ASAPP

eBook

Generative AI agent use cases for property and casualty insurance contact centers

Table of contents

Introduction 03

What is a generative AI agent 05

The shifting legal and regulatory landscape 06

Our methodology 06

Property and casualty insurance use cases 08

Automate policyholder service without compromising satisfaction 14

Customer service teams in property and casualty (P&C) insurance face a challenging mix of demands—complex policies, high-stakes claims, and policyholders who need support at some of the most stressful moments in their lives. At the same time, contact center leaders are under pressure to manage increasing volume and rising costs, while meeting policyholder expectations for accuracy, empathy, and speed.

Traditional automation has offered incremental gains in efficiency. But it's not enough. Deterministic flows can't handle the variability of claims scenarios, policy questions, and emotionally charged interactions.

Implementing automation across customer touchpoints—particularly in claims processing and servicing—can

reduce service costs by up to 30% and accelerate process cycle times by more than 40%[1].

That's why generative AI agents are quickly emerging as a transformative solution in property and casualty insurance. Unlike scripted bots, generative AI agents can interpret context, adjust in real time, and take secure, policy-aligned actions on behalf of the policyholder. By scaling complex service interactions without sacrificing compliance or empathy, they help insurers deliver a higher standard of support while managing costs.

1. McKinsey & Company. (2023). Automation in claims processing: Reducing turnaround time. ResearchGate.

ASAPP

Generative AI agents create new opportunities for insurers to:

- ✓ Resolve claims and policy issues faster while preserving accuracy and trust
- ✓ Maintain security, compliance, and empathy during sensitive interactions
- ✓ Deliver more consistent, reliable service experiences across every channel

With generative AI agents, insurers can deliver the kind of service policyholders want. And they can do it at scale without compromising the care, trust, and accuracy the industry demands. This guide presents use cases for generative AI agents that reduce costs, keep policyholders satisfied, and drive genuine value for insurance companies.

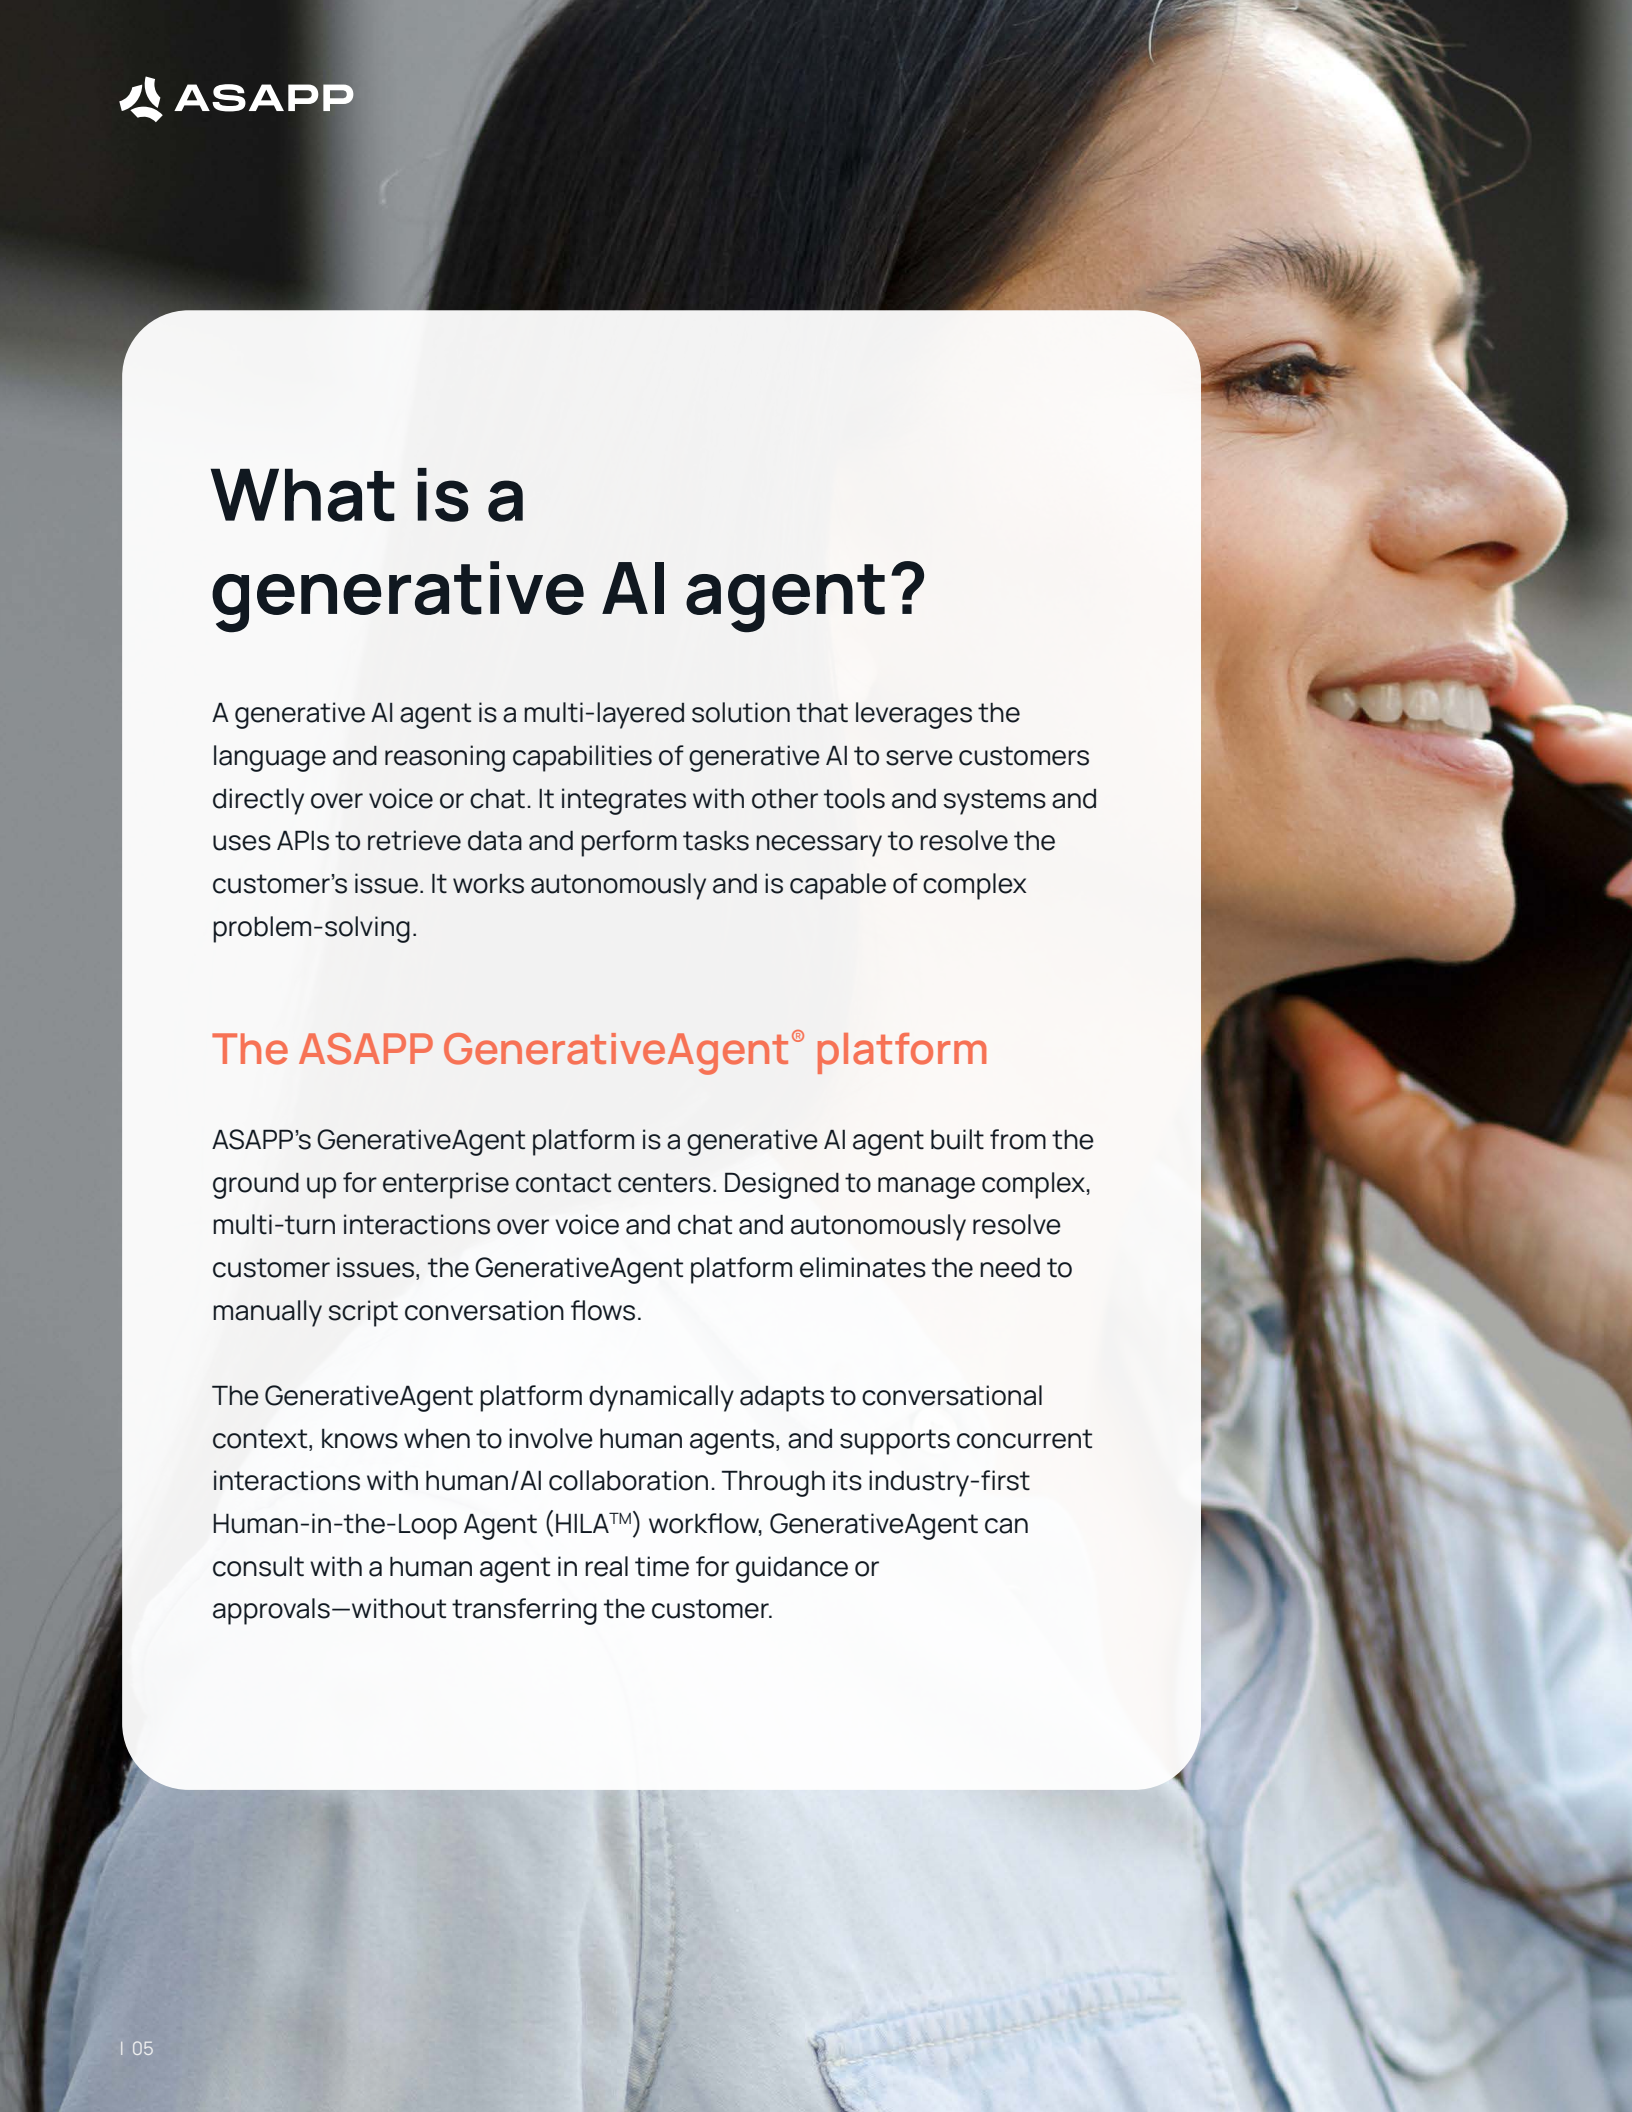

What is a generative AI agent?

A generative AI agent is a multi-layered solution that leverages the language and reasoning capabilities of generative AI to serve customers directly over voice or chat. It integrates with other tools and systems and uses APIs to retrieve data and perform tasks necessary to resolve the customer's issue. It works autonomously and is capable of complex problem-solving.

The ASAPP GenerativeAgent® platform

ASAPP's GenerativeAgent platform is a generative AI agent built from the ground up for enterprise contact centers. Designed to manage complex, multi-turn interactions over voice and chat and autonomously resolve customer issues, the GenerativeAgent platform eliminates the need to manually script conversation flows.

The GenerativeAgent platform dynamically adapts to conversational context, knows when to involve human agents, and supports concurrent interactions with human/AI collaboration. Through its industry-first Human-in-the-Loop Agent (HILA™) workflow, GenerativeAgent can consult with a human agent in real time for guidance or approvals—without transferring the customer.

The shifting legal and regulatory landscape

When evaluating generative AI agents for property and casualty insurance, compliance is non-negotiable. Data security and privacy are just the start. In many jurisdictions, AI agents must clearly disclose that they are AI and obtain customer approval before proceeding. You'll also need to consider the complex regulatory landscape and be certain that your AI agent won't run afoul of varied state licensing requirements. Regulations around the use of AI in customer service are evolving rapidly. Any AI solution you adopt must not only meet today's compliance standards but also adapt as the regulatory landscape continues to change.

Our methodology

With each use case, we've included an estimated deployment time, value drivers, and relevant metrics.

Deployment time

The deployment times here are estimates based on our experience deploying the GenerativeAgent platform and other AI solutions in enterprise contact centers. They represent typical durations from scoping to live production, derived from ASAPP benchmarks and industry studies. You'll want to keep in mind that your specific deployment time could vary depending on your CX technology infrastructure, the availability of your IT and development resources, the AI agent vendor you choose, whether you work with a system integrator or other strategic partner, and other factors.

With that in mind, the deployment time estimates should be viewed only as a guide to the relative ease and speed of implementing each use case.

- 2–4 weeks (Quick win)
- 1–2 months (Structured)
- 2+ months (Complex)

Value drivers

A successful AI agent deployment can drive value in a number of ways, affecting costs, revenue, operational efficiency, and customer satisfaction. The mix of value drivers will vary from one use case to the next.

For each use case included here, we've listed the value drivers that will impact your customer service operations:

- **Efficiency Gain:** Reduces average handle time (AHT), manual work, or after-call effort.
- **CSAT improvement:** Increases customer satisfaction through faster, clearer, or more consistent interactions.
- **Revenue gain:** Drives incremental sales via better cross-sell/upsell or conversion support.
- **Cost reduction:** Lowers operational expenses by automating high-volume or low-value interactions.
- **Quality assurance:** Improves compliance and consistency at scale, and reduces risk.

Relevant metrics

Real success with a generative AI agent depends on outcomes that have a positive and measurable impact on your business. So, your goals for any use case deployment should go far beyond the mere containment you might expect with legacy automation. The relevant metrics listed for each use case provide a starting point for measuring genuine business value.

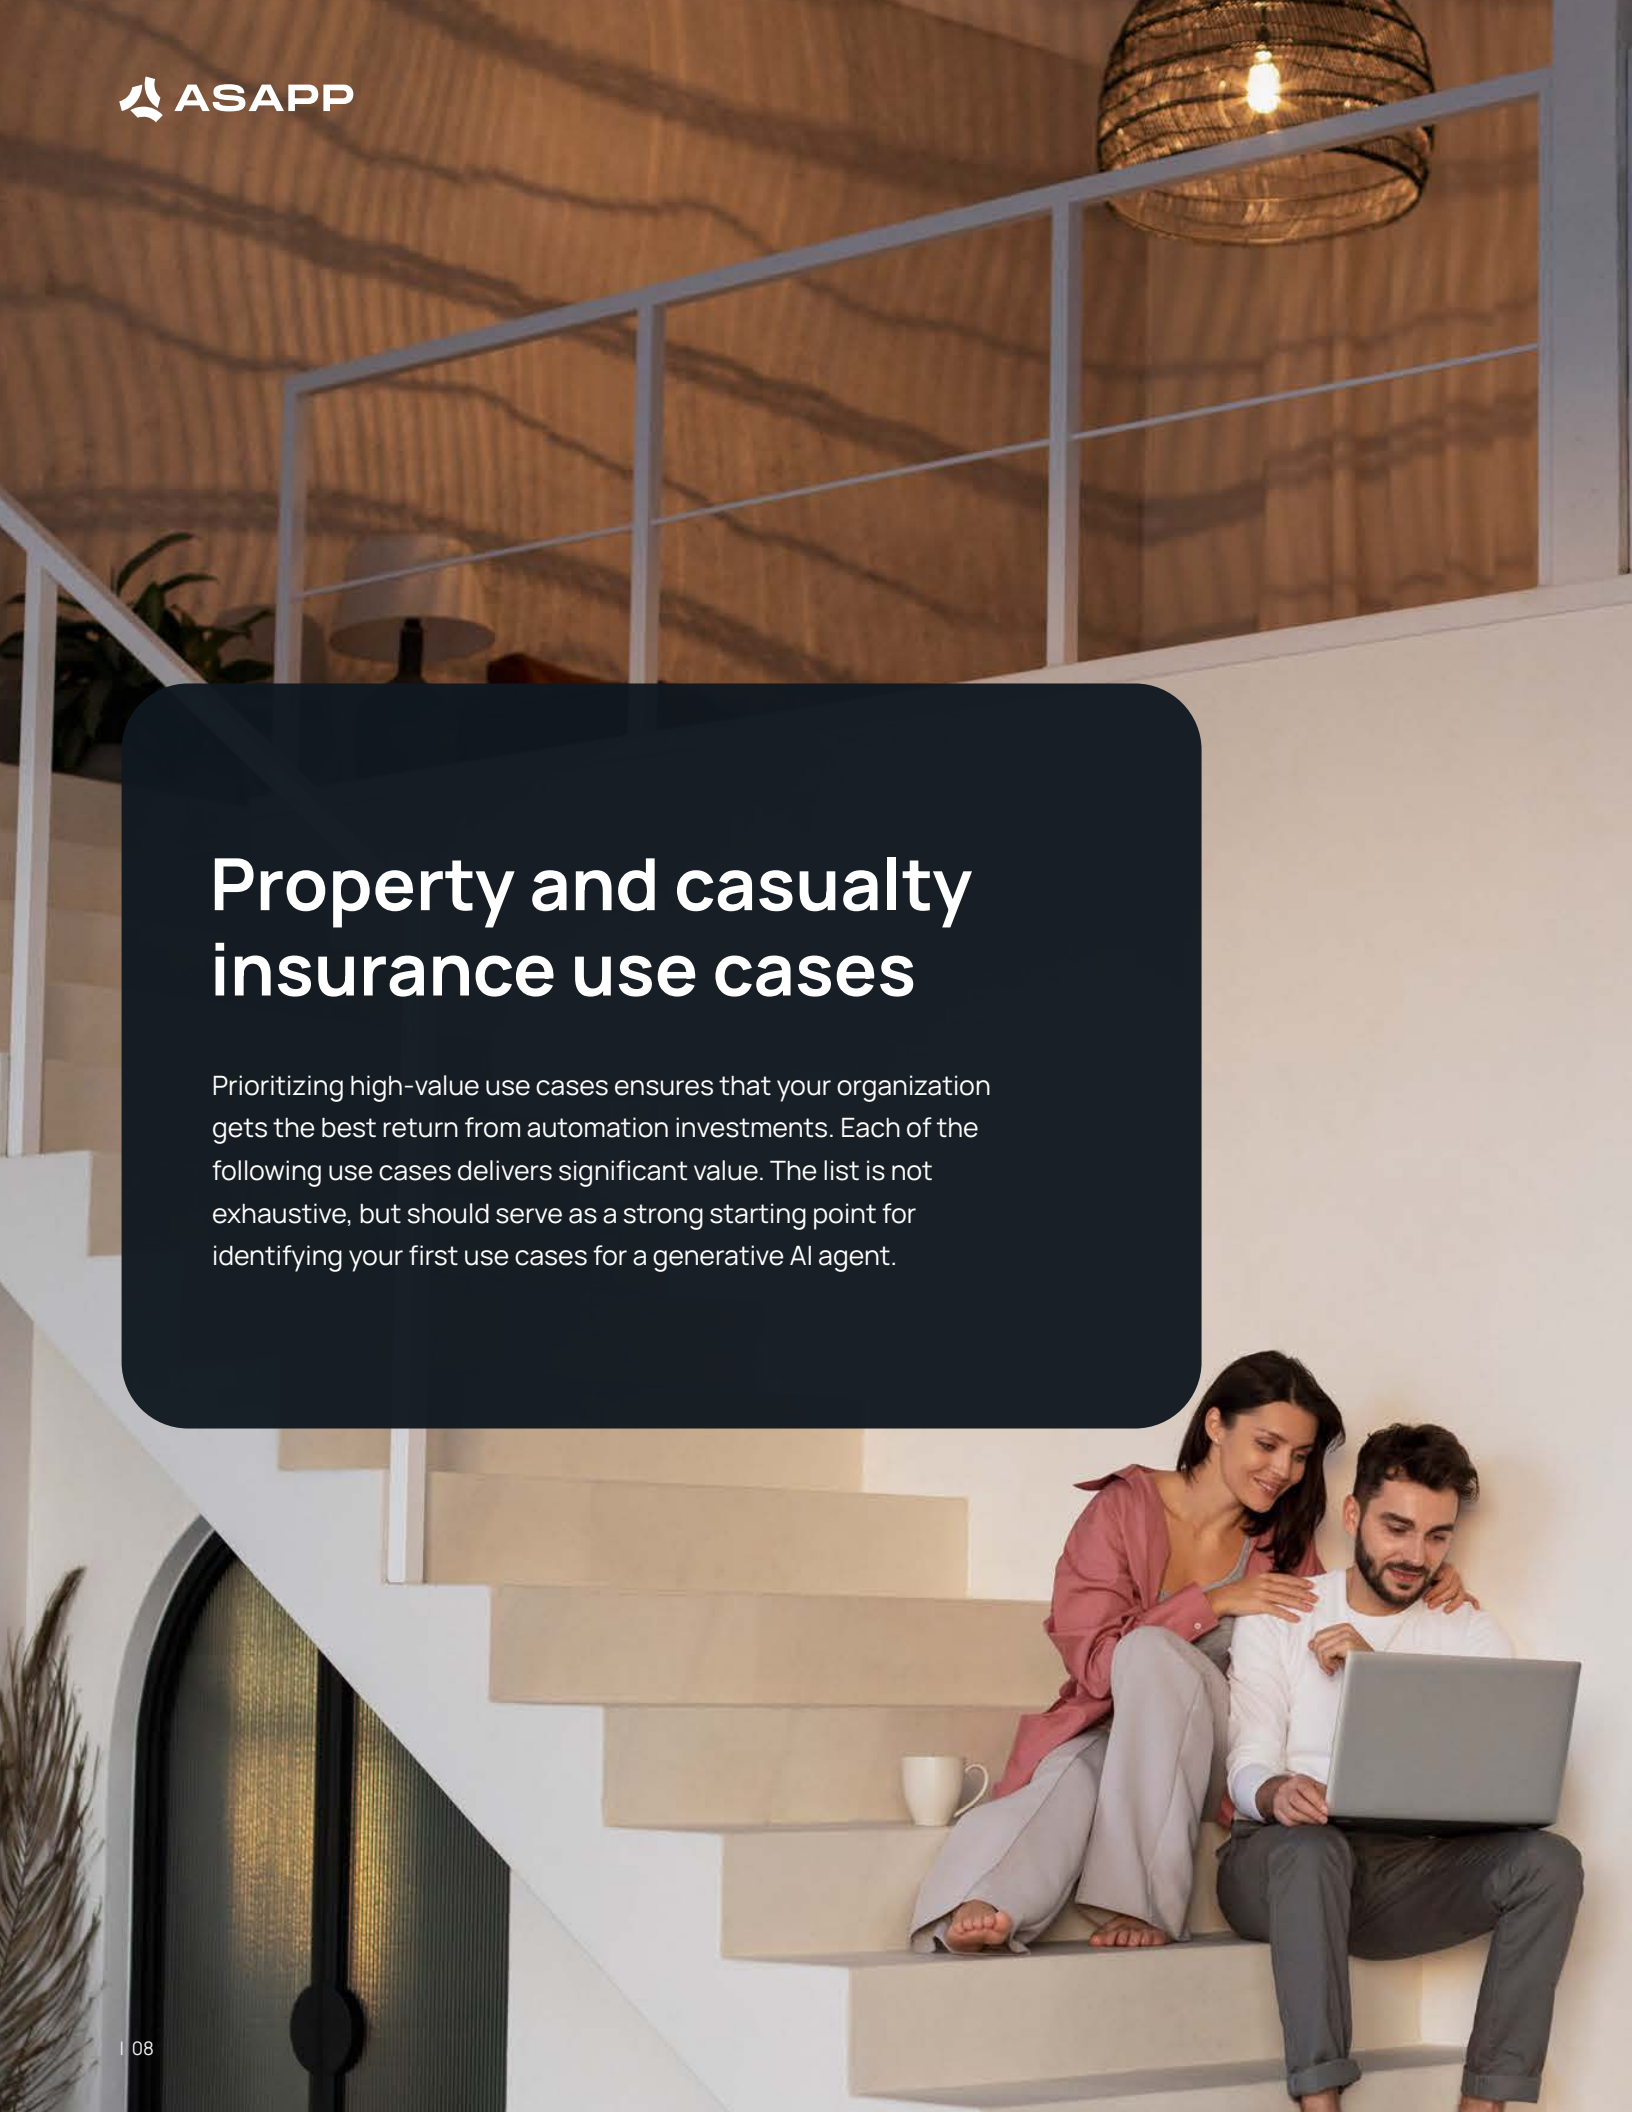

Property and casualty insurance use cases

Prioritizing high-value use cases ensures that your organization gets the best return from automation investments. Each of the following use cases delivers significant value. The list is not exhaustive, but should serve as a strong starting point for identifying your first use cases for a generative AI agent.

Policy information

When policyholders call with policy questions like coverage limits or deductibles, a generative AI agent retrieves and summarizes relevant policy details. Because the AI can retrieve the information instantly, it reduces the time policyholders wait on hold when compared to a service representative who might need some time to dig for the right information. This ensures policyholders get the information they need faster.

- **Deployment time:** 2–4 weeks
- **Value drivers:** Efficiency gain
- **Relevant metrics:** Average handle time (AHT) reduced, first contact resolution (FCR) improved

Billing inquiry

For calls about billing errors, premium due dates, or other billing-related inquiries, the generative AI agent can pull up the policyholder's billing history, investigate issues, and explain things like a recent rate change. Policyholders get clear answers quickly, without waiting on hold for a service rep.

- **Deployment time:** 4–6 weeks
- **Value drivers:** Efficiency gain
- **Relevant metrics:** High containment and resolution rate for billing inquiries

New policy quote

During cost inquiries, such as auto insurance quotes, a generative AI agent gathers initial information from the policyholder and suggests a coverage package. To ensure compliance with state licensing requirements, the AI agent can pass all information, including recommended coverage, to a licensed agent, if necessary. This accelerates the quoting process and improves the policyholder's experience with faster service.

- **Deployment time:** 4–6 weeks
- **Value drivers:** Efficiency gain, revenue gain
- **Relevant metrics:** Handle time for live agents reduced, sales conversion rate improved (faster quotes)

Accident assistance

Immediately after an accident, policyholders might call in distress. A generative AI agent for voice can autonomously walk them through crucial steps, such as safety checks or gathering information for a claim, and provide reassurance. For more complex cases, the AI agent can consult a human service rep for guidance. This human-in-the-loop mode of automation ensures timely guidance in emergencies where a human agent engages only as needed.

- **Deployment time:** 1–2 months
- **Value drivers:** CSAT improvement, efficiency gain
- **Relevant metrics:** Simple accident calls handled without full human involvement, policyholder feedback indicates improved post-accident support experiences

Coverage guidance

Policyholders often ask, "Am I covered if XYZ happens?" A generative AI agent can handle these FAQs by referencing policy documents and the customer's policy context. It asks clarifying questions as needed and provides personalized answers based on the policyholder's coverage. This 24/7 self-service reduces the need for live agents on basic questions, increasing convenience.

- **Deployment time:** 4–6 weeks
- **Value drivers:** CSAT improvement, cost reduction
- **Relevant metrics:** High call containment for coverage FAQs, increased customer satisfaction

Claims filing

When a policyholder reports a new claim, such as for an auto accident, the generative AI agent helps the claims representative or First Notice of Loss (FNOL) specialist by pre-filling claim forms and summarizing details from the conversation, such as location, time, and incident description. The human specialist can then validate and finalize the claim report. The generative AI agent speeds the FNOL process and ensures critical details aren't missed.

- **Deployment time:** 4–6 weeks
- **Value drivers:** Efficiency gain, quality assurance
- **Relevant metrics:** Claims intake time reduced, error rate in initial claim data decreased

Document submission

After an insurance claim is initiated, policyholders may need to submit documents, including photos or a police report. A generative AI agent for chat guides the policyholder through this process. For example, it can provide a secure link and instructions, and confirm when documents are received. The AI handles routine queries about document requirements and nudges the policyholder if something is missing, freeing up agents from follow-up calls.

- **Deployment time:** 4–6 weeks
- **Value drivers:** Efficiency gain
- **Relevant metrics:** Document submission cycle time reduced, claim processing delays due to missing docs decreased

Policy update

For simple policy changes, like an address change or adding a vehicle, a voice generative AI agent can authenticate the policyholder and process the update end-to-end. This frees customer service reps from routine data-change calls. A generative AI agent can be integrated with policy admin systems to execute transactions directly with minimal or no human oversight.

- **Deployment time:** 1–2 months
- **Value drivers:** Cost reduction
- **Relevant metrics:** High containment rate for simple requests, operational cost per policy update reduced

Fraud prevention

Policyholders often call in to check the status of claim reimbursements after they submit images of receipts, checks, or invoices for refunds or claims processing. Visual agentic AI detects digital manipulation, duplicate images, or AI-generated visuals and compares submissions against past transactions for anomalies. This prevents a generative AI agent from providing any information to a caller engaging in fraudulent reimbursements or insurance claim fraud.

- **Deployment time:** 4–6 months
- **Value drivers:** Quality assurance
- **Relevant metrics:** Decrease in fraud loss (value of fraudulent transactions prevented)

Fully automated claim

A generative AI agent handles most aspects of a straightforward claim—from first notice of loss (FNOL) to adjudication. Policyholders can report an incident, answer all questions through the generative AI agent, and can even receive an instant approval or payment for qualifying claims (e.g., small auto glass repair). Human adjusters only review edge cases or large claims. This transforms the claims experience with real-time service and drastically lower handling costs.

- **Deployment time:** 2–3 months
- **Value drivers:** Cost reduction
- **Relevant metrics:** High percentage of simple claims handled without human intervention, claim cycle time for minor claims cut from days to minutes, substantial cost savings from reduced manual processing

Compliance assistance

Insurance calls often require compliance disclosures. A generative AI agent can listen for required phrases and alerts and can provide the correct scripted language if necessary with accuracy. This ensures regulatory compliance at scale and reduces after-call corrections or even fines or penalties.

- **Deployment time:** 1–2 months
- **Value drivers:** Quality assurance
- **Relevant metrics:** 100% compliance script adherence (from ~95%); reductions in compliance penalties or QA failures

Policies and guidelines

Generative AI agents handling complex insurance queries, such as unusual coverage scenarios or policy exceptions, can search the insurer's knowledge base and underwriting guidelines. The AI quickly provides the relevant excerpt or summarized answer in understandable language, for faster resolution and consistent accuracy. This boosts policyholder satisfaction and contact center compliance.

- **Deployment time:** 4–6 weeks
- **Value drivers:** Efficiency gain, CSAT improvement
- **Relevant metrics:** Handle time reduced, high first-contact resolution (FCR)

Cross-sell recommendation

A policyholder contacts the insurer's chat agent to update their auto policy after buying a new car. While handling the update, the generative AI agent reviews the policyholder's profile and conversational cues to identify potential coverage gaps, such as no roadside assistance or umbrella coverage). It then offers relevant add-ons or complementary policies, explaining benefits in clear, compliant language. This turns routine service interactions into personalized sales opportunities without being pushy, increases policy bundling rates, and boosts customer lifetime value.

- **Deployment time:** 2-3 months
- **Value drivers:** Revenue gain, revenue retention
- **Relevant metrics:** Increased cross-sell conversion, incremental premium growth, higher retention for policyholders with multiple products

Complaints and escalation

When a policyholder is upset (e.g. disputing a claim decision), the AI monitors the call sentiment and content. It notices patterns in tone, pauses, even escalation triggers like mentioning a competitor. It also connects the dots across past interactions. That means the conversation starts from where the customer is in the moment, not where a script begins. This understanding of intent, sentiment and context all help to de-escalate the situation and speed a satisfactory resolution. The AI can also summarize the key complaint points for quick escalation to a manager if needed and for future analytics to improve performance. This results in more consistent, high-quality handling of complaints.

- **Deployment time:** 1–2 months
- **Value drivers:** CSAT improvement, cost reduction
- **Relevant metrics:** Improved FCR with complaint calls, supervisory escalations reduced due to AI-guided de-escalation

Multilingual support

A generative AI agent for voice that supports multiple languages lets policyholders get help in their preferred language without waiting for a specialist. For example, a Spanish-speaking policyholder can interact with the generative AI agent, which then either resolves the issue or hands off to a Spanish-speaking customer service rep with context. This broadens service reach by reducing language barriers.

- **Deployment time:** 1–2 months
- **Value drivers:** CSAT improvement, cost reduction
- **Relevant metrics:** Increased number of callers who successfully self-serve in non-English languages, CSAT increases among non-English callers due to reduced wait for language-specific agents

Emergency claims triage

Following natural disasters or other unexpected events that cause inbound volume to spike, a generative AI agent can gather preliminary claim information from callers, provide immediate guidance on next steps, and help to prioritize high-severity cases for fast-tracking by human adjusters. This ensures no policyholder has to wait to begin the claims process during a crisis.

- **Deployment time:** 1–2 months
- **Value drivers:** CSAT improvement, cost reduction
- **Relevant metrics:** High percentage of disaster-related claim intakes handled by the AI agents, time commitment for live agents dramatically reduced, improved policyholder confidence and satisfaction

Enhanced risk assessment and underwriting

When a homeowner contacts the insurer for a quote, a generative AI agent can dynamically conduct the underwriting interview—clarifying details, probing for risk factors, and ensuring data is complete and consistent before it reaches the underwriter. This results in faster underwriting turnaround by reducing back-and-forth, higher first-contact resolution through complete, accurate data capture, and a better customer experience with conversational, personalized guidance instead of rigid questionnaires.

- **Deployment time:** 2-3 months
- **Value drivers:** Efficiency gains, quality assurance
- **Relevant metrics:** Significant reduction in data collection time, increase in first-pass submission accuracy

Automate policyholder service without compromising satisfaction

Each of the use cases listed here demonstrates how a generative AI agent can automate policyholder interactions in property and casualty insurance contact centers, delivering benefits ranging from cost savings and efficiency gains to improved customer satisfaction, quality assurance, and new revenue opportunities. By selecting the right initial use cases and gradually expanding AI automation, insurance companies can modernize their policyholder service while tracking metrics to ensure each deployment delivers real value.

For more information on the ASAPP process for identifying the best use cases for your business, check out this guide:
Identifying the Ideal Use Cases for GenerativeAgent.

Get in touch with us

+1 (646) 386-8639

hello@asapp.com

www.ASAPP.com

About ASAPP

ASAPP is an artificial intelligence solution provider committed to solving the toughest problems in customer service. Our flagship product, GenerativeAgent®, is a platform built from the ground up to handle complex, multi-turn conversations with enterprise-grade performance, safety, and control. Because we automate what was previously impossible to automate, our AI-native® solutions deliver more than efficiency gains. They redefine the role of AI in the contact center and lay the groundwork for businesses to reimagine their customer experience delivery for the age of AI. Leading enterprises rely on ASAPP's generative and agentic AI solutions to dramatically expand contact center capacity and transform their contact centers from cost centers into value drivers. To learn more about ASAPP, visit **www.asapp.com.**