBUTED, ON THE RECORD

The Hartford: 75% of small business quotes with no human touch.
AIG: three to four weeks to under a day. Bind-to-submit 15 to 20%.
Travelers: submission ingestion hours to minutes.

5 carriers

AIG, Zurich, The Hartford, Travelers, W.R. Berkley. That is the whole set.

VENDOR-PUBLISHED

Zurich intake straight-through 10% to 95%, triage 75 to 15 minutes.
From Cytora's case study, not a Zurich filing.

The most repeated numbers in the discourse sit here.

UNSOURCED, IN WIDE CIRCULATION

"3 to 5 point loss ratio improvement." "40% lift across eight lines." "5x compression."

None traceable to a carrier statement.

MISSING ENTIRELY

An isolated loss ratio effect attributable to AI underwriting.

At any carrier. In any line. Also missing: any headcount effect, any published shadow-mode result, any life auto-issue rate.

Figure 08 . A defensible answer on profitability arrives in 2027 and 2028, once the 2025 deployments have two to three developed years behind them.

Measure it before you install it.

Four measurements almost nobody is taking.

01

Hold some back

Run the machine alongside your own people on the same business, for long enough to see which chose better. No insurer has published one of these.

02

Watch the ones you turned away

A rejection earns nothing, so a bad rejection never shows up in the numbers and never gets corrected. There is no feedback on saying no.

03

Watch for quiet drift

The more capable system was measurably worse at obeying a simple hard rule. It can reason its way around a line it should have just stopped at.

04

Keep every override

Every time a person overrules the machine is a lesson. It is also how a system ends up learning only from its own past choices.

16:40 WED — THE SIGNATURE

09

In the end, a person signs.



Thirty one hours after it arrived, a human being puts their name on it. Everything the machine did was preparation for that one moment. **And the rulebook for that moment is four hundred years old.**

Nobody is asking you to explain the model.

There is no single rule about explaining AI decisions. There are five, for five different audiences. Almost none of them asks how the model works.

They ask what you told the customer, which facts you used, where those facts came from, and **who was allowed to make the call**. That is a filing problem, not a mathematics problem.

ONE THING YOU NEED TO KNOW

The four hundred year old answer

Insurance has always let people who do not work for it commit its money. The instrument is a written grant: these kinds of business, up to this size, in these countries. Go outside it and you must ask. You report what you did. You get audited. The insurer stays liable throughout. It has a name. It is

WHAT FOLLOWS

The market regulator spent 2026 tightening the rules on permission, not on AI.

NEXT . WHAT HAD TO BE TRUE FOR THAT SIGNATURE

The rulebook already exists.

What insurance already does for people, line by line, against what a machine needs. The match is almost exact.

WHAT INSURANCE ALREADY DOES		WHAT THE AGENT NEEDS
The coverholder		The agent
Written binding authority agreement	→	**A versioned authority grant in code**
Class, limit, territory, pricing basis	→	**Line of business, maximum premium, confidence fl**
Named individual who holds it	→	**The licensed underwriter the agent acts for**
Mandatory referral above the box	→	**Escalation on breach or low confidence**
Bordereaux reported periodically	→	**A decision record written per run**
Policed by coverholder audit	→	**Shadow mode, drift monitoring, file review**
Carrier remains liable throughout	→	**The insurer remains liable throughout**

An automated underwriting system is a grant of authority to software. Insurance has governed that for a cent

Figure 09 . This mapping is our argument, not a regulator's instruction. It lines up with what supervisors are asking for and with a process the market already knows how to run.

Write the permission before the software.

Before the model, before the data, before the pilot.

01 **Say what it may do, in writing**

The same document a human gets. Kinds of business, size limit, countries, and the point at which it has to stop and ask.

02 **Name who answers for it**

One US state already requires a named officer to sign that the controls work, and a remediation plan if they cannot. Decide now who that is.

03 **Make the limits testable**

Not escalate when uncertain. A number, a list, a ceiling, and a log entry every time one is reached.

04 **Three problems nobody has solved**

A machine that quietly never answers leaves no trace of a rejection.
One mistake made ten thousand times is a different kind of problem.
And nobody knows who is personally accountable.

16:40 WED — SIGNED

The thirty one hours were not waste.

They were the time it took for the liability to find somebody to sit with. Here is what the file had to carry before anyone would put their name on it.

01 **The spreadsheet was believed**

2,800 rows read, nine flood exposures corrected after challenge, and a note of where every figure came from.

02 **The yes or no was traceable**

Not a score out of a hundred. A rule, a version of that rule, and the state of the book it was checked against.

03 **The price had two numbers**

What the model said, and what the person chose. Both recorded, with the difference explained.

04 **Someone held the permission**

A named person, inside a written limit, with the line they did not cross logged next to the one they did.

05 **The record outlives everybody**

Readable in eighteen months, after the reorganisation, after whoever championed it has left.

06 **Nobody knows if the bet was good**

And will not until 2028. Everything above is proof of process, not proof of judgement.

WHY THIS MATTERS OUTSIDE INSURANCE

You can hire somebody brilliant. You still cannot let them sign cheques on day one.

That is the whole thing. Not a limit on intelligence. A limit on **permission**.

For twenty years, the hard part of automating knowledge work was comprehension. Machines could not read what people send each other. That is now solved, and a lot of people are about to discover it was the easy half.

The half that remains is answering who allowed this, on what evidence, within whose limit, in a way that still stands up eighteen months later when the person who approved it has left.

That is not a model problem. It is a **filing problem**, and insurance happens to be the industry that has been solving it the longest.

Now run that file a hundred thousand times.

One file handled well proves very little. Here is what the human signature does not protect you from.

01 **A document that lies convincingly**

Someone sends a forged safety report. It contains no instruction and no trick. It simply reads well. This is the one attack that does not close.

02 **Coverage nobody meant to give**

Roughly three in a hundred answers still contain an invented detail. In a contract, an invented detail is a promise.

03 **Quiet drift**

A machine clever enough to reason can reason its way around a rule it was supposed to just obey.

04 **Learning from its own choices**

It only ever sees how the business it accepted turned out. It never finds out about the business it turned away.

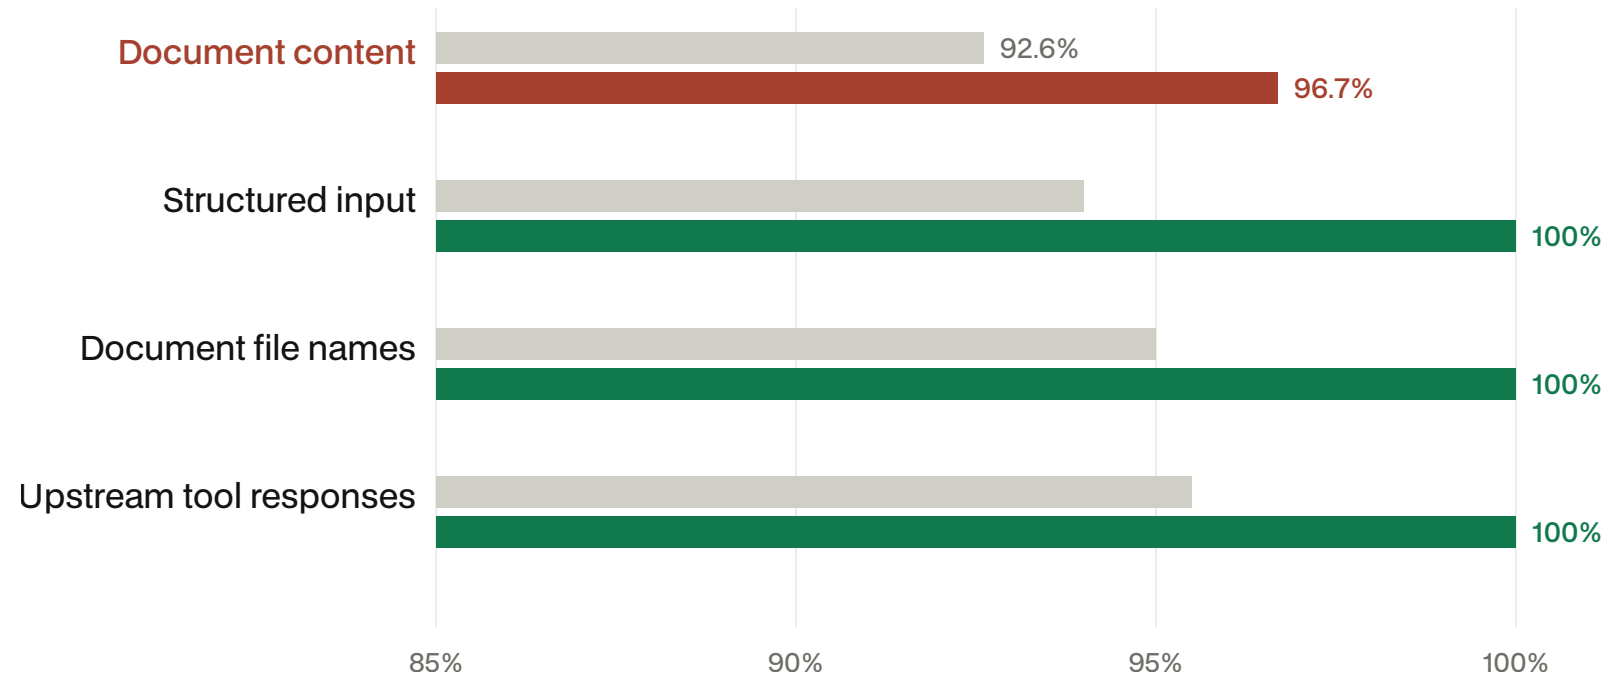
05 **Nobody learns the job**

Sorting the post is how junior people learn to judge risk. Automate it and the path to senior disappears.

One way in that will not close.

Marking everything the machine did not write itself raises the defence rate from 93.2 percent to 97.7 percent, with no false alarms on honest files. It closes every route except one.

DEFENCE SUCCESS RATE, BY WHERE THE ATTACK ENTERS



Grey: no defence. Colour: untrusted content wrapped in boundary tags. Every vector closes except the one the underwriting process exists to receive.

Figure 10 . The sophisticated version contains no instruction at all. It is a convincing forged document, in exactly the form the machine was built to trust. There is nothing for a filter to catch, because nothing is being smuggled in. Source: published benchmark of 78 cases across 16 models, June 2026.

WHAT TO DO

Five things, in the order they actually bind.

If you are building this, or buying it, this is the order that survives contact with a regulator.

01

Write down what it may do

Before the model. Kinds of work, size limit, where it must stop and ask, and who signs that the limits hold.

02

Build the record before the machine

What it read, where that came from, which version decided, who approved. It cannot be added afterwards.

03

Map the ground before you search it

Settle what things are once, in one structured place, rather than working it out again on every question.

04

Automate the checker first

Reviewing work already done is lower risk than doing the work, and it is what makes everything else defensible.

05

Mark everything you did not write

Incoming files, data from outside, anything the machine did not author itself. It costs nothing and closes almost every way in.

© www.gaiforum.com | 2026

WHAT WE COULD NOT ESTABLISH

Six things this deck does not know.

Stated here rather than left out, because the gaps are part of the finding.

01 **Whether any of this improves results**

No insurer has published a profitability effect attributable to AI. Not one, in any market.

02 **What share of insurers actually run it**

The regulatory surveys that measure this were taken between 2021 and 2023 and are now a baseline, not a reading.

03 **How accurately these tools read**

No independent benchmark for spreadsheet or claims-history extraction exists in public.

04 **Whether the machine beats the person**

No company has published a side by side test against its own underwriters.

05 **What it costs to run properly**

No credible public figure exists for the cost of reviewing a single piece of business, before or after.

06 **Where the jobs go**

Every company on the record says it handled more work without hiring. None has published a headcount effect either way.

The file is illustrative. The numbers are not.

The piece of business followed through this deck is assembled from the pattern rather than drawn from one account. Every figure attached to it is sourced, and sits at one of three levels.

AUDITED

Company financial statements

One company in the whole field, and it works in a narrow corner of the market.

COMPANY STATED

A named executive, on the record

Five insurers. Earnings calls, investor days, conference remarks.

SUPPLIER PUBLISHED

Published by the vendor

Including the most repeated statistic in the whole category.

LEFT OUT ON PURPOSE

Three figures in wide circulation do not appear here. A claimed improvement in profitability that traces to no named company. A claimed increase in business won that covers one product line and is usually quoted as covering eight. And a claimed speed-up that appears in no company transcript we could find.

DRAWN ON

Company filings and earnings call transcripts, regulatory instruments from the NAIC, New York, Colorado, Texas, the European supervisory authorities and IRDAI, vendor release notes and availability tiers, the Ivans connectivity survey, Datos Insights, Celent, Capgemini and Munich Re industry surveys, and published technical research on agentic underwriting.

We check what AI actually did inside the enterprise.

Global AI Forum is an independent enterprise research body. Our whole focus is enterprise AI adoption: what was deployed, what it produced, and what the evidence will and will not support.

IN THE ROOM

Workshops

Sessions that turn a deployment question into a decision the room can defend.

IN THE ROOM

Briefings

Private briefings for leadership teams and their boards.

UNDERNEATH

Research

Deployment studies and regulation tracking that support the work above.

> Somewhere a file is sitting in an inbox right now. The question was never whether a machine can read it. It is who will be able to defend what happens next.

research@gaiforum.com · www.gaiforum.com · San Francisco and Chennai

© www.gaiforum.com | 2026