

eBook

Generative AI agent use cases for **retail** 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
Retail use cases for a generative AI agent	09
Expand customer service capacity without compromising customer satisfaction	15

Customer service teams in retail face constant pressure – high volume, complex interactions, and ever-growing customer expectations. Meanwhile, customer service leaders grapple with the need to maintain service speed and quality without increasing operating costs.

Basic automation and AI-powered agent copilots have yielded modest improvements in efficiency and productivity. But these technologies can't keep up with ongoing budgetary pressures and increasing customer demands.

70% of customer journeys use self-service channels

but

only 9-14% of issues are resolved through self-service¹

That's why a growing number of retail customer service leaders are exploring the possibilities of generative AI agents. Unlike rigid bots that run on deterministic flows, generative AI agents can adapt to the customer conversation, make contextual decisions, and take actions to resolve customers' problems.

1. Gartner, Inc. (2024, August 19). Survey finds only 14% of customer service issues are fully resolved in self-service. Gartner Newsroom.

Generative AI agents expand capacity without compromising customer satisfaction.

- ✓ Scale effortlessly when demand fluctuates
- ✓ Deliver a more consistent customer experience
- ✓ Reduce the cost to serve

In other words, generative AI agents offer a real solution to the challenges of retail contact centers—as long as you choose the right use cases.

This guide is your starting point for enterprise use cases that drive measurable results.

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.

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 legal and regulatory compliance issues. Data security and privacy are just the start. In some jurisdictions, the agent must disclose that it's AI and specifically ask for the customer's consent to continue. In some countries, all customer data must reside in that country and cannot be transferred elsewhere. These regulations are still evolving and expanding into new states and countries. So, you'll need to be sure any AI agent you're considering will enable you to adapt and maintain compliance even as the regulatory landscape changes.

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:

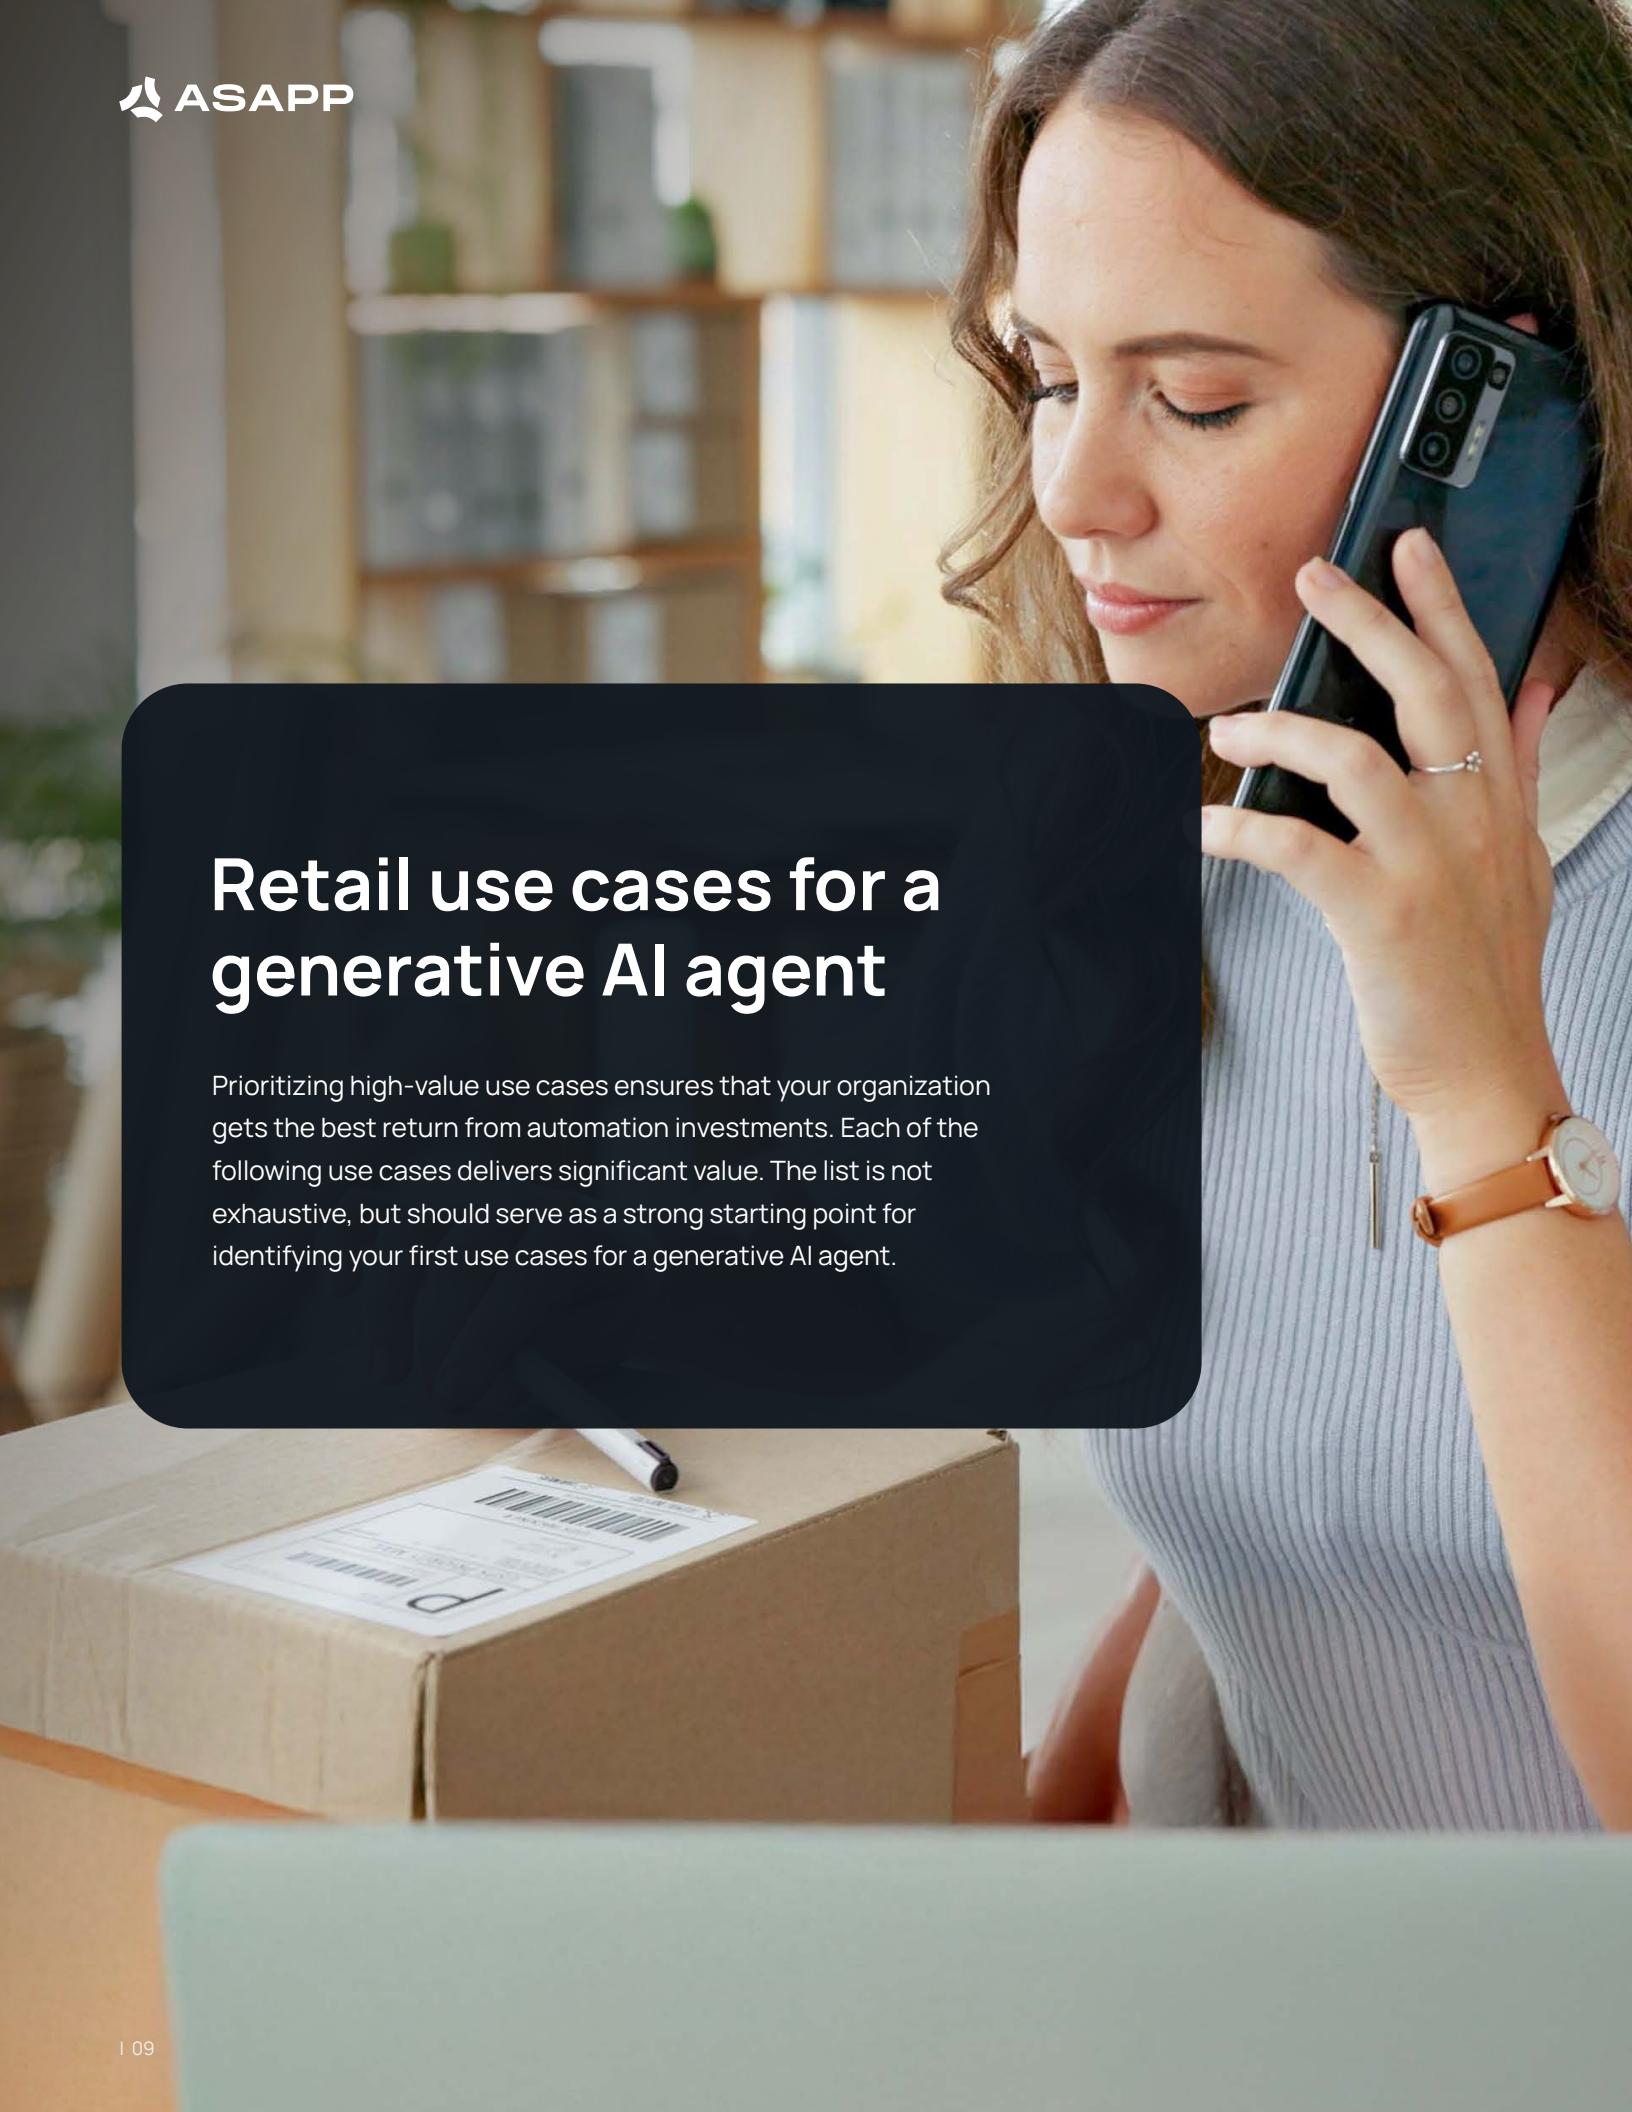
- **Efficiency Gain:** Reduces average handle time (AHT)
- **CSAT improvement:** Increases customer satisfaction through faster, clearer, and more consistent, personalized interactions.
- **Revenue gain:** Drives incremental sales via better cross-sell/upsell or conversion support.
- **Cost reduction:** Lowers operational expenses by automating high-volume or low-value interactions.
- **Quality assurance:** Improves compliance and consistency at scale, and reduces risk.
- **Revenue retention:** Helps reduce churn through improved save rates and proactive service offers.

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.

Retail 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.



Order delivery issues

Customers frequently contact retailers asking, “Where is my order?” (WISMO). A generative AI agent integrated with order management can instantly provide shipping status, tracking numbers, and delivery ETA when given a name or order number. And when delivery is delayed or a package is lost, the AI agent can go beyond basic order tracking to troubleshoot the issue and provide an on-policy resolution. This quick, 24/7 on-demand service boosts customer confidence in the retailer.

- **Deployment time:** 4–6 weeks
- **Value drivers:** CSAT improvement, cost reduction
- **Relevant metrics:** High containment for complex WISMO queries, reduction in live agent volume, improvement in CSAT for delivery information interactions

Returns and refunds

When a customer has questions about initiating a return, a normally simple process can get complicated quickly. A generative AI agent can guide the customer through return eligibility questions (“Within 30 days? Original packaging?”), provide a return merchandise authorization (RMA) or return shipping label via email, and trigger a refund process once it’s in transit or received—all without human intervention. For ineligible returns, it politely explains the policy or offers alternatives (store credit, exchange). This provides immediate resolution to customers on return requests without burdening live agents.

- **Deployment time:** 1–2 months
- **Value drivers:** Cost reduction
- **Relevant metrics:** A large portion of return requests handled end-to-end by the AI agent, support cost for returns down significantly

Conversational shopping assistants

A generative AI agent can act as a virtual shopping assistant on websites, mobile apps, or messaging channels to help customers find products, compare options, and answer questions in natural language. For instance, a customer might say, “I need a gift for my 10-year-old nephew who loves science.” The AI agent will ask a few questions to narrow down the options and then recommend some items with reasoning, such as, “This chemistry set is popular for that age.” By mimicking the conversational approach of an in-store sales associate, the AI agent increases engagement and basket size. It can also explain product features, upsell complementary products, and offer bundled deals. This real-time support enhances the online shopping experience and drives sales.

- **Deployment time:** 2+ months
- **Value drivers:** Revenue gain, cost reduction
- **Relevant metrics:** Decrease in abandoned carts, increase in average order value with AI-driven upselling/cross-selling, high containment rate

Product information

Shoppers often have questions about product specifications, compatibility, or usage before purchase. For example, "Is this laptop compatible with XYZ software?" or "Does this dress have pockets?" A generative AI agent for voice or chat can answer them in plain, conversational language by accessing product descriptions, manuals, and prior Q&A. It can also clarify policies for warranties, returns, and other customer concerns in context. Instant answers reduce purchase anxiety and prevent calls or emails from reaching live agent queues.

- **Deployment time:** 4–6 weeks
- **Value drivers:** CSAT improvement, revenue gain
- **Relevant metrics:** Reduction in pre-sale support volume for live agents, higher conversion rate with frictionless support, improved CSAT

Store information and inventory

Customers often call to ask about store hours and whether a specific item is in stock at a nearby location. A generative AI agent can check the inventory system and respond with accurate information. It can even place the item on hold or help the customer complete the purchase for a quicker pick-up at the store. This gives customers quick answers and helpful real-time service without tying up live agents or in-store sales associates.

- **Deployment time:** 2–4 weeks
- **Value drivers:** Efficiency gain, cost reduction
- **Relevant metrics:** High containment for inventory inquiries, measurable conversion rate for product holds and purchases

Loyalty program support

The value of a loyalty program depends on customers using it. A generative AI agent can boost program engagement with convenient 24/7 conversational support. They can help customers manage their accounts, check their balances, understand program policies, and even apply discounts and redeem rewards. Instant gratification and clarity around program rewards encourage program engagement and keep loyal customers happy.

- **Deployment time:** 4–6 weeks
- **Value drivers:** CSAT improvement, cost reduction
- **Relevant metrics:** Measurable increase in loyalty program engagement, decrease in loyalty program inquiry volume for live agents, increase in CSAT among loyalty members

Post-purchase support

A generative AI agent can walk customers through product setup, maintenance, troubleshooting, or warranty claims. It can also escalate to video tutorials or to a human agent for detailed technical support. This results in reduced product returns, increased satisfaction, and fewer support calls.

- **Deployment time:** 4–6 weeks
- **Value drivers:** Efficiency gain, cost reduction
- **Relevant metrics:** High containment, reduction in returns attributed to setup/usage issues, faster time to resolve warranty claims

Order modification

If a customer needs to change an online order (shipping address, quantity, or cancel an item shortly after ordering), a generative AI agent can handle the common parts automatically and involve a human agent if needed. It might automatically check if the order has shipped. If not, it can execute the address change or cancellation and confirm to the customer. If it's too late (order already shipped), it informs the customer and offers a return process. The AI reduces agent involvement by doing the system checks and only involving agents for edge cases or approvals.

- **Deployment time:** 4–6 weeks
- **Value drivers:** Efficiency gain
- **Relevant metrics:** AHT for order change and cancellation calls reduced, percentage of modifications done automatically (prior to shipping) high, leading to fewer manual interventions, CSAT improvement

Warranty claim

For products under warranty, customers can initiate a claim through a generative AI agent for chat or voice. Provided with a product serial number or order number, the AI agent verifies warranty coverage, captures issue details ("My device won't power on"), and if eligible, sets up a repair or replacement. It generates shipping labels or directs the customer to a service center, and provides a return merchandise authorization (RMA) number. Automating this process speeds up service for the customer and reduces the burden on warranty support teams.

- **Deployment time:** 1–2 months
- **Value drivers:** Cost reduction
- **Relevant metrics:** High percentage of basic warranty claims resolved by the AI agent, lower support cost per claim processed

Price match inquiry

Some retailers offer a price match guarantee. A customer might contact support saying, "I found this item cheaper at Competitor X." A generative AI agent can navigate that conversation by asking for the competitor price and a link, then verifying the information. If it meets the criteria for a price match based on company policy, the AI agent can apply a price adjustment or provide the customer with a promo code for the difference. If it's not eligible, the AI agent can clearly and politely explain why per policy. Automating this inquiry makes the process quick and fair, enhancing customer perception that the retailer honors its promises.

- **Deployment time:** 4–6 weeks
- **Value drivers:** CSAT improvement, cost reduction
- **Relevant metrics:** High containment and high CSAT for price match inquiries, fewer of these inquiries in live agent queues

Payment issue

An AI agent can troubleshoot and resolve a range of payment issues, like duplicate charges or discounts that were not applied correctly. The AI agent can quickly access the transaction and verify the error by comparing system records. Then, it can apply a refund, offer store credit, or apply the discount according to company policy. Customers get quick resolutions, and the AI agent keeps routine payment issue inquiries out of live agent queues.

- **Deployment time:** 4–6 weeks
- **Value drivers:** Efficiency gain, cost reduction
- **Relevant metrics:** High containment and FCR for routine payment issues, decrease in escalations to live agents, improved CSAT

Gift card inquiry

When a customer has questions or runs into problems with a gift card, basic automation isn't enough. A generative AI agent can troubleshoot and resolve issues with balance discrepancies, lost cards, fraudulent usage, and cases in which the customer reports that the card is not working. It can also answer questions about gift card terms, such as when it expires or where it can be used. Automating these inquiries can considerably reduce live agent volume during holiday seasons when gift card use is high.

- **Deployment time:** 2–4 weeks
- **Value drivers:** Efficiency gain
- **Relevant metrics:** High containment for gift card inquiries, peak-season support load reduced

Delivery rescheduling

For scheduled deliveries, like furniture and appliances, or signature-required packages, customers sometimes need to change the delivery date or time window. A generative AI agent can handle this request by checking available delivery slots and rescheduling accordingly, or coordinating a hold at a pickup location. It updates logistics systems and confirms the new schedule with the customer. This flexibility reduces failed delivery attempts and ensures customers receive orders at a convenient time without multiple calls.

- **Deployment time:** 4–6 weeks
- **Value drivers:** CSAT improvement, cost reduction
- **Relevant metrics:** Reduction in failed deliveries, improved CSAT for the delivery process

Scaling for seasonal volume

During seasonal spikes (e.g., Black Friday or holiday shipping disruptions), generative AI agents can preemptively deflect common queries like order status, cancellations, or policy questions. A generative AI agent scales instantly to handle these high-frequency topics, reducing agent load while maintaining response speed and accuracy. Enterprises can handle 60–80% of surge-related volume automatically, reduce the need for temporary staff hiring or overtime, and maintain response quality and CSAT during peak periods.

- **Deployment time:** 1–2 months
- **Value drivers:** Cost reduction
- **Relevant metrics:** Improved containment of peak-period interactions, lower seasonal staffing requirements, maintained or improved CSAT during high-demand windows

Product recalls and incident response automation

Contact centers experience sudden volume spikes during recalls, safety advisories, or class-action events. Generative AI agents can absorb this surge by automating common questions about eligibility, next steps, and documentation requirements. They deliver consistent, compliant messaging while triaging cases that require human escalation.

- **Deployment time:** 1–2 months
- **Value drivers:** Cost reduction, quality assurance
- **Relevant metrics:** High containment and FCR for recall-related inquiries, reduction in response time during incident surges, improved compliance through standardized responses and documentation handling

Expand customer service capacity without compromising customer satisfaction

Each of the use cases listed here demonstrates how a generative AI agent can automate customer-facing interactions in retail 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, retailers can modernize their customer 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.