

eBook

Generative AI agent use cases for health insurance contact centers



Table of contents

| | |
|-------------------------------------------------------------------------|----|
| Introduction | 03 |
| What is a generative AI agent | 05 |
| GenerativeAgent® and the ASAPP CXP | 06 |
| The shifting legal and regulatory landscape | 07 |
| Our methodology | 07 |
| Health insurance use cases for a generative AI agent | 09 |
| Automate service without compromising compliance or member satisfaction | 14 |

Health insurance contact centers do much more than customer service. They're critical to patient access, satisfaction, care plan adherence, and even patient outcomes. It's complex work that comes with a multitude of challenges—highly personal and sensitive interactions, strict privacy regulations, and members who expect clear, timely answers when their health and well-being are on the line. At the same time, healthcare payors are under constant pressure to reduce overhead and manage costs without compromising members' access to care.

Traditional automation and basic AI chatbots have offered incremental efficiencies, but they struggle with the nuance of healthcare interactions, which can involve complicated coverage details, sensitive health data, and complex care plans.

Healthcare organizations deploying automation for routine tasks saw a

17% reduction

in average call wait time and up to a

22% decrease

in abandoned calls¹.

Generative AI agents are quickly emerging to fill the service gaps that exist between basic bots that can't handle complexity and human agents whose queues are often backed up, leading to long hold times. Unlike rigid bots, AI agents can interpret context, personalize responses, and securely act within care and benefits systems. They're capable of helping members navigate insurance coverage while maintaining HIPAA compliance and patient confidentiality.

¹. Healthcare IT News, "Healthcare Contact Center Automation Survey Results," 2024.

Generative AI agents create new opportunities for health insurance organizations to:

- ✓ Resolve member inquiries more quickly without sacrificing privacy or security
- ✓ Provide empathetic, consistent support tailored to each unique care or coverage scenario
- ✓ Reduce operational costs without eroding trust or quality of service

In short, generative AI agents unlock the ability to scale member support with the accuracy, safety, and empathy that healthcare requires.

This guide introduces some of the most impactful use cases for AI agents for health insurance member support. It's intended to help you choose use cases that will deliver measurable improvements in patient access, operational efficiency, and outcomes for both patients and your business.

What is a generative AI agent?

For contact centers, 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.

GenerativeAgent[®] and the ASAPP CXP

ASAPP's **GenerativeAgent** 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, GenerativeAgent eliminates the need to manually script conversation flows.

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

But GenerativeAgent is more than just a customer-facing AI agent. It's also the core of the **ASAPP CXP (Customer Experience Platform)**. The CXP brings every interaction, workflow, and customer signal into one intelligent system that resolves issues, enforces policies, and acts across enterprise systems. Unlike CCaaS or conversational AI tools that stop at simple deflection or routing, the CXP handles complex, multi-step workflows with accuracy, safety, and control while tailoring every step to the individual customer's context.

Leading enterprises use the ASAPP CXP to cut operating costs, accelerate resolution, modernize their CX stack, and build the foundation for an agentic enterprise where each member or patient has their own personalized AI agent.

The shifting legal and regulatory landscape

As you consider generative AI agents, you'll need to be mindful of a range of legal and regulatory compliance requirements. Data security and patient privacy requirements mandated by HIPAA and PIPEDA are just the beginning. Beyond that, various federal and state laws mandate clear disclosures when AI is used in patient interactions, including explicit consent in some cases. Several states also impose restrictions on how and where patient data can be stored and transferred.

The regulatory landscape is changing rapidly, with an eye toward transparency, bias prevention, and provider oversight. With that in mind, any AI agent solution you choose must enable you to adapt and maintain compliance as regulations evolve.

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, operational efficiency, and patient/member 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, and more consistent, personalized interactions.
- **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.

Health insurance use cases for a generative AI agent

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.

Coverage inquiries

Members often call to ask if a particular test or procedure is covered or whether they need a referral to see a specialist. A generative AI agent looks up the member's plan details and the relevant coverage rules to provide accurate, policy-compliant answers. This avoids the typically lengthy process of a human agent putting the member on hold to manually look up coverage details.

- **Deployment time:** 4–6 weeks
- **Value drivers:** Efficiency gain
- **Relevant metrics:** Time to answer insurance coverage queries reduced, fewer callbacks needed after verifying coverage (increasing FCR)

Billing and payment questions

Patients often call about medical bill charges, and set up payment plans, or make payments. A generative AI agent can handle these tasks by explaining bill components, or checking if insurance has paid its portion, and can even take a payment or schedule a payment plan. This automation provides quick service, while reducing the load on the billing staff.

- **Deployment time:** 4–6 weeks
- **Value drivers:** Efficiency gain
- **Relevant metrics:** ~30% of billing inquiries resolved by AI without escalation; average call length for bill explanations reduced (e.g., 8 min to 5 min); increase in on-time payments due to ease of self-service

Care access and coordination

A generative AI agent can help members navigate the healthcare system, from finding the right in-network provider to coordinating next steps. Members can describe their needs naturally (“My child needs a specialist for recurring ear infections nearby who speaks Spanish”), and the AI interprets intent, recommends the right type of care, checks availability, and helps schedule appointments. It can also close care gaps by reminding members of preventive visits or follow-up care. This improves access, streamlines coordination, and reduces staff workload.

- **Deployment time:** 4–6 weeks
- **Value drivers:** Efficiency gain, CSAT improvement
- **Relevant metrics:** High containment and faster resolutions for care access inquiries, improved care plan adherence

Multilingual member support

A generative AI agent capable of handling inquiries in multiple languages allows a health insurance contact center to serve diverse populations without staffing a full team for each language. The AI can converse in, say, English, Spanish, and Mandarin seamlessly. Human agents are on standby for medical issues or edge cases, but coverage questions in those languages are handled by the generative AI agent. This greatly expands accessible service and reduces wait times for non-English speakers.

- **Deployment time:** 1–2 months
- **Value drivers:** CSAT improvement, cost reduction
- **Relevant metrics:** High percentage of non-English calls resolved without a human translator, member satisfaction in non-English segments on par with English

Vaccine information

During flu season or public health campaigns, many people call with questions about which vaccinations are covered by their insurance plan. A generative AI agent can provide up-to-date answers about vaccine coverage, eligibility, and estimated costs. Because an AI agent scales instantly, it can handle high volumes of calls that might otherwise overwhelm staff, ensuring members get fast, accurate, and compliant information.

- **Deployment time:** 4–6 weeks
- **Value drivers:** Efficiency gain, quality assurance
- **Relevant metrics:** High first-contact resolution for vaccine queries, low rate of transfers to a human agent

Preauthorization status inquiries

Members frequently contact their health plan to check the status of a preauthorization for a procedure, test, or medication. A generative AI agent can securely access internal systems to provide real-time updates—whether a request is approved, pending, or requires additional information. It can explain next steps, such as what documentation is missing or when to expect a decision. This automation shortens wait times, prevents unnecessary callbacks, and frees human agents to focus on complex cases.

- **Deployment time:** 4–6 weeks
- **Value drivers:** Efficiency gain, cost reduction
- **Relevant metrics:** Average resolution time for preauthorization status queries reduced, fewer follow-up calls because patients get real-time answers

Telehealth coverage questions

Members often contact their health plan to understand what virtual care services are covered and what costs apply. A generative AI agent can answer questions about the type of telehealth services that are covered and whether there's a limit on the number of visits. It provides personalized answers based on the member's plan details, copays, and eligible providers. It can also guide members to in-network telehealth platforms or help them set up an appointment. This ensures members can access care quickly while reducing repetitive coverage calls.

- **Deployment time:** 4–6 weeks
- **Value drivers:** CSAT improvement, cost reduction
- **Relevant metrics:** High containment and resolution of telehealth inquiries, reduction in average handle time for telehealth coverage questions

Accumulators and cost estimation

Members sometimes struggle to understand how much they'll pay for care given their plan's deductible, out-of-pocket maximum, and copays. An AI agent can compute cost estimates for common services based on the member's accumulators, contracted rates, and network status, giving clear "what you will owe" answers ahead of scheduled care.

- **Deployment time:** 1–2 months
- **Value drivers:** Efficiency gain, CSAT improvement
- **Relevant metrics:** High containment and resolution for cost estimate interactions, reduction in post-service billing disputes

Automated preauthorization processing

Generative AI agents streamline preauthorization workflows by collecting patient details, pre-filling forms, and validating insurance requirements during inbound calls or chats. They guide patients and providers through necessary steps in real time, check documentation completeness, and escalate exceptions to human specialists. This reduces delays, minimizes manual effort, and accelerates patient care access—all while improving regulatory compliance in healthcare environments.

- **Deployment time:** 3–6 months
- **Value drivers:** Efficiency gain, cost reduction
- **Relevant metrics:** Significant reduction in per-transaction administrative costs, reduction in average processing time per preauthorization

Coordination with Medicare or other coverage

Many members have multiple coverage sources (Medicare, employer-sponsored, supplemental). They often call for clarification on which payer is primary for a claim or how a benefit will be split. A generative AI agent can access eligibility systems, documentation, and coordination-of-benefits rules to distill relevant information and provide accurate guidance in real time. It can also trigger notifications if primary coverage has changed, for example, after retirement, birthday, or other qualifying life event.

- **Deployment time:** 1–2 months
- **Value drivers:** Efficiency gain, CSAT improvement
- **Relevant metrics:** Reduction in misrouted claims and manual coordination of benefits (COB) case reviews, increased percentage of member inquiries answered without transfer

Chronic condition management guidance

Members newly diagnosed with chronic conditions, such as diabetes, heart disease, or asthma, often seek help understanding the condition, recommended care options, and covered services. An AI agent can deliver personalized educational guidance, remind members about required labs or specialist visits, and help them enroll in disease management programs offered by the insurer—all while confirming benefit coverage.

- **Deployment time:** 2+ months
- **Value drivers:** CSAT improvement, cost reduction
- **Relevant metrics:** Program enrollment uplift, decreased avoidable ER utilization among engaged members

Compliance and policy guidance

Health insurers must meet strict regulatory requirements, from HIPAA and PIPEDA privacy rules to mandated benefits and state-specific notices. Members sometimes have compliance-related questions about data sharing, privacy rights, or coverage mandates. An AI agent can deliver approved, compliant responses and personalized policy-specific explanations to ensure members remain informed while reducing compliance risk.

- **Deployment time:** 2–4 weeks
- **Value drivers:** Cost reduction, quality assurance
- **Relevant metrics:** Reduction in escalations to compliance and legal teams, improvement in audit quality reviews

Personal information and account updates

Members frequently need to update addresses, phone numbers, preferred pharmacies, communication preferences, or authorized representative information. An AI agent can authenticate the member, update demographic fields directly in core systems, and confirm changes immediately, bringing in a human agent for required approvals or policy-driven tasks as needed. This timely service prevents coverage gaps, returned mail, and compliance flags.

- **Deployment time:** 2–4 weeks
- **Value drivers:** Efficiency gain, CSAT improvement
- **Relevant metrics:** High containment and resolution for account update interactions

Pharmacy and drug benefit navigation

Members want to know if a medication is covered, what cost tier it falls under, the copay amount, and whether a generic or alternative medication is an option. An AI agent can run formulary checks, explain tiering rules in plain language, identify less costly covered alternatives, and direct members to in-network pharmacies.

- **Deployment time:** 1–2 months
- **Value drivers:** Cost reduction, CSAT improvement
- **Relevant metrics:** High containment and resolution for pharmacy benefit-related interactions, increase in generic or preferred drug utilization when appropriate

Open enrollment and plan selection

During open enrollment, call volumes surge with questions on benefit differences, premiums, and provider networks. An AI agent can provide personalized plan comparisons based on the member's utilization history and stated preferences, such as budget, care needs, and specialists they want to keep. Because it can collect and offer information, bring in a human as necessary, or even guide members through enrollment or coverage changes without human intervention, an AI agent can provide faster service for members and ease the burden on human agents during open enrollment.

- **Deployment time:** 2+ months
- **Value drivers:** Efficiency gain, CSAT improvement
- **Relevant metrics:** Reduction in open enrollment call backlogs, increase in self-service plan changes, improved member satisfaction during open enrollment season


Automate service without compromising compliance or member satisfaction

Each of the use cases listed here demonstrates how a generative AI agent can automate member interactions in health insurance contact centers, delivering benefits ranging from cost savings and efficiency gains to improved member satisfaction and quality assurance. By selecting the right initial use cases and gradually expanding AI automation, healthcare payors can modernize their member 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. Its AI-native Customer Experience Platform, powered by GenerativeAgent® integrates with existing systems and uses generative, personalized interaction to bring radical efficiency to every customer workflow. 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 agentic enterprise. ASAPP delivers personalized, context-aware interactions by giving every customer their own AI agent powered by their interaction history and enterprise data. 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.