

PROPERTY AND CASUALTY INSURANCE

Trusted AI in Property and Casualty Insurance

Turning imagery and claim files into verified underwriting and claim decisions

THE P&C AI VALUE GAP

38%

of P&C insurers generate value at scale from AI in core workflows
BCG, 2026

3x

expected growth in industry AI spending as a share of revenue in 2026
BCG AI Radar, 2026

15% to 25%

potential reduction in operating cost per dollar of premium for AI leaders
BCG, 2026

25

states have adopted the NAIC AI Model Bulletin, with 8 more in process
Openlayer, July 2026

BUILT FOR HIGH-CONSEQUENCE DECISIONS

- Property and casualty carriers
- MGAs with underwriting or claims authority
- Third-party administrators
- Self-insured programs

EXECUTIVE SUMMARY

AI can read the claim. Trust decides the outcome.

The question for property and casualty (P&C) carriers is no longer whether AI can read a claim file or provide an underwriting score for a roof. It is whether those outputs can be trusted at the moment a premium is set, a reserve is booked, or a payment is released.

Investment is accelerating faster than results. BCG expects that insurance industry AI spending as a share of revenue will triple in 2026, yet finds that only 38% of P&C insurers are generating value at scale from AI in core workflows (BCG). The same research estimates that operating costs per dollar of premium could fall by 15% to 25% for carriers that redesign workflows around AI, rather than layering it onto processes built for human-led execution.

The timing matters. AM Best expected the industry combined ratio to improve to 95.0 in 2025, but projects commercial lines a couple of points higher in 2026 as rates soften and social inflation keeps loss severity elevated in commercial auto and general liability (AM Best). As profits come under pressure, insurers will be looking to further reduce expenses using AI. Regulators are already moving in parallel: As of July 2026, 25 states have adopted the NAIC Model Bulletin on insurers' use of AI, with eight more in process (Openlayer). The margin for unverified AI is narrowing from both sides.

Most carriers, MGAs, and third-party administrators (TPAs) now run AI in some way throughout the policy lifecycle: new business, claim and policy servicing intake, document extraction, imagery analysis, and underwriting prefill. Few have converted those deployments into systems that adjusters and underwriters rely on and that claims, actuarial, and compliance leaders can attest to for consistent decision-making. BCG frames the obstacle as strategic rather than technical. We agree and would add that we discern one capability underneath the operating model. Trust.

Trust in insurance AI is usually treated as a gate a model passes once, when it clears an extraction accuracy target on a curated sample of files. That framing is why so many programs plateau and do not scale. A book of business is not a test set. Demand letters change format, imagery vendors change capture angles, and a hailstorm produces damage no training set anticipated. A model that was trustworthy at launch can silently develop incorrect risk scoring by renewal season, with errors not surfacing until they appear as leakage or adverse reserve development.

The carriers moving fastest treat trust as an operating capability: the continuous ability to detect where AI is uncertain, route those cases to the resource with the right expertise, correct them, and return the correction to the system so the same error does not recur. In this model, production readiness is not a threshold a model crosses once. It is the rate at which an operation converts exceptions into verified knowledge.

This paper relates how unverified AI shows up in the pure loss and expense ratios that drive the combined ratio, defines what trusted AI requires in underwriting and claims, and describes how CloudFactory builds that capability through the trust loop, running on CloudFactory's AI Platform. It closes with five moves executives can make now, regardless of the platform they choose.

WHY P&C AI STALLS

P&C AI stalls after the pilot, not in the model.

The pattern is familiar. A claims intake agent extracts fields from first notice of loss (FNOL) submissions, ACORD forms, police reports and other relevant artifacts with high accuracy on a curated validation sample. The pilot team reports minutes per file instead of half an hour. Leadership approves expansion to two more lines.

Six months later, adjusters re-check most of what the agent populates, supervisors keep a manual file review running in parallel, but no one can say whether the system is improving or degrading. The pilot did not fail. It never earned the right to be relied on by the organization.

Three mechanisms keep P&C AI out of production

01 / ACCURACY Accuracy is measured once; the book changes continuously A validation score describes performance against the files and images available at launch. Production introduces covariate shift (a new plaintiff firm's demand letter template, a new imagery provider, a different roofing material in a growth territory) and concept drift (a revised policy form, a new state notice rule, an emerging injury pattern in litigated claims). Without continuous evaluation against live files, accuracy becomes a historical claim rather than a current fact.	02 / UNCERTAINTY Uncertainty has nowhere to go Most deployments return an extraction or a score with a confidence value, then stop. When the model is unsure, the output is either accepted, which creates silent leakage, or every file is touched by a cautious adjuster, which erases the expense case. What is missing is a routing layer that sends each uncertain case to the right level of expertise based on claim value, coverage complexity, injury severity, and regulatory exposure, and that captures the corrected decision that results.	03 / CORRECTIONS Corrections do not compound When an adjuster overrides an injury code or an underwriter corrects a prefilled roof attribute, that judgment lands in a claim note, a referral email, or memory. The model never learns from it. Each new line, state, or program starts from zero because the institutional knowledge generated by the last deployment was never captured.
---	--	--

Much of today's claims automation reinforces the gap by measuring the wrong thing. Document automation tools report time saved per file and extraction rates. Evaluation and observability platforms report drift and regressions, then leave remediation to the carrier's own team. Both are useful. Neither closes the loop between detecting an error and making sure it never recurs.

AI errors land in pure loss and expense ratios.

In insurance, AI errors rarely show up on a model dashboard. They show up in the financials, often long after the decision that caused them, as direct impacts on the pure loss ratio and the expense ratio. Those two ratios are the path through which AI, for better or worse, influences the combined ratio.

EXHIBIT 1 AI errors surface in pure loss and expense ratios, often years after the decision

Combined ratio what the market sees	=	Pure loss ratio accident year losses + prior-year reserve development	+	Expense ratio including loss adjustment expense
Component	How unverified AI shows up	When it surfaces	What the loop controls	
Pure loss ratio: current accident year	Risk priced on stale or unverified property attributes; coverage terms and limits missed at intake, so severity leaks into payments	At claim, renewal, or the next catastrophe	Verified attributes behind every rate; coverage match checked before payment	
Pure loss ratio: prior-year reserve development	Inconsistent injury coding and reserve calls across adjusters, offices, and TPAs	One to several years later, as adverse development in calendar year results	Consistent, expert-verified severity and reserve inputs on every file	
Expense ratio	Re-keying, parallel manual review, unnecessary site inspections, and reopened files	Immediately, but spread across headcount and vendor spend	Human effort concentrated on the exceptions that need it	
Regulatory and audit exposure	Missed state notices or Medicare Section 111 reporting; AI-influenced decisions that cannot be explained	At market conduct exam, audit, or complaint	Every step logged, with a traceable decision path	

Source: CloudFactory analysis. Calendar year pure loss ratio equals accident year losses plus prior-year reserve development, divided by net earned premium. Loss adjustment expense is shown in the expense ratio.

THE CORE INSIGHT

CloudFactory influences the combined ratio through direct impacts on the pure loss ratio and the expense ratio, and the two have to improve together.

A model that saves 20 minutes per file but shifts reserve accuracy by a few points has not done that. It has moved cost from the expense ratio into the pure loss ratio through future reserve development, where it is harder to see and more expensive to fix.

Trust architecture, not model quality, separates pilots from production.

EXHIBIT 2 Trust architecture, not model quality, separates pilots from production

Aspect	Piloted AI	Trusted AI in production
How performance is known	An extraction or accuracy score at launch	Continuous evaluation against live files, imagery, and a governed evaluation set
What happens to uncertain outputs	Accepted silently or reworked manually on every file	Routed by confidence, claim value, coverage complexity, and regulatory exposure to the right expertise
Where adjuster and underwriter judgment goes	Claim notes, referral emails, and individual memory	Captured, validated, and redeployed as training and evaluation data
How the system changes	Periodic retraining after leakage or a complaint surfaces	Every exception strengthens the model and the knowledge graph
How it scales	Each line, state, and program rebuilds rules and thresholds	Validated knowledge transfers across lines, states, and programs
Who can defend it	The vendor or data scientist who built it	Claims, underwriting, actuarial, compliance, and regulators, with traceable decisions

Source: CloudFactory analysis

A VISUAL LANGUAGE OF CONFIDENCE

Every output carries a score.

Confidence is shown, not hidden. Each state has one route and one owner, so an adjuster or underwriter knows exactly which outputs to rely on.

The distinction matters because the cost of a wrong answer in P&C is financial, regulatory, and personal to a customer. A missed exclusion becomes leakage. An inconsistent reserve becomes an adverse development. A mis-scored roof can become a mispriced risk for the life of the policy. An unexplained adverse action becomes a regulatory finding.

■ 0.97 · passed

Goes straight through to payment, reserve, or bind

■ 0.62 · in review

Routed to a specialist agent or an Evaluator

■ 0.31 · held

Contained and held back from the decision

And a slow or confusing claim becomes a lost customer: Deloitte's analysis of nearly 4,000 P&C customer responses found that the claims process ranked lowest of the drivers measured in creating positive sentiment, while human factors (customer support, knowledge, communication, and attitude) were the top four drivers (Deloitte).

Adjusters and underwriters know this, which is why they withhold trust from systems they cannot interrogate. Their skepticism is not resistance to change. It is a rational response to a system that cannot show its work.

WHAT IT TAKES

What trusted AI means in underwriting and claims.

Trusted AI is that which can demonstrate current performance as fact, whose decisions can be traced, whose authority is bounded, and whose errors make it better. Each property is operational, not aspirational, and each can be measured.

01 / KNOWN PERFORMANCE Evaluation that runs continuously, not once That means a governed evaluation set that reflects the real book: rare policy endorsements, litigated bodily injury files, unusual roof geometries, low-resolution or obstructed imagery, and post-catastrophe damage patterns. The set grows as new conditions appear. It also means monitoring for drift and confidence decay in live production, so a degrading model is caught before an actuary or a regulator catches it.	02 / TRACEABLE DECISIONS Every outcome can be reproduced When a model flags a claim for coverage review, codes an injury, or prefills a roof attribute, the carrier needs to know which document or image, which model version, which threshold, and which human, if any, produced that outcome. As BCG describes carriers redesigning workflows so that orchestrated AI agents handle core tasks, agent tracing and evaluation of multistep tool calls become as important as single-model accuracy. Traceability lets claims, compliance, and audit teams defend a decision after the fact, and lets a carrier explain an AI-influenced adverse action to a consumer.
03 / BOUNDED AUTHORITY Decision rights defined before autonomy goes live Straight-through processing within guardrails, human validation before payment or bind, escalation to a named role at defined thresholds, and a return to manual handling when coverage, severity, or fraud indicators appear. BCG describes the same shift, with humans refocused on validating outputs and resolving exceptions. These modes only work if the system can tell, file by file, which mode applies. That requires calibrated confidence scoring and risk-based routing, not a single global threshold.	04 / ERRORS THAT MAKE IT BETTER The property most deployments lack A trusted system treats every override, escalation, and correction as a learning signal. Validated corrections flow back into the evaluation set and the training set. Thresholds and rules are updated. The next occurrence of the same condition is handled correctly without human intervention. Over time, the share of files the system resolves straight through grows because it has earned that authority, case by case.

This aligns with the direction of regulation. The NAIC Model Bulletin requires insurers to maintain a written AI systems program covering how systems are selected, validated, monitored, and retired, to test systems before deployment and monitor them after, and to keep documentation that reflects current practice rather than aspirational policy ([Openlayer](#)). Insurers remain responsible for vendor-supplied model outputs; in the bulletin's framing, vendors cannot serve as a compliance shield.

The fourth property is what turns the first three from compliance overhead into a source of compounding value. It is also the one that depends most on the data underneath.

WHERE THE EVIDENCE LIVES

Imagery and claim files are the trust gap.

The hardest data problem in P&C AI is not the data in the policy administration or claims system. It is everything attached to it. Aerial, street-level, and drone imagery. FNOL emails and ACORD forms. Police reports, medical records, bills, and demand letters. Inspection reports, adjuster notes, and the reasoning behind every reserve change. This unstructured data carries most of the context a model needs to make a trustworthy decision, and most of it is never made AI-ready.

Core systems increasingly hold transactions well. Premiums, losses, expenses, and claim counts flow into the claims system, the data warehouse, and the actuarial models that produce frequency, severity, and combined ratio analysis. The evidence behind each of those numbers sits in attachments and image archives that the same platforms do not govern. CloudFactory's work with carriers, MGAs, TPAs, and imagery providers points to four patterns.

01 / IMAGERY Imagery is not yet usable data Carriers can access more aerial and street-level imagery than ever, but turning images into verified roof, structure, and hazard attributes is still largely manual. The image and the rating variable rarely meet.	02 / SELF-REPORTED DATA Self-reported data drifts Applicant answers and inspections from prior years age quickly as roofs wear, structures change, and rebuild costs climb. Rates end up resting on facts nobody has verified recently.
03 / CLAIM FILES Claim files are read once, by one person The chronology, injury coding, and coverage reasoning in a casualty file live in one adjuster's notes. Consistency across adjusters, offices, and TPAs depends on who picked up the file.	04 / SCALE Smaller carriers are left behind Small and mid-size carriers, MGAs, TPAs, and self-insured programs cannot fund the in-house computer vision, labeling, and claim AI teams that the largest carriers run.

This matters for trust because a model can only be evaluated, traced, and corrected against data that has been structured to support those tasks. An unlabeled image archive cannot become an evaluation set. An adjuster's undocumented reserve rationale cannot become a training example. The work of turning unstructured insurance data into AI-ready data is not a preprocessing step that ends at launch. It is the same work, performed continuously, that keeps a production model trustworthy.

Carriers do not have an AI accuracy problem. They have a verification problem, and trust is what it looks like when that problem is solved.

HOW CLOUDFACTORY CLOSES THE GAP

One loop, from claim file to trusted decision.

If trust is a verification problem, the solution is an architecture that captures knowledge at the exact moment AI is uncertain. That moment is the most valuable signal a production system generates, and most systems discard it. CloudFactory builds edge case detection, the automatic identification of data, inputs, or outputs a model cannot yet handle reliably, into the trust loop.

EXHIBIT 3 The trust loop



The trust loop starts by **knowing when AI is wrong**.

01

Automated error detection

AI Trust evaluates your AI against your ontology, policies, and rules to flag low-confidence extractions, edge cases, and likely non-compliance.

02

Intelligent routing

CloudFactory's AI Platform sends each flagged item to the Evaluator with the right expertise, weighted by confidence, claim value, coverage complexity, and regulatory risk.

03

Validation and remediation

Evaluators confirm the issue and correct the label, attribute, or rule. Corrective action is orchestrated through machines or people, whichever fits.

04

Continuous feedback

Validated examples enter the evaluation set first and the training set second. Thresholds update, and the knowledge graph grows.

Data Foundation

01

Built on data you can trust

Imagery, FNOL submissions, ACORD forms, medical records, and adjuster notes become AI-ready. AI does the work, a second AI layer scores it, and data analysts validate only where confidence falls short.

Model & Agent Evaluation

02

Proof it works, before it ships

Evaluation against your definitions and your book of business, with trace-level grading of agent reasoning, tool calls, and decision paths.

AI Trust

03

Catch what it gets wrong, prove you caught it

Underwriting and claims policies become enforceable guardrails, with full lineage across prompt, inference, tool call, and correction.

Managed AI Operations

04

We run it, at your SLA

Hosting, pipelines, and retrieval run together, with coverage up to 24x7 for catastrophe surges. Drift routes to a person, not an alert queue.

Human expertise is a confidence mechanism built into the system, not a labor pool bolted on.

The trust loop on a roof and in a casualty file.

Consider a carrier using aerial imagery to prefill roof condition at quote, and a casualty team reading demand packages at intake. The same loop makes both defensible.

The loop in a property workflow

After a hail season, the imagery provider refreshes coverage in a growth territory where stone-coated steel roofing is common, and the model begins returning lower-confidence condition scores on roofs that resemble weathered shingle under certain sun angles. In a typical deployment, those scores pass into prefill or trigger a blanket inspection order, and the issue goes unnoticed until loss experience or inspection spend reveals it.

With the loop in place, AI Trust detects the confidence shift within the affected territory. CloudFactory's AI Platform routes the uncertain scores by risk: low-value risks continue under existing guardrails, while high-value or hazard-flagged properties go to CloudFactory Evaluators. They confirm the material, verify condition, correct the attributes, and document the visual signature. The validated examples enter the evaluation set, so every future model version is tested against the condition, then the training set. Thresholds are adjusted, and the knowledge graph links the condition to the territory, the imagery source, and the roof type. The same verified record becomes the pre-loss baseline if a claim arrives later.

The loop in a casualty file

A plaintiff firm adopts a new demand package format, and the extraction model's confidence on treatment chronologies and billed amounts drops. The AI Platform routes the affected files, prioritizing those with severe injury indicators or approaching state notice deadlines. Evaluators build the verified chronology, confirm injury coding, and correct the extraction schema. The next demand from that firm is read correctly, and reserve inputs stay consistent across adjusters, offices, and TPAs.

Why the loop moves AI from pilot to production

The loop changes what production readiness means. Instead of asking whether a model is accurate enough to deploy, claims and underwriting leaders can ask a better question of their programs: whether the system is getting more reliable every week, and whether the carrier can prove it. That has a measurable answer.

It also gives a deployment path that adjusters, underwriters, and compliance teams can accept. A model can enter production in shadow mode, with every output reviewed. As validated corrections accumulate and evaluation confirms performance on the conditions that matter, the AI Platform narrows routing to the cases that still require human judgment. Autonomy expands where it has been earned and stays bounded where it has not. Trust is built in the open, with evidence, rather than assumed at launch.

TWO SOLUTIONS FOR P&C CARRIERS

Two solutions, one AI Platform.

CloudFactory applies the trust loop through two solutions built around the workflows where P&C carriers price and pay risk. Both run on CloudFactory's AI Platform, alongside existing core systems rather than replacing them, and are designed for carriers that cannot build the pipelines and AI claim tooling the largest insurers run in-house.

EXHIBIT 4 Two solutions, one AI Platform

Property Data Intelligence

Turns imagery into verified attribute data and delivers it into the underwriting and claims workbenches.

MARKET PRESSURES

- Rising claim severity as rebuild costs and weather losses climb
- Stale, self-reported data from applicants and old inspections
- Imagery that is not yet usable data
- Smaller carriers unable to build the pipelines large carriers run

HOW IT WORKS

- Ecosystem data: aerial, street-level, and drone imagery plus property records, via partners
- Attribute conversion: AI scores roof, structure, and hazard attributes; Evaluators verify edge cases
- Integration layer: APIs and connectors push verified attributes into carrier workbenches
- Workbench inputs: underwriting prefill, risk flags, inspection triage, and pre-loss claim baselines

OUTCOMES MEASURED

Pure loss ratio, prefill rate, inspections avoided, CAT triage time, and prior damage identified at claim

BUILT FOR

Property carriers, MGAs with property programs, TPAs, independent adjusters, and imagery and data providers

Casualty Claim Intelligence

An AI-assisted, human-verified solution that carries casualty claims from first notice through coverage, severity handling, and regulatory reporting.

MARKET PRESSURES

- Rising claim severity as loss costs climb while the market softens
- Regulatory pressure: notices, deadlines, and reporting on every file
- Smaller carriers unable to fund in-house AI claim tooling
- TPA margin squeeze as self-insured clients demand more

HOW IT WORKS

- Intake and coverage match: AI extracts every file, and loss facts are checked against policy terms and limits
- Severity management: chronologies, injury coding, and reserves within a proven adjusting workflow
- High-frequency automation: simple claims handled straight through; Evaluators review the exceptions
- Regulatory outputs: state notices, Medicare Section 111 reporting, and a complete audit trail

OUTCOMES MEASURED

Adjusting expense, severity leakage, cycle time, straight-through rate, and compliance exceptions

BUILT FOR

Small and mid-size carriers, TPAs, self-insured programs, and MGAs with claims authority

Source: CloudFactory. Benefit figures are being set in design-partner pilots.

What distinguishes both solutions from document automation and model monitoring tools is what happens after the AI is uncertain. Automation tools optimize for speed on the files they can read. Monitoring tools report when a model is drifting. The loop adds the step that makes a result defensible: expert verification inside the live workflow, with the correction returned to the system as permanent knowledge and every decision logged for audit.

CloudFactory's experience in roof and aerial annotation at scale, and its elastic 24x7 capacity for catastrophe surges, means the human side of the loop scales when the book does. At the critical SLA tier, low-confidence and high-risk outputs are resolved inside the live workflow in 1 to 5 minutes, across Data Foundation, Model & Agent Evaluation, AI Trust, and Managed AI Operations.

WHERE TRUSTED AI CHANGES OUTCOMES

Five domains where the loop pays off.

The loop applies wherever an insurance AI system makes decisions with financial or regulatory consequences and encounters conditions its training data did not cover. Five domains show the pattern most clearly.

EXHIBIT 5 Five domains where the loop changes outcomes

Domain	Where trust breaks today	What the loop captures	Outcome to measure
Property underwriting and prefill	Roof, structure, and hazard attributes rest on self-reported answers or aging inspections	Imagery-verified attributes, edge-case signatures by roof type and territory	Prefill rate, inspections avoided, and pure loss ratio by cohort
Catastrophe triage and property claims	Surge volumes outpace field capacity, and prior damage is hard to separate from new loss	Validated damage assessments linked to pre-loss baselines	CAT triage time, field visits avoided, and prior damage identified
Casualty intake and coverage	Extraction from FNOL, ACORD forms, police reports, and medical bills is fast but unverified against policy terms	Corrected extractions and coverage-match rules by form, firm, and line	Coverage issues caught at intake, straight-through rate, and cycle time
Severity and reserving	Injury coding, chronologies, and reserve calls vary by adjuster, office, and TPA	Expert-verified chronologies, injury codes, and reserve rationale	Severity leakage, reserve consistency, and prior-year development
Agentic claims workflows and regulatory reporting	Multistep agent decisions and filing obligations are hard to trace, evaluate, or defend	Graded agent traces, failure patterns, and corrected tool-call chains, with full logs	Compliance exceptions, escalation rate, and verified-run rate

Source: CloudFactory analysis

The common thread is that each domain produces a steady stream of exceptions. Without a loop, those exceptions are cost, recorded in the pure loss ratio, through current losses and future reserve development, and in the expense ratio. With one, they are the raw material of a more capable operation, and the knowledge graph they build is specific to the carrier's own book, which makes it difficult for competitors to replicate.

Five moves to make before the next renewal cycle.

None of these moves requires a new core system decision. Each changes how a carrier funds, measures, and governs AI so that pilots can earn their way into production.

01 Measure AI on pure loss and expense ratios, not the dashboard.

Tie every AI program to the ratio it directly moves: severity leakage, pricing accuracy, and reserve consistency in the pure loss ratio, and adjusting expense in the expense ratio. Those direct impacts are how AI influences the combined ratio. Replace the single launch accuracy gate with a trajectory: performance against a governed evaluation set, measured weekly, with a target for the share of files resolved straight through. Ask every program owner to show that curve, not a point.

02 Fund the exception path, not only the agent.

Budget for what happens when a model is uncertain: routing rules, review capacity for catastrophe surges, and the work of converting reviewed cases into training and evaluation data. A deployment with no exception path is a deployment designed to stay a pilot.

03 Make imagery and claim files an asset class.

Inventory the imagery, inspection reports, demand packages, medical records, and adjuster notes the organization already holds. Decide which should become labeled, searchable, and evaluable, and assign an owner. This is the foundation that core systems and data warehouses do not yet cover.

04 Write decision rights before autonomy goes live.

For each line and workflow, define which decisions run straight through, which require human validation, which escalate and to whom, and what triggers a return to manual handling, by claim value, injury severity, coverage complexity, and state. Then make sure the system can apply those rules file by file, and that the documentation satisfies the AI systems program regulators now expect.

05 Capture the judgment of your best adjusters and underwriters.

Every override by an experienced adjuster or underwriter is knowledge the organization is about to lose. Build the mechanism that turns those decisions into institutional memory, so the next file, the next office, the next TPA, and the next model start from what the best people already know.

CONCLUSION

Trust compounds. So does the gap.

The carriers that scale AI fastest will not be the ones with the fastest extraction at launch. They will be the ones whose systems become more trustworthy every week, because every uncertain decision is detected, reviewed, corrected, and returned as knowledge. That capability compounds. Each line and program that runs the loop strengthens the knowledge graph the next one inherits, and the gap between carriers who have built it and those who have not widens with every cycle.

THE GOAL

The unit of progress in P&C AI is not a file processed faster. It is a decision that holds up at reserve review, in the regulator's file, and at renewal.

Three ways in. One accountable partner.

WAYS TO ENGAGE

PATH 01

AI Advisory & Engineering

Forward Deployed Engineers (FDEs) embed with your team to design, build, and ship, transferring knowledge as they go.

PATH 02

AI Applications

Packaged AI that delivers a business outcome from day one, led by our flagship, Invoice Intelligence.

PATH 03

CloudFactory's AI Platform

The software behind the loop, run by your own team, with CloudFactory Evaluators on call for exceptions.

Forward deployed to build. Managed to run. One accountable partner across both.

SOURCES

BCG, [The AI-First Property and Casualty Insurer](#), March 2026. Cited for AI spending growth, the share of P&C insurers generating value at scale, operating cost reduction potential, and the shift of human roles to validating outputs and resolving exceptions.

Deloitte Center for Financial Services, [A poor claims experience can drive customers away. How can insurers use AI to help?](#), July 2026. Cited for drivers of P&C customer sentiment in claims, based on nearly 4,000 customer responses across 12 carriers.

AM Best, [Best's Market Segment Report: Rate Actions, Investment Gains Drive US Property/Casualty Insurance Segments' 2025 Results; Headwinds May Pressure Carriers in 2026](#), as reported by Business Wire, February 2026. Cited for combined ratio expectations, softening rates, and social inflation.

Openlayer, [NAIC AI Model Bulletin: What Insurers Must Prepare for in July 2026](#), July 2026. Cited for state adoption of the NAIC Model Bulletin and its program, testing, documentation, and third-party requirements.

CloudFactory, Property Data Intelligence and Casualty Claim Intelligence solution briefs, and CloudFactory analysis of P&C underwriting, claims, and combined ratio workflows, 2026.

ABOUT CLOUDFACTORY

CloudFactory is an end-to-end AI trust partner, ensuring high-value AI processes and decisions run on data you can trust. Since 2010, CloudFactory has combined AI technology and human expertise to help enterprises and AI technology companies develop, deploy, and operate reliable AI across industries where trust is critical to operations.

CloudFactory's AI Platform, our FDEs, and our global team of CloudFactory Evaluators cover the full AI lifecycle, from Data Foundation, Model & Agent Evaluation, and AI Trust to Managed AI Operations. More than 700 of the world's most ambitious AI companies trust CloudFactory to keep their AI accurate in production.

START THE CONVERSATION

See the trust loop running on your underwriting and claims workflows.

cloudfactory.com

Trusted AI at Scale.