

eBook

Generative AI agent use cases for financial services contact centers



Table of contents

Introduction	03
What is a generative AI agent	05
The shifting legal and regulatory landscape	06
Our methodology	06
Financial services use cases for a generative AI agent	08
Automate customer service without compromising customer satisfaction	15

Customer service teams in financial services face a mix of complex challenges—complicated products, strict compliance requirements, and customers who expect fast, accurate support when it matters most. At the same time, contact center leaders are under pressure to manage rising volumes and keep costs under control without sacrificing trust or service quality.

Traditional automation and AI-powered copilots have made minor efficiency gains at the margins. But those deterministic solutions break down when faced with the depth and variability of customer service needs in a large bank or financial institution.

Automation in financial services delivers
25% decrease
In cost to serve¹

That's why generative AI agents are quickly making inroads in financial services contact centers. Unlike rules-based bots, generative AI agents can understand context, adapt in real time, and take secure actions on behalf of the customer. Their ability to handle complex scenarios at scale without sacrificing compliance translates into better customer service with lower costs.

1. DECTA, "Banking Automation in Digital Banking," 2025: "Automation reduces banking costs by up to 20–30%. Modern fraud detection systems significantly improve banking security."

Generative AI agents create new opportunities for financial services organizations to:

- ✓ Resolve customer issues faster while maintaining compliance guardrails
- ✓ Maintain security, safety, and compliance while navigating nuanced situations
- ✓ Deliver a more consistent, trustworthy experience across every interaction

In short, generative AI agents unlock the ability to serve customers at scale without diluting the security and care that financial services demand.

This guide introduces the most valuable use cases for AI agents in financial services—so you can prioritize the ones that deliver measurable impact.

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 financial services, compliance is non-negotiable. Data security and privacy are only the foundation—firms must also consider strict financial regulations, auditability, and customer consent requirements. In many jurisdictions, AI agents must clearly disclose that they are AI and obtain customer approval before proceeding. 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.

Financial services 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.

Account balance and transaction inquiries

Customers frequently call banks for simple queries like checking account balances or recent transactions. Simple automation can already handle this task well. But a conversation that starts with checking an account balance can easily move on to more complex topics, so it's a good idea to consider this use case for a generative AI agent. A generative AI agent can adapt as the conversation shifts to other intents. This fully automates routine inquiries that would otherwise occupy live agent time, and it's available 24/7.

- **Deployment time:** 4–6 weeks
- **Value drivers:** Cost reduction
- **Relevant metrics:** High containment rate for multi-intent conversations that start with balance inquiries.

Lost/stolen card reporting

When a customer needs to report a credit/debit card lost or stolen, time is critical. A generative AI agent can securely verify the customer, immediately freeze the card, and initiate a replacement card shipment – all without waiting for a human. It can also answer common questions about liability for fraudulent charges and other related topics, reassuring the customer.

- **Deployment time:** 4–6 weeks
- **Value drivers:** Efficiency gain, cost reduction
- **Relevant metrics:** High percentage of card loss reports fully handled by AI (no live agent needed) with faster average speed to answer, shorter handle time, and quicker resolution.

Fraud alert verification

Banks often contact customers to verify potentially fraudulent transactions. A generative AI agent can either call the customer or handle inbound responses, asking them about recent charges. The AI collects the information and either clears the alert or initiates the fraud dispute process. Human-in-the-loop agents assist with exceptions and complex cases. This speeds up fraud resolution and reduces false declines.

- **Deployment time:** 1–2 months
- **Value drivers:** Efficiency gain
- **Relevant metrics:** Fraud verification calls resolved faster 24/7, high resolution rate, low transfer rate for fraud alert interactions.

Fraud detection and escalation

The AI instantly flags suspicious activity during support calls or digital chats, automatically verifies customer identity with adaptive questioning, and triggers alerts to human specialists. By acting as an early-warning system, generative AI agents protect both customers and institutions, helping meet evolving fraud-prevention requirements.

- **Deployment time:** 2–3 months
- **Value drivers:** Quality assurance
- **Relevant metrics:** High fraud detection accuracy, mean time to respond reduced.

Password reset troubleshooting

For most customers locked out of their accounts, simple automation can fix the problem. But for some, a call or chat is needed to resolve the issue. A generative AI agent can authenticate the user and guide them through resetting their password or unlocking the account. This extra guidance can be especially helpful for customers who were unable to reset their password successfully with the standard self-service option. A generative AI agent can improve the customer experience by quickly restoring access, while freeing human staff to handle more complex requests.

- **Deployment time:** 4–6 weeks
- **Value drivers:** Cost reduction
- **Relevant metrics:** High resolution rate and shorter handle time for interactions related to password reset issues.

Credit card bill and payment issues

Customers may have questions about their credit card bill and payment options, including whether they can get an extension or pay less than the minimum amount this month. For simple billing inquiries, a generative AI agent can answer questions and process payments. For other requests, like modified payment amounts or due dates, it can follow your company policies and request guidance from a human-in-the-loop for approval, then advise the customer if granted and set it up for them. Automating these high-volume transactions and requests improves efficiency.

- **Deployment time:** 4–6 weeks
- **Value drivers:** Efficiency gain
- **Relevant metrics:** High containment for billing-related interactions, reduction in live-agent-handled payment calls.

Product recommendation

A generative AI agent on a bank's website or mobile app helps customers choose financial products. For example, with a customer who needs help choosing the right credit card, the AI agent could ask about spending habits, whether the customer typically carries a balance or pays in full each month, what types of rewards they would prefer, and what their primary goal is for getting the card (building credit, earning rewards, etc.). By combining this familiar conversational approach with customer data, the AI can recommend the credit card that would best meet the customer's needs. This drives digital sales with a guided, personalized experience.

- **Deployment time:** 1–2 months
- **Value drivers:** Revenue gain
- **Relevant metrics:** Increased conversion rate for online product inquiries.

KYC (Know Your Customer) document collection

If a customer's account requires KYC updates or additional documents (ID proof, etc.), the generative AI agent assists in the outreach and collection. For inbound calls, the generative AI agent sends the customer a secure link and instructions to upload documents, confirms receipt, and can check completeness. The AI understands which required documents are still pending and can remind the customer. This streamlines compliance-related document gathering without lengthy call handling by a human agent.

- **Deployment time:** 4–6 weeks
- **Value drivers:** Efficiency gain, quality assurance
- **Relevant metrics:** Turnaround time for KYC compliance reduced, significant decrease in agent time spent on chasing documents.

Loan application status

Applicants often call to check the status of a loan or mortgage application. A generative AI agent, integrated with the bank's loan system, can provide real-time updates about the application's status and whether the customer needs to take any action, like submitting additional documentation. It can also answer questions about next steps. This speeds status updates to customers to reduce their anxiety and reduces call volume for agents, so they can focus on more complex inquiries.

- **Deployment time:** 4–6 weeks
- **Value drivers:** Efficiency gain
- **Relevant metrics:** High containment for loan status queries (no human needed), improved NPS or CSAT from loan applicants.

Transaction dispute

When a customer calls to dispute a charge on their account, the generative AI agent collects details (charge date, amount, merchant, etc.), asks the required questions ("Did you attempt to resolve with the merchant?"), ensures all necessary information is captured, and initiates the dispute process. This results in faster service with high accuracy and ensures fully traceable compliance with dispute procedures.

- **Deployment time:** 4–6 weeks
- **Value drivers:** Efficiency gain, quality assurance
- **Relevant metrics:** Time to log a dispute case reduced with shorter handle time and no waiting for the next available agent, error/omission rate in dispute filings near zero (AI prompts ensure complete information).

New account onboarding

When a new customer calls to set up a new account, or has questions about the process, a generative AI agent can provide answers and guide customers through onboarding. For example, if a customer just opened an account online and is calling to finish setup, the AI can help them set up direct deposit, go paperless, or download the mobile app. This conversational support speeds up onboarding and ensures a smooth first experience.

- **Deployment time:** 4–6 weeks
- **Value drivers:** CSAT improvement, efficiency gain
- **Relevant metrics:** Faster setup of new accounts and features, high containment for onboarding inquiries, improved CSAT for new customers.

Mortgage process

Mortgage applicants often have questions about the process (“What does conditional approval mean?” or “How do I lock my rate?”). A generative AI agent can give instant answers to a wide range of questions based on internal mortgage guidelines and educational content. It can explain complex mortgage concepts and financial terms (like debt-to-income ratio) in simple language, ensuring consistency. Customers get clear answers quickly, improving their experience during a stressful process.

- **Deployment time:** 1–2 months
- **Value drivers:** CSAT improvement, Efficiency gain
- **Relevant metrics:** High containment and FCR for mortgage-related questions, improved satisfaction with the mortgage application process.

Wealth portfolio inquiries

In wealth management or brokerage services, clients may call with complex questions about their portfolio or market conditions. A generative AI agent quickly summarizes the client’s portfolio performance, recent market news, or relevant research, suggests potential rebalancing options or products (within compliance limits and suitability rules), and can schedule an appointment with an advisor to discuss.

- **Deployment time:** 1–2 months
- **Value drivers:** Efficiency gain, CSAT improvement
- **Relevant metrics:** Good containment for wealth portfolio inquiries, improved customer satisfaction with quick resolutions and no hold time.

Verifications and disclosures

In banking calls, there are strict regulations (verification steps, legal notices, disclosure of fees, etc.). The generative AI agent knows what is required and dynamically adapts for accuracy and compliance, with full visibility for all responses and actions captured in audit logs. For example, it can provide exact wording for legal disclosures (like FDIC insurance details or privacy notices). This protects the bank from compliance breaches.

- **Deployment time:** 1–2 months
- **Value drivers:** Quality assurance
- **Relevant metrics:** Compliance checklist adherence 100%; reduced supervisory interventions or audit findings related to call compliance.

Virtual banking

A generative AI agent can be deployed across voice and chat to handle a wide range of banking inquiries (balance, transfers, card issues, FAQs, basic troubleshooting) at scale. It autonomously resolves routine and complex issues, learning and improving over time. By handling a large portion of inquiries, it turns the contact center into a high-quality, self-service experience. Banks have found that generative AI agents can automate a majority of calls, yielding significant savings.

- **Deployment time:** 2–3 months
- **Value drivers:** Cost reduction, CSAT improvement
- **Relevant metrics:** High percentage of inbound call types automated for significant savings, CSAT maintained or improved despite high automation due to fast service.

Loan pre-approval

Potential borrowers can interact with the generative AI agent to see what loans or credit cards they might pre-qualify for. The AI can ask for some basic information and use the bank's criteria to give a preliminary answer about how much they qualify for, or direct them to apply within the chat. This automated guidance encourages customers to pursue products in a compliant, friendly manner, without needing a human agent to walk them through options.

- **Deployment time:** 1–2 months
- **Value drivers:** Efficiency gain, revenue gain
- **Relevant metrics:** Increase in self-service loan inquiries, reduction in calls to sales agents for basic eligibility questions, improved conversion rate by providing instant answers.

Card activation troubleshooting

When customers run into problems activating a new debit or credit card, a generative AI agent can troubleshoot the issue. It can guide the customer through alternate activation options step by step, determine whether the customer's information needs to be updated, and see if the account is flagged for possible fraud. Once it's identified the issue, the AI agent can work with the customer to get it resolved, in most cases without support from a human agent.

- **Deployment time:** 2–4 weeks
- **Value drivers:** Efficiency gain, cost reduction
- **Relevant metrics:** High percentage of card activation issues resolved through self-service, significant live agent hours saved.

Personalized member experience

A generative AI agent enables credit unions to deliver hyper-personalized member engagement by analyzing member profiles, transaction history, and behavior patterns to recommend tailored financial products, loans, and savings plans. A generative AI agent provides 24/7 personalized support, handling inquiries from routine account questions to complex loan application guidance, improving member satisfaction and boosting engagement without increasing staff burden.

- **Deployment time:** 3-6 months
- **Value drivers:** Efficiency gain, revenue gain
- **Relevant metrics:** Increase in member engagement rates, uplift in cross-sell conversion, reduction in routine inquiry call volume.

Personalized member/customer education

A member or customer logs in to the mobile app or website and engages the generative AI agent to understand financial topics, such as budgeting, credit building, or first-time home buying. The goal is to deliver educational content that's relevant, actionable, and linked to the institution's products. For personal guidance, the AI agent can engage a financial advisor in real-time. This builds financial literacy, increases engagement, positions the institution as a trusted advisor, strengthens loyalty, and reduces routine inquiries to staff by pre-educating members and customers.

- **Deployment time:** 3-4 months
- **Value drivers:** Efficiency gain, CSAT
- **Relevant metrics:** Increase in relevant product enrollment after educational interactions, higher retention among members completing educational journeys, improved CSAT for members interacting with educational offers.

Automate customer service without compromising customer satisfaction

Each of the use cases listed here demonstrates how a generative AI agent can automate customer-facing interactions in financial services 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, banks and other financial institutions 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. 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.