



The Generative AI Agent 100

100 use cases for contact centers by industry

Table of contents

Introduction	03
What is a generative AI agent	04
The shifting legal and regulatory landscape	05
Generative AI agent 100 methodology	05
Use cases	
Property and casualty	07
Financial services	13
Healthcare	20
Travel and hospitality	26
Telecommunications	32
Retail	38
Conclusion	45

This guide presents 100 enterprise use cases where a generative AI agent can drive measurable automation in contact centers. **Each use case is categorized by industry.** We also included estimated deployment times and sample business metrics to help customer experience and operations teams prioritize their generative AI agent automation effectively. Both the estimated deployment time and sample metrics are based on our experience deploying the GenerativeAgent® platform and other AI-powered solutions in enterprise contact centers. Results may vary based on a variety of factors that affect the implementation timeline and deployment value drivers.

Industry verticals covered in this eBook

- ✓ Property and casualty
- ✓ Financial services
- ✓ Healthcare
- ✓ Travel and hospitality
- ✓ Telecommunications
- ✓ Retail

What is a generative AI agent?

A generative AI agent assists—or fully handles—customer interactions by generating human-like responses, surfacing insights, automating tasks, and adapting in real time based on the conversation. It uses large language models (LLMs) to understand context, access systems, and provide accurate, personalized service—improving efficiency, consistency, and customer satisfaction.

ASAPP's GenerativeAgent® platform is an implementation of this concept—built from the ground up for enterprise contact centers. Designed to autonomously manage complex, multi-turn interactions over voice and chat, the GenerativeAgent platform eliminates the need for manual scripting by integrating directly with enterprise APIs and knowledge bases. It dynamically adapts to conversational context, knows when to involve human agents, and supports concurrent interactions with human/AI collaboration. Key capabilities include reasoning, acting on behalf of a customer (e.g., processing refunds or rebooking flights, via API), Human-In-the-Loop Agent (HILA™) support, robust safety guardrails to prevent hallucinations and ensure compliance, and integration with existing IVR/telephony, CCaaS, messaging, and CRM systems.

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 industries, like property and casualty insurance, additional regulations stipulate that only a licensed professional can perform certain services. So, you'll need to be sure any AI agent you're considering will maintain compliance even as the regulatory landscape changes.

Generative AI agent 100 methodology

Each use case includes **deployment time and value drivers**.

Deployment Time

Deployment Time represents typical durations from scoping to live production, derived from ASAPP benchmarks and industry studies.

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

Value Drivers

Value Drivers reflect the potential impact a generative AI agent use case can have on contact center performance in the following areas:

- **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, script adherence, and consistency at scale, and reduces risk.
- **Revenue Retention** – helps reduce churn by improving save rates and proactive service.

Prioritizing high-value use cases ensures
that your organization gets the best
return from automation investments.

Property and casualty

Contact centers in property and casualty (P&C) insurance experience significant pressure due to the **high volume of claims, billing, and policy inquiries**. These core service components constitute over half of all interactions, leading to agent overload with repetitive requests. The average duration of **claims-related calls exceeds seven minutes**, and workforce shortages exacerbate wait times, impacting customer satisfaction.

According to 'McKinsey, 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%**. This enables scalable, compliant handling of high-volume scenarios, allowing human agents to focus on complex claims and dispute resolutions.



1. Policy Information

Description and Value: When customers call with policy questions (coverage limits, deductibles, etc.), a generative AI agent instantly retrieves and summarizes relevant policy details. This reduces hold time compared to agents searching multiple manuals with complex details, enabling faster resolutions. It ensures customers get accurate, consistent answers without long hold times.

- **Deployment Time:** 2–4 weeks
- **Value Drivers:** Efficiency gain
- **Sample Metrics:** Average handle time (AHT) reduced by ~10%; First-call resolution (FCR) improved by 5%.

2. Billing Inquiry

Description and Value: For billing-related calls (e.g. premium due, billing errors), the generative AI agent suggests next steps and responses to the agent. It might pull up billing history or identify common solutions (like explaining a recent rate change). This guidance speeds up call resolution and reduces errors in billing explanations.

- **Deployment Time:** 4–6 weeks
- **Value Drivers:** Efficiency gain
- **Sample Metrics:** Quote process time cut by ~30%; conversion +5%; AHT reduced ~15%; call transfers on billing issues down 20%.

3. New Policy Quote

Description and Value: During cost inquiries (e.g. auto insurance quotes), a generative AI agent gathers initial information from the customer via chat or IVR and suggests a coverage package to the human customer service representative (CSR). The CSR remains in control but spends less time on data entry. This accelerates the quoting process and improves the customer's experience with faster service.

- **Deployment Time:** 4–6 weeks
- **Value Drivers:** Efficiency gain
- **Sample Metrics:** Quote processing time cut by ~30%; sales conversion rate improved (faster quotes) by 5%.

4. Accident Assistance

Description and Value: Immediately after an accident, policyholders might call in distress. A generative AI agent for voice can autonomously walk them through crucial steps (safety checks, gathering information for a claim) and provide reassurance, while a human customer service rep is available to handle any complex deviations (edge cases). This human-in-the-loop mode 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
- **Sample Metrics:** ~40% of simple accident calls handled without full human involvement; customer feedback indicates improved post-accident support experiences.

5. Coverage Guidance

Description and Value: Customers 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 provides personalized answers (or clarifying questions) based on the customer’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
- **Sample Metrics:** Call containment ~25% (coverage questions handled by the generative AI agent); customer satisfaction (CSAT) scores +5%.

6. Claims Filing

Description and Value: When a customer reports a new claim (e.g., 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 (location, time, 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
- **Sample Metrics:** Claims intake time reduced ~20%; error rate in initial claim data down 30%.

7. Document Submission

Description and Value: After an insurance claim is initiated, customers may need to submit documents (photos, police report). A generative AI agent for chat guides the customer 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 customer if something is missing, freeing up agents from follow-up calls.

- **Deployment Time:** 4–6 weeks
- **Value Drivers:** Efficiency gain
- **Sample Metrics:** Document submission cycle time reduced ~50%; claim processing delays due to missing docs down by 30%.

8. Policy Update

Description and Value: For simple policy changes (address change, adding a vehicle or dependent), a voice generative AI agent can authenticate the customer and process the update end to end. This frees customer service reps from mindless 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
- **Sample Metrics:** Containment rate ~60% for address/vehicle change requests (no human needed); operational cost per policy update reduced ~50%.

9. Fraud Prevention

Description and Value: Policyholders often call in to check 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
- **Sample Metrics:** 30–50% decrease in fraud loss (\$ value of fraudulent transactions prevented)

10. Cross-Sell Recommendation

Description and Value: 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 customer's profile and conversational cues to identify potential coverage gaps (e.g., no roadside assistance, no renters 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
- **Sample Metrics:** Increased cross-sell conversion (5–8% of routine service contacts result in an additional product sold); incremental premium growth (\$50–\$150 per customer per year from added coverage); 2–4% higher retention for customers with multiple products.

11. Enhanced Risk Assessment and Underwriting

Description and Value: 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
- **Sample metrics:** 30–50% reduction in data collection time (dynamic questioning captures complete information in one conversation); 20–35% increase in first-pass submission accuracy (all required fields captured, risk factors correctly identified).

12. Compliance Assistance

Description and Value: Insurance calls often require compliance disclosures. A generative AI agent can listen for required phrases and alerts and can prompt the human agent with the correct scripted language if necessary. 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
- **Sample Metrics:** 100% compliance script adherence (from ~95%); reductions in compliance penalties or QA failures.

13. Policies and Guidelines

Description and Value: Generative AI agents handling complex insurance queries (unusual coverage scenarios, policy exceptions) can search the company's knowledge base and underwriting guidelines across multiple articles. The AI quickly provides the relevant excerpt or summarized answer in understandable language, for faster resolution and consistent accuracy. This boosts customer satisfaction, and contact center compliance.

- **Deployment Time:** 4–6 weeks
- **Value Drivers:** Efficiency gain
- **Sample Metrics:** Agent after-call work time down ~25%; first-contact resolution up (agents find correct answers more often on the first try).

14. Emergency Claims Triage

Description and Value: In the wake of natural disasters or other unexpected events where calls spike, a generative AI agent can gather preliminary claim information from callers (location, damage description), provide immediate guidance on next steps (safety, temporary repairs), and help to prioritize high-severity cases for fast-tracking by human adjusters. It ensures no customer has to wait to at least begin the claims process during crises.

- **Deployment Time:** 1–2 months
- **Value Drivers:** CSAT improvement, cost reduction
- **Sample Metrics:** >50% of disaster-related claim intakes handled via AI in peak surge; average response time for initial contact < 1 minute (versus hours), improving customer confidence.

15. Multi-Lingual Support

Description and Value: 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 customer can interact with the generative AI agent, which then either resolves the issue or hands off to a Spanish-speaking human customer service rep with context. This broadens service reach by reducing language barriers.

- **Deployment Time:** 1–2 months
- **Value Drivers:** CSAT improvement, cost reduction
- **Sample Metrics:** Increased number of callers who successfully self-serve in non-English languages (coverage of 5+ languages); CSAT increases among non-English callers +10% due to reduced wait for language-specific agents.

16. Fully Autonomous Claim (End to End)

Description and Value: 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
- **Sample Metrics:** Up to ~89% of simple claims handled without human intervention; claim cycle time for minor claims cut from days to minutes; ~\$4M annual cost savings observed via reduced manual processing in one deployment.

17. Complaints and Escalation

Description and Value: When a customer is upset (e.g. disputing a claim decision), the AI monitors the call sentiment and content. It provides the live CSR real-time coaching, such as suggesting empathy statements or offering a supervisor callback, to de-escalate the situation. It can also summarize the key complaint points for quick escalation to a manager if needed. This results in more consistent, high-quality handling of complaints.

- **Deployment Time:** 1–2 months
- **Value Drivers:** CSAT improvement, cost reduction
- **Sample Metrics:** Complaint call FCR +10% (fewer repeat calls); supervisory escalations reduced by ~20% due to AI-guided de-escalation.

Financial Services

Financial institutions face high contact center volumes driven by sensitive transactions, account issues, and fraud alerts. According to ²McKinsey, most banking-related inquiries are repetitive, rules-based, and prime candidates for automation. Customers expect instant resolutions without compromising data security.

Generative AI agents are uniquely positioned to help banks scale support for account access, payments, and transactional errors. ³DECTA reports that automation in financial services reduces the cost to serve by up to 25%, while maintaining compliance and enhancing fraud detection protocols.

18. Account Balance and Transaction Inquiries

Description and Value: 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 for voice can authenticate the caller (via PIN or voice biometrics), provide the information instantly, and 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
- **Sample Metrics:** Containment rate ~70% for multi-intent conversations that start with balance inquiries (handled entirely by self-service)

19. Password Reset

Description and Value: 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 (using one-time passwords or security questions via the system) 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
- **Sample Metrics:** Up to 60–70% of password reset/unlock requests fully self-serviced; average resolution time < 3 minutes, versus 6–10 minutes with an agent (including hold time).

20. Mortgage Process

Description and Value: Mortgage applicants often have questions about the process (“What does conditional approval mean?” or “How do I lock my rate?”). A generative AI agent integrated into the support center gives instant answers from a trove of mortgage guidelines and educational content, and can explain complex mortgage concepts and financial terms (e.g. “debt-to-income ratio”) in simple terms, ensuring consistency. Customers get clear answers quickly, improving their experience during a stressful process.

- **Deployment Time:** 1–2 months
- **Value Drivers:** CSAT improvement
- **Sample Metrics:** FCR for mortgage-related questions +8% (fewer call-backs needed after clear answers); call handle time on explanatory answers reduced by ~20%.

21. Fraud Alert Verification

Description and Value: 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 information (e.g., “did you attempt a \$500 purchase at X store?”) and either clears the alert or flags it for human follow-up. Human-in-the-loop agents assist with exceptions or complex cases. This speeds up fraud resolution and reduces false declines.

- **Deployment Time:** 1–2 months
- **Value Drivers:** Efficiency gain
- **Sample Metrics:** Fraud verification calls resolved 24/7 with minimal wait; ~40% of alerts cleared automatically without agent, reducing fraud team workload and customer friction.

22. Loan Application Status

Description and Value: 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 (e.g., “your application is in underwriting” or “approved pending document X”). 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 challenging inquiries.

- **Deployment Time:** 4–6 weeks
- **Value Drivers:** Efficiency gain
- **Sample Metrics:** Containment ~50% for loan status queries (no human needed); customer update frequency increased (timely information provided), contributing to +5% NPS (Net Promoter Score) from loan applicants.

23. New Account Onboarding

Description and Value: When a new customer calls to set up or has questions about a new account (checking, brokerage, etc.), the 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 surfaces their application information and suggests it can help them with linking their direct deposit or downloading the mobile app. This context speeds up onboarding and ensures a smooth first experience.

- **Deployment Time:** 4–6 weeks
- **Value Drivers:** CSAT improvement
- **Sample Metrics:** New customer call duration down ~15%; activation of online banking or other features up (more customers complete setup during first call).

24. Product Recommendation

Description and Value: A generative AI agent on a bank’s website/app helps customers choose financial products. For example, a customer might ask “what credit card is best for me?” and the AI will ask a few questions about spending habits, combined with customer data, then recommend a suitable card or loan product. This drives digital sales with a guided, personalized experience.

- **Deployment Time:** 1–2 months
- **Value Drivers:** Revenue gain
- **Sample Metrics:** Conversion rate of online product inquiries to applications +10%.

25. Credit Card Bill and Payment

Description and Value: Customers may inquire about their credit card payment due date, minimum due, or want to make a payment by phone. A generative AI agent for voice can provide due dates, amounts, and even take a payment securely. If a customer requests a payment extension, the agent can ask a human-in-the-loop for approval, then advise the customer if granted and set it up for them. Automating a large volume of calls that are straightforward transactions improves efficiency.

- **Deployment Time:** 4–6 weeks
- **Value Drivers:** Efficiency gain
- **Sample Metrics:** Containment rate ~50% for payment and due date queries; reduction in agent-handled payment calls, saving significant call time per month.

26. KYC Document Collection (Know Your Customer)

Description and Value: 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 the human agent.

- **Deployment Time:** 4–6 weeks
- **Value Drivers:** Efficiency gain, quality assurance
- **Sample Metrics:** Turnaround time for KYC compliance reduced (e.g., majority of customers complete within same call/day); agent time spent on chasing documents reduced by ~40%.

27. Personalized Member/Customer Education

Description and Value: 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 improvement
- **Sample Metrics:** +8–12% increase in relevant product enrollment after educational interactions; 2–4% higher retention among members completing educational journeys; CSAT/NPS +5–8 points for members interacting with educational offers.

28. Wealth Portfolio Inquiries

Description and Value: 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
- **Sample Metrics:** Advisor prep time per call reduced by ~30%; improvement in client satisfaction as advisors can answer detailed questions faster (possibly reflected in CSAT or retention rates).

29. Lost/Stolen Card Reporting

Description and Value: 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 (e.g., about liability for fraudulent charges), speed transaction completion, and reassure customers.

- **Deployment Time:** 4–6 weeks
- **Value Drivers:** Efficiency Gain, cost reduction
- **Sample Metrics:** ~50% of card loss reports fully handled by AI (no agent needed); average time to block card down to <1 minute from call start.

30. Transaction Dispute

Description and Value: 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
- **Sample Metrics:** Time to log a dispute case reduced ~30%; error/omission rate in dispute filings near zero (AI prompts ensure complete information).

31. Real-Time Compliance

Description and Value: 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
- **Sample Metrics:** Compliance checklist adherence 100%; reduced supervisory interventions or audit findings related to call compliance.

32. Comprehensive Virtual Banking

Description and Value: 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 continuously learns from interactions and autonomously resolves routine as well as complex issues. 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
- **Sample Metrics:** ~60% of inbound call types automated within ~60 days of deployment; significant ROI (e.g., tens of millions saved annually in some cases); CSAT maintained or improved despite high automation due to fast service.

33. Loan Pre-Approval

Description and Value: 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 (income, etc.) and use the bank's criteria to give a preliminary answer (e.g., "you may qualify for up to \$X") or direct them to apply within the chat. This automated guidance encourages customers to pursue products without needing a human agent to walk them through options, and can do it in a compliant, friendly manner.

- **Deployment Time:** 1–2 months
- **Value Drivers:** Efficiency gain
- **Sample Metrics:** Increase in self-service loan inquiries; ~20% reduction in calls to sales agents for basic eligibility questions; improved lead conversion rate by providing instant answers (prospects kept engaged).

34. Fraudulent Transaction Detection and Escalation

Description and Value: In banking, generative AI monitors customer interactions in real-time for patterns that may indicate fraud or account compromise. 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
- **Sample Metrics:** 96% fraud detection accuracy, with false positives as low as 0.8%; mean time to respond reduced from hours/days to seconds/minutes; \$35 million saved in fraud loss.

35. Card Activation Troubleshooting

Description and Value: When customers run into problems activating a new debit or credit card, a generative AI agent can troubleshoot and help the customer get the problem resolved. 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
- **Sample Metrics:** >80% of card activation issues resolved through self-service; thousands of agent hours saved per year.

36. Personalized Member Experience

Description and Value: 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
- **Sample Metrics:** 20–30% increase in member engagement rates; 10–15% uplift in cross-sell conversion; 25% reduction in routine inquiry call volume.

Healthcare

Healthcare contact centers are critical to patient access and satisfaction. ⁴The majority of inbound volume relates to appointment scheduling, prescription refills, insurance questions, and billing disputes, according to Healthcare Finance News. Yet many systems rely on outdated IVRs and long wait queues for live agents.

⁵According to a 2024 survey by Healthcare IT News, organizations deploying automation for routine tasks saw a 17% reduction in average call wait time and up to a 22% decrease in abandoned calls. ⁶Similarly, an Accenture study found that automated self-service features can improve first contact resolution while enhancing operational efficiency—key drivers of patient satisfaction and operational throughput.

37. 24/7 Appointment Scheduler

Description and Value: Patients can call or chat at any hour to schedule, reschedule, or cancel appointments using generative AI agents. The AI handles this by accessing open slots in the provider's calendar, offering choices, and booking the appointment. It can also answer basic prep questions ("Do I need to fast before this blood test?"). Eliminating hold times and offering after-hours scheduling greatly improves patient convenience.

- **Deployment Time:** 4–6 weeks
- **Value Drivers:** Efficiency gain, cost reduction
- **Sample Metrics:** Significant reduction in hold times (from an average of 5 minutes to ~0); call containment ~40% for scheduling requests (no human needed); no-show rates improved with instant confirmations.

38. Prescription Refill Line

Description and Value: Patients often call to request medication refills or check refill status. The generative AI agent can verify patient identity, gather prescription details (Rx number or drug name), and submit the refill request to the pharmacy system for processing. It can also inform the caller when the prescription will be ready. This automation streamlines a very common task, freeing pharmacy staff and reducing wait times for patients.

- **Deployment Time:** 4–6 weeks
- **Value Drivers:** Efficiency gain
- **Sample Metrics:** ~45% of refill calls fully handled by the AI; average call duration for refills cut in half; pharmacy staff call volume reduced, allowing more focus on in-person patients.

39. Healthcare Questions

Description and Value: A generative AI agent on a hospital or clinic website can answer common patient questions, like, "What are your clinic hours?" "Do I need a referral to see a specialist?" "What's the procedure for X test?" The AI provides accurate, policy-compliant answers drawn from the health system's knowledge base. This instant self-service reduces the number of calls for basic information and ensures patients get answers even when offices are closed.

- **Deployment Time:** 4–6 weeks
- **Value Drivers:** Cost reduction, CSAT improvement
- **Sample Metrics:** Call deflection for FAQ topics ~50% (patients get answers via chat instead of calling); positive feedback on instant answers (CSAT for website support +10%).

40. Facility Information

Description and Value: Patients can ask a generative AI agent for clinic addresses, hours of operation, or directions, and also whether the clinic is accepting walk-ins or has holiday hours – all updated in real time. By automating these common queries, callers get immediate answers and contact center lines are free for more complex issues.

- **Deployment Time:** 2–4 weeks
- **Value Drivers:** CSAT improvement
- **Sample Metrics:** Containment >80% for facility information queries; overall call volume to live agents reduced by ~5% with no service quality loss.

41. Insurance Coverage

Description and Value: Patients often call insurers asking, “Does my insurance cover X treatment?” The generative AI agent looks up the patient’s insurance details and the relevant coverage rules to provide the patient with the information (or asks additional questions if needed). This avoids a typically lengthy process of putting the patient on hold to manually look up coverage details.

- **Deployment Time:** 4–6 weeks
- **Value Drivers:** Efficiency gain
- **Sample Metrics:** Time to answer insurance coverage queries reduced (e.g., from ~5-6 minutes of hold to 1 minute); fewer call-backs needed after verifying coverage, improving FCR by ~10%.

42. Post-Discharge Instructions

Description and Value: When a discharged patient calls with questions about their care instructions (e.g., “How do I change my bandage?” or “What side effects should I watch for after surgery?”), a generative AI agent accesses specific discharge instructions for that patient’s procedure and relays them to the patient, also answering any questions and involving a human-in-the-loop agent (a nurse or other specialist) as needed. This ensures the patient gets accurate information matching what they were given at discharge, and improving safety and confidence.

- **Deployment Time:** 4–6 weeks
- **Value Drivers:** CSAT improvement, cost reduction
- **Sample Metrics:** Reduction in call time for discharge-related questions (nurse spends less time searching records, perhaps 20% faster); improved patient understanding scores (as measured by follow-up surveys).

43. Billing and Payment

Description and Value: 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
- **Sample 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.

44. Provider Finder Assistant

Description and Value: A generative AI agent can help patients looking for a doctor (e.g., “I need a pediatrician in ZIP code 12345 who speaks Spanish”). The AI accesses the provider directory, filters by specialty, location, language, etc., and presents matching doctors. It can even help initiate scheduling an appointment. This quick “matchmaking” improves patient access and reduces the need for staff to manually search databases while the caller waits.

- **Deployment Time:** 4–6 weeks
- **Value Drivers:** Efficiency gain, cost reduction
- **Sample Metrics:** A high self-service provider success rate (users find a doctor without human help ~70% of time); average time to locate a suitable provider cut drastically (from several minutes to under 1 minute).

45. Multi-Lingual Patient Support

Description and Value: A generative AI agent capable of handling patient inquiries in multiple languages allows a healthcare 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 routine scheduling, directions, and FAQs 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
- **Sample Metrics:** % of non-English calls resolved without a human translator up significantly (e.g., 90%); patient satisfaction in non-English segments on par with English (closing previous CSAT gap of, say, 10% lower to near equal).

46. Vaccine Information

Description and Value: During flu season or public health campaigns, many people call with questions like “Can I get the flu shot at your clinic without an appointment?” or “What are the side effects of a shingles vaccine?” A generative AI agent hotline provides up-to-date answers about vaccine availability, eligibility, scheduling, and general medical information (sourced from approved health information). It handles high volumes of calls that otherwise overwhelm staff, ensuring consistent and accurate information dissemination.

- **Deployment Time:** 4–6 weeks
- **Value Drivers:** Efficiency gain, quality assurance
- **Sample Metrics:** Calls handled per hour massively increased during vaccine season (AI can scale instantly to support volume surges); first-call resolution for vaccine queries ~95% (very few need human transfer); public information accuracy maintained at 100% (scripts updated centrally).

47. Telehealth Tech Support

Description and Value: If a patient is having trouble joining a telehealth appointment (e.g., camera/microphone issues, login problems), a generative AI agent can guide them through troubleshooting steps. It might send the patient a link to test their connection or walk them through installing the telehealth app. The AI can handle common tech issues, while involving a human support specialist for truly technical failures. This ensures patients can access care without long delays due to technical difficulties.

- **Deployment Time:** 1–2 months
- **Value Drivers:** CSAT improvement, cost reduction
- **Sample Metrics:** Successful self-resolution of telehealth access issues for 65% of patients (reducing appointment no-show/late start due to technical issues by, say, 20%); patient feedback on virtual visit ease improved.

48. Member Benefits Questions

Description and Value: When members contact the insurer to understand their benefits, accuracy is paramount, but information is typically complex, confusing and spread across multiple documents within a knowledge base. A generative AI agent can interpret highly detailed and complicated plan information and provide accurate answers to members.

- **Deployment Time:** 2–4 weeks
- **Value Drivers:** Efficiency gain, quality assurance
- **Sample Metrics:** After-call documentation time reduced ~50%; call wrap-up (disposition coding, notes) accuracy improved, as evidenced by fewer omissions in records.

49. Insurance Pre-Auth Status Checker

Description and Value: Patients or providers often call to check if a pre-authorization for a procedure or test has been approved by insurance. The generative AI agent can interface with the insurance portal or internal systems to get the latest status (approved, pending, or more information needed) and present that directly to the authenticated patient. This eliminates long hold times and call-backs that usually happen when staff have to manually check and follow up.

- **Deployment Time:** 4–6 weeks
- **Value Drivers:** Efficiency gain, cost reduction
- **Sample Metrics:** Average resolution time for pre-authorization status queries reduced (from perhaps 2–3 days with callbacks to immediate on call); reduction in follow-up calls by ~30% because patients get real-time answers.

50. Front Desk Triage

Description and Value: Generative AI agent “receptionists” handle the majority of incoming calls to clinics and hospitals for routine tasks—appointment scheduling, directions, lab hour inquiries, prescription refills, and basic triage questions. Human staff are then focused only on complex or urgent calls. The generative AI agent determines the caller’s need(s) and either addresses them or routes the call appropriately. This effectively scales patient support without equivalent staffing increases.

- **Deployment Time:** 2–3 months
- **Value Drivers:** Cost reduction, efficiency gain
- **Sample Metrics:** As much as 45% of calls handled entirely by AI across a healthcare network; call abandonment dropped dramatically (one hospital saw 3× reduction in hang-ups after implementing AI, along with ~30% service improvement); operational savings from lower call center staffing, while maintaining quality.

51. Patient Registration

Description and Value: When a patient calls to register or provide intake information ahead of a visit, the generative AI agent can collect the required information, verify insurance eligibility in real time and ask the patient for any missing information (e.g., updated insurance information). By streamlining registration, the patient spends less time on the phone and check-in is faster on the appointment day.

- **Deployment Time:** 4–6 weeks
- **Value Drivers:** Efficiency gain
- **Sample Metrics:** Registration call time reduced ~20%; fewer errors in patient data entry (e.g., fewer insurance claim rejections due to incorrect info); higher completion of preregistration tasks before visits.

52. Automated Appointment Reminders and Rescheduling

Description and Value: The AI places outbound reminder calls or texts to patients for upcoming appointments and allows them to confirm or reschedule via voice or chat. If a patient says they can't make it, the generative AI agent can offer a new date/time and update the schedule. This reduces no-shows and fills open slots promptly. Patients appreciate the interactive reminder instead of just a one-way message.

- **Deployment Time:** 1–2 months
- **Value Drivers:** CSAT improvement, cost reduction
- **Sample Metrics:** No-show rates reduced (e.g., from 10% to 7% due to easy rescheduling); ~3x improvement in appointment confirmation rates with interactive AI versus static messages; staff time saved by automating what used to be manual reminder calls.

53. Automated Prior Authorization Processing

Description and Value: Generative AI agents streamline prior-authorization 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
- **Sample Metrics:** >98% reduction in per-transaction administrative costs; average provider processing time per prior-authorization falls from hours to minutes.

Travel and Hospitality

Travel brands contend with volatile and high-stress contact center volumes stemming from cancellations, itinerary adjustments, loyalty program questions, and refund processing. ⁷According to Google Cloud data, travel and hospitality companies can automate ~15–30% of incoming contact volume through generative AI agents.

Deploying generative AI agents has resulted in up to a 23% reduction in average handle time and a significant drop in queue abandons for leading travel platforms, according to Google. ⁸A global Bain & Company executive survey found that more than two-thirds of business leaders expect AI to boost efficiency by 20% or more—and early implementations in customer contact centers, including travel and hospitality, are already exceeding these efficiency gains.

54. Voluntary Flight Changes

Description and Value: Customers can contact an airline and book a flight through a voice or chat generative AI agent that can handle the entire transaction: gathering travel details (dates, destination), searching flights, quoting prices, and even securing payment details securely. The AI can upsell options like seat selection or baggage, apply loyalty miles, and confirm the booking by email or text. This provides fast service during peak hours when hold times are long, and frees the agents to handle more difficult calls.

- **Deployment Time:** 1–2 months
- **Value Drivers:** Cost reduction
- **Sample Metrics:** Call containment ~50% for simple booking calls (fully handled by AI); reduction in average booking call length by 30% when AI handles vs. human (no hold or typing delays).

55. Flight Status and Notifications

Description and Value: Travelers often contact airlines or airports for flight status (“Is my flight on time?”). A generative AI agent can instantly provide real-time flight status, gate information, or delay estimates via chat or voice. It can also proactively notify passengers of delays or cancellations and offer rebooking options via a conversational interface. This keeps customers informed and can trigger rebooking workflows without waiting for an agent, which is critical, especially during weather disruptions.

- **Deployment Time:** 2–4 weeks
- **Value Drivers:** CSAT improvement
- **Sample Metrics:** Huge surge capacity – AI handles thousands of status queries per hour during storms, deflecting up to 80% of such calls from agents; customer update speed improved (notifications sent immediately, contributing to higher customer trust scores during irregular operations).

56. Guest Service Request Automation for Hotel Contact Centers

Description and Value: Hotel brands often centralize customer service across properties via contact centers, handling common guest requests like late checkouts, special accommodations, Wi-Fi issues, and billing questions. A generative AI agent can triage and resolve these high-frequency inquiries at scale, accessing reservation data and loyalty profiles to personalize responses and initiate workflows (e.g., scheduling housekeeping or issuing room upgrades). The value of the generative AI agent is in deflecting calls from front desk staff and property teams, improving service speed and consistency for repeat guests, and handling >60% of non-urgent guest inquiries autonomously.

- **Deployment Time:** 1–2 months
- **Value Drivers:** Efficiency gain, cost reduction
- **Sample Metrics:** 60–70% of guest requests handled via automation; reduced average wait time during high-occupancy periods; guest satisfaction scores improved for service responsiveness.

57. Pet Policy

Description and Value: Traveling with pets has become much more common; however, pet policies can be complex and highly varied based on travel mode, country of origin/destination, and the pet itself. A generative AI agent can handle this high-volume interaction type for airlines, hotels and other travel contact centers with speed and accuracy, ensuring travelers are well-prepared and avoid surprises on their journey.

- **Deployment Time:** 1–2 months
- **Value Drivers:** Efficiency gain, cost reduction
- **Sample Metrics:** AI agent is 3.5x faster and 26x less likely to make a mistake vs a human agent; 62% lift in containment (with no repeat interactions).

58. Multi-Lingual Hotel Booking

Description and Value: A hotel chain deploys a generative AI agent for voice that can handle reservation calls in many languages simultaneously. The AI checks room availability, makes bookings, applies loyalty discounts, answers questions (pool hours, parking, etc.), and easily handles high call volume. This dramatically cuts the need for multilingual human staff while improving service accessibility.

- **Deployment Time:** 1–2 months
- **Value Drivers:** Cost reduction, efficiency gain
- **Sample Metrics:** Thousands of calls handled per day by AI across languages; call center operating costs for bookings reduced ~30%; customer wait times for non-English service nearly eliminated.

59. Lost Baggage

Description and Value: A generative AI agent creates reports of lost or delayed baggage after it collects all relevant information (last flight, bag description, delivery address) and provides a tracking ID. The AI can offer immediate small compensations (e.g., toiletry stipend as per policy) or status updates. The AI can involve a human-in-the-loop agent as needed to handle unusual issues or approvals. Customers feel heard immediately, rather than waiting in long claim lines or holds.

- **Deployment Time:** 1–2 months
- **Value Drivers:** Efficiency gain
- **Sample Metrics:** ~60% of baggage reports fully handled by the generative AI agent with no human intervention; average claim report time down to 3–5 minutes (from ~10); improved post-flight CSAT despite baggage issues because the initial response was immediate.

60. Upsell/Cross-Sell Eligibility Monitoring

Description and Value: On calls for high-value travel products (cruise packages, all-inclusive resorts, etc.), a generative AI agent not only answers questions, but provides real-time options if the customer expresses interest in a premium offering (“We’d love an ocean view if possible”). AI provides details and pricing, and can mention add-ons (travel insurance, tours). This ensures there aren’t missed revenue opportunities.

- **Deployment Time:** 1–2 months
- **Value Drivers:** Revenue gain, efficiency gain
- **Sample Metrics:** Incremental revenue per call up by 5–10%; conversion rate on optional add-ons improved.

61. Trip Support and Change Management for Travel Contact Centers

Description and Value: Travel contact centers often handle flight changes, rebooking, missed connections, and special requests. A generative AI agent can dynamically assist customers with itinerary updates, cancellations, or changes by referencing airline, hotel, or tour APIs. It can also notify customers proactively when disruption is detected. Enterprises benefit by automating high-volume itinerary change requests, reducing agent handling time for rebooking scenarios, and enhancing CX with fast, accurate self-service options during travel disruptions.

- **Deployment Time:** 1–2 months
- **Value Drivers:** CSAT improvement, cost reduction
- **Sample Metrics:** Containment ~60% for pre-trip queries; +7% improvement in CSAT for international travelers; reduction in repeat calls from travelers unsure about documentation or policies.

62. Travel Cancellation and Refund

Description and Value: When unforeseen events happen, customers flood contact centers to cancel or get refunds on trips. A generative AI agent can handle standard refund processes (applying cancellation policies, initiating refunds or credits) for eligible cases. For instance, if a fully refundable hotel booking is canceled, the AI processes it end to end. For partial refunds or exceptions, it gathers information and consults a human-in-the-loop agent or offers the best approach per policy. This dramatically speeds up resolution during high-volume cancellation events.

- **Deployment Time:** 1–2 months
- **Value Drivers:** Efficiency gain, cost reduction
- **Sample Metrics:** ~30–40% of cancellation requests fully automated; call volumes during travel disruptions managed with 20% fewer live agents due to AI support; average refund processing time (from request to confirmation) improved by days in many cases.

63. Reservation Support for Travel and Hospitality Contact Centers

Description and Value: Travel and hospitality contact centers frequently field customer calls about restaurant, spa, or excursion bookings. A generative AI agent integrated with booking systems can check availability, make or modify reservations, and suggest upsells (e.g., wine pairings, cabana rentals). This reduces agent workload and increases conversion on ancillary services. The value to the enterprise is in automating a large volume of repetitive reservation requests, driving incremental revenue through smart upsells, and improving CSAT by providing fast, self-serve confirmations.

- **Deployment Time:** 1–2 months
- **Value Drivers:** CSAT improvement, revenue gain
- **Sample Metrics:** ~50% of service reservations handled without agent handoff; increased CSAT due to real-time confirmations; higher conversion on ancillary services through upsell prompts.

64. Customer Sentiment and Recovery

Description and Value: During a travel customer support call (for example, a complaint about a hotel stay or flight delay), the generative AI agent can analyze the caller's tone and words to gauge sentiment. If it detects high frustration or key phrases like "cancel" or "never again," it dynamically adapts to retention or recovery tactics (perhaps offering loyalty points, a voucher, or asking a human-in-the-loop agent for help). This ensures approved, timely save offers are made to appease the customer and prevent churn.

- **Deployment Time:** 1–2 months
- **Value Drivers:** CSAT improvement
- **Sample Metrics:** Save rate of at-risk customers improved (e.g., retention offers made in >90% of high-churn-risk calls identified by AI); overall complaint call satisfaction scores improved by a few points due to proactive recovery actions.

65. Loyalty Program Self-Service

Description and Value: Frequent travelers often have questions about their loyalty points balance, how to redeem rewards, status tier benefits, etc. A generative AI agent can tell a customer their points balance, help them redeem points for a booking or upgrade, and answer policy questions ("How many points to get a free night?"). By integrating with the loyalty database, it even performs transactions like issuing an upgrade. This instant service keeps loyal customers happy and eases the burden on loyalty program call centers.

- **Deployment Time:** 4–6 weeks
- **Value Drivers:** Efficiency gain, cost reduction
- **Sample Metrics:** ~50% of loyalty-related inquiries contained within self-service; increase in reward redemption (as the bot makes it easier); reduced emails and calls to loyalty desk, saving staff hours.

66. Self-Service Reaccommodation / Rebooking

Description and Value: During unplanned events (e.g., weather delays or cancellations), travelers overwhelm contact centers with frantic requests to rebook flights. A generative AI agent can instantly scale to efficiently handle these complex interactions, with no waiting, and high-quality, efficient service. AI can provide accurate options, maintain airline policy and complete the booking process.

- **Deployment Time:** 2–3 months
- **Value Drivers:** Efficiency gains, cost reduction
- **Sample Metrics:** Faster resolution (GenerativeAgent® rebooked passengers for a major US airline in an average of 8 minutes, compared to 29 minutes with a live agent); 20+ point increase in containment, even for previously low-automation tasks (now reaching 70% self-service).

67. Rental Car

Description and Value: A generative AI agent can manage incoming reservation calls or chats. It can quote rates, make new reservations, modify existing bookings (change pickup time or location), and answer common questions about insurance or requirements (driver's license, age). This is done without human intervention for standard cases. The AI's speed and 24/7 availability can improve the customer experience, especially for international travelers dealing with different time zones.

- **Deployment Time:** 4–6 weeks
- **Value Drivers:** Cost reduction, revenue gain
- **Sample Metrics:** Call containment ~55% for rental bookings and changes; labor cost per booking reduced; customers able to secure or change cars after hours, lifting overall rental bookings slightly.

68. Travel Advisory Info

Description and Value: Travelers often ask about destination-specific advisories (weather, safety, etc.). A generative AI agent, kept up to date with official travel advisories and news, can answer these questions accurately. For example, "Is there a hurricane warning in the Bahamas?" The AI provides answers and can proactively push notifications for major advisories. This helps customers make informed decisions quickly.

- **Deployment Time:** 4–6 weeks
- **Value Drivers:** CSAT improvement
- **Sample Metrics:** Improved customer trust; reduction in misinformation-related calls or complaints.

69. Flight Delay Transparent Communication

Description and Value: Airlines use generative AI agents to create detailed, empathetic SMS and app notifications explaining flight delays and disruptions. The AI synthesizes data from operations, crew reports, and external sources to generate transparent and reassuring messages at scale, boosting passenger trust and allowing human agents to focus on complex issues.

- **Deployment Time:** 3–6 months
- **Value Drivers:** CSAT improvement, efficiency gains
- **Sample Metrics:** 6% improvement in customer satisfaction; over 100,000 automated push notifications sent to passengers explaining flight changes; reduction in manual message-drafting workload, enabling quicker and more accurate passenger communications during peak travel.

Telecommunications

Telecom providers manage some of the highest customer service volumes in the enterprise world—often driven by outages, billing cycles, plan changes, and device issues. The industry consensus is that the majority, ~65% according to ⁹Accenture, of telecom inquiries are repetitive and ripe for automation. Generative AI agents can address common customer pain points instantly and at scale.

Industry data shows that telecom companies deploying automation reduced average handle times by ¹⁰20–30% and improved ¹¹NPS by 15 points, all while lowering support costs by over 25%. In fact, at Vodafone, low-code & AI-powered agents delivered **~25% increases in service speed and agent productivity**, while reducing AHT and elevating NPS.

70. Outage Troubleshooting

Description and Value: When customers suspect an Internet or phone outage, they typically flood the call center. A generative AI agent for voice can handle these calls by checking the customer's line status and known outage information by address. If there's a known outage, it informs the customer and provides an ETA for the fix. If not, it guides the customer through basic troubleshooting (modem reset, etc.). The generative AI agent rapidly addresses customer concerns and decreases frustration during outages, all without waiting for a human.

- **Deployment Time:** 4–6 weeks
- **Value Drivers:** CSAT improvement, cost reduction
- **Sample Metrics:** During outages, up to 70% of calls deflected from live agents with automated status updates; average call time for outage inquiries < 2 minutes; improved CSAT in outage scenarios (customers appreciate proactive info).

71. New Service Signup

Description and Value: A prospective customer calls to set up a new service (Internet, TV, mobile). The generative AI agent checks service availability at the customer's address, recommends optimal plans based on the customer's needs (speed, channels, etc.), and calculates any bundle discounts. The AI also checks to see if the new customer is eligible for any promotions, and if so, applies them, then confirms monthly payment and billing cycle and shipping address for any new device. This leads to faster, more accurate orders.

- **Deployment Time:** 1–2 months
- **Value Drivers:** Efficiency gain, revenue gain
- **Sample Metrics:** Order processing time reduced ~20%; fewer errors in service orders (AI validation); sales conversion rate improved a few percentage points due to immediate plan recommendations.

72. New Customer Onboarding

Description and Value: After a customer signs up for a new service, a generative AI agent reaches out (via call or chat) to welcome them, verify they received equipment, and help with setup. It can answer common first-week questions (Wi-Fi name/password changes, initial bill expectations) and ensure the service is working properly. While mostly autonomous, it can loop in a human agent if it detects unresolved issues (no signal, etc.). This proactive onboarding reduces early frustration and returns/cancellations.

- **Deployment Time:** 1–2 months
- **Value Drivers:** Efficiency gain, CSAT improvement
- **Sample Metrics:** New customer call-ins about setup issues are reduced by ~25% (because generative AI agent addresses them first); improved 90-day retention by several percentage points (smoother onboarding leads to fewer early disconnects); higher initial satisfaction ratings as welcome calls appreciated.

73. Plan Upgrade Suggestion

Description and Value: If a customer calls frequently about overage fees or additional services, the generative AI agent recognizes this pattern and suggests a more suitable plan. For example: "I see you've incurred overage fees; it might be more cost-effective to upgrade to our premium plan, and avoid the additional charges. This may generate more revenue, and more importantly, helps customers move to plans that better fit their usage, which may improve CSAT long-term.

- **Deployment Time:** 4–6 weeks
- **Value Drivers:** Revenue gain
- **Sample Metrics:** Conversion rate on plan upgrades +10% with AI suggestions; reduction in future overage complaint calls from those customers (improved experience); incremental ARPU (average revenue per user) lift identified.

74. Device Troubleshooting Virtual Tech

Description and Value: Customers often call with device issues. A generative AI agent can walk customers through troubleshooting steps conversationally. It asks questions about indicator lights, instructs the user to reboot or check cables in natural language, and dynamically adapts the dialogue based on responses. Many issues (like Wi-Fi not working) can be resolved without a technician or human agent. This lowers support costs and resolves problems faster for the customer.

- **Deployment Time:** 1–2 months
- **Value Drivers:** Cost reduction
- **Sample Metrics:** First-call resolution improvement (many tech issues solved by AI—FCR +15%); reduced truck rolls (field tech dispatches) as minor issues fixed on call; support cost per subscriber down by 30% through automation.

75. Data Usage and Upgrades

Description and Value: A mobile carrier's generative AI agent can handle questions like, "How much data have I used this month?" or "Why is my data slow now?" by accessing the customer's usage and plan details. It can explain any throttling policies if they exceed limits and, if appropriate, offer an upgrade to a higher data plan on the spot via a quick consent. This real-time personalized service addresses customer concerns about data overages or slowdowns and presents a solution proactively.

- **Deployment Time:** 4–6 weeks
- **Value Drivers:** Cost reduction, revenue gain
- **Sample Metrics:** Containment ~50% (many usage queries handled fully by generative AI agent); upsell conversion of heavy users to bigger plans increased (some percentage of users accept upgrade offers via AI); fewer complaints about "unknown" data slowdowns as AI transparently communicates policy triggers.

76. Billing Explainer and Invoice

Description and Value: Telecom bills can be confusing. A generative AI agent can take a customer's query ("Why is my bill higher this month?") and break down the charges in simple language: e.g., "Your data usage exceeded your plan by 2GB, resulting in \$15 extra charges, and a \$5 equipment fee was applied." It can email or text a summary. By handling these explanations and even emailing a simplified bill, the generative AI agent reduces call escalations and would likely improve CSAT.

- **Deployment Time:** 4–6 weeks
- **Value Drivers:** CSAT improvement, cost reduction
- **Sample Metrics:** Repeat calls about the same billing issue are reduced (better understanding); faster resolution for billing questions (generative AI agent instantly fetches usage details); CSAT regarding billing transparency is also improved. .

77. Proactive Outage Alert System

Description and Value: Instead of waiting for customers to call in about outages, the generative AI agent automatically calls or texts customers in affected areas to inform them of the issue and an estimated time to fix. It allows customers to say if they're also experiencing issues (feeding back into detection systems) or ask questions, which the AI can answer ("Yes, this affects TV and Internet"). By proactively communicating, the telco reduces incoming call volume and keeps customers informed, which improves trust.

- **Deployment Time:** 4–6 weeks
- **Value Drivers:** CSAT improvement, cost reduction
- **Sample Metrics:** During outage events, incoming calls are reduced by up to 30% due to proactive notifications; customer feedback indicates higher satisfaction with communication (CSAT scores during outages improve vs. past events without AI outreach).

78. Payment Arrangement

Description and Value: When customers cannot pay their full bill, they often need to call to arrange a payment plan or extension. A generative AI agent can handle this discreetly: verifying account identity, offering eligible payment plan options (e.g., splitting into installments or extending due date by X days per company policy), and setting it up automatically. This reduces bad debt by making it easier for customers to commit to a plan, and it frees agents from lengthy, sensitive calls.

- **Deployment Time:** 4–6 weeks
- **Value Drivers:** Cost reduction
- **Sample Metrics:** 85% of payment extension requests self-served, reducing agent workload; decrease in involuntary churn or late disconnects, as more customers engage with payment plans; improved call handle time for those that do go to agents because the generative AI agent knows the pre-qualification options.

79. Churn Risk Retention Agent

Description and Value: When a customer calls to cancel their service, a generative AI agent analyzes their account (tenure, usage, previous complaints) and suggests a tailored retention offer. For example, “Offer \$10 off for 12 months + upgrade to higher speed tier, since the customer had recent issues and has been with us 5 years.” A generative AI agent ensures consistency in retention efforts and maximizes the likelihood of keeping valuable customers.

- **Deployment Time:** 1–2 months
- **Value Drivers:** Revenue retention
- **Sample Metrics:** Save rate on cancellation calls increased (e.g., from 30% to 45%); reduction in churn among high-value segments; retention offer cost efficiency improved (offers tailored so as not to over-discount for customers likely to stay anyway).

80. Field Technician Dispatch

Description and Value: If an issue can’t be solved remotely, scheduling a technician is next. The generative AI agent assists by automatically checking the customer’s address, matching it with technician availability and skill, and confirming an appointment slot. It also lists likely tools or equipment needed based on the troubleshooting information gathered for inclusion on the service order. This reduces the back-and-forth scheduling effort and increases first-visit resolution rates (tech arrives informed and equipped).

- **Deployment Time:** 4–6 weeks
- **Value Drivers:** Efficiency gain, reduced cost
- **Sample Metrics:** Scheduling handled ~30% faster; first-time fix rate for dispatched jobs improved (tech arrives with context); customer downtime shortened by scheduling at call time instead of later follow-up.

81. Voice Biometric Authentication

Description and Value: The generative AI agent uses the customer’s unique voiceprint to authenticate their identity within seconds of speaking, instead of asking security questions or PINs. On a support call, as the customer describes their issue, the generative AI agent matches their voice to the stored print on file. Once verified, the AI proceeds with account-specific actions. This reduces handle time and improves account security (harder to fake a voice).

- **Deployment Time:** 1–2 months
- **Value Drivers:** Efficiency gain, quality assurance
- **Sample Metrics:** Average verification time reduced from ~30-60 seconds of Q&A to ~5 seconds; call drop-offs during tedious ID verification reduced; fraud attempts caught (voice not matching) that might have bypassed Q&A.

82. Visual-Assisted Troubleshooting & Guided Repair

Description and Value: During a voice or text conversation, the customer can upload photos or stream video. The generative AI agent uses visual AI to detect issues and provides step-by-step voice or text guidance. For example: The AI tech support agent asks the customer to point their camera at a router. The AI identifies loose cables and provides visual overlays in an app showing where to plug them in.

- **Deployment Time:** 3-4 months
- **Value Drivers:** Cost reduction
- **Sample Metrics:** 15-25% truck roll avoidance for field services due to ability to resolve issues remotely.

83. Coverage Map and Service Availability

Description and Value: Customers commonly ask, "Do you have 5G coverage in my area?" or "Can I get fiber at my house?" A generative AI agent can answer these questions by checking coverage maps and service eligibility based on address or zip code. It can also inform the customer which technologies or speeds are available and even suggest upsell or promotional offers to eligible customers.

- **Deployment Time:** 2-4 weeks
- **Value Drivers:** Cost reduction, revenue gain
- **Sample Metrics:** High self-service success for coverage queries (containment ~85%); increased lead capture for areas with upcoming service (AI collects information and makes an offer instead of missing an opportunity).

Retail

Retailers face volatile seasonal demand, high return rates, and complex loyalty programs. ¹²According to Gartner, while self-service channels are used in 70% of customer journeys, only 9–14% of issues are fully resolved via those channels—and yet up to 40% of live agent contact volume could be deflected through better structured self-service experiences.

Generative AI agents offer scalable, automated handling of returns, order tracking, and product inquiries. ¹³McKinsey's analysis of digital transformations across retail and adjacent sectors suggests **customer service costs can fall by 20-40%**, and customer satisfaction metrics can improve significantly after deploying automation across key interaction points.

84. Order Delivery Issues

Description and Value: 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, on-demand service 24/7 boosts customer confidence in the retailer.

- **Deployment Time:** 4–6 weeks
- **Value Drivers:** CSAT improvement, cost reduction
- **Sample Metrics:** Containment >70% for complex WISMO queries; reduction in live agent volume; CSAT for delivery information interactions +5% (due to immediacy).

85. Returns and Refunds

Description and Value: 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 an RMA (return merchandise authorization) or return shipping label via email, and process the refund 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 streamlines the reverse logistics pipeline and provides immediate resolution to customers on return requests.

- **Deployment Time:** 1–2 months
- **Value Drivers:** Cost reduction
- **Sample Metrics:** A large portion of return requests (e.g., 50%) handled end to end by AI; average time to issue refund reduced (customers get confirmations faster); support cost for returns down significantly (fewer touches by staff).

86. Product Information

Description and Value: Shoppers often have questions about products (specifications, compatibility, usage) before purchase. A generative AI agent for voice or chat can answer these. For example, “Is this laptop compatible with XYZ software?” or “Does this dress have pockets?” The AI accesses product descriptions, manuals, and prior Q&A. Instant answers reduce purchase anxiety and prevent calls or emails to customer service. It can also clarify policies (warranties, returns) in context.

- **Deployment Time:** 4–6 weeks
- **Value Drivers:** CSAT improvement, revenue gain
- **Sample Metrics:** Reduction in pre-sale support contacts (customers get answers on site, leading to perhaps 10% fewer calls/emails); higher conversion rate on product pages with active Q&A (customers more likely to buy when questions are resolved); improved CSAT as customers feel informed.

87. Personalized Product Recommender

Description and Value: A conversational generative AI agent for shopping assistance helps customers find products. For instance, a customer can chat, "I need a gift for my 10-year-old nephew who loves science." The AI will ask a couple of questions and then recommend a few items (with reasoning: "This chemistry set is popular for that age."). By mimicking an in-store sales associate, it increases engagement and basket size. It can upsell complementary products ("You might also need batteries for that toy."). This enhances the online shopping experience and drives sales.

- **Deployment Time:** 1–2 months
- **Value Drivers:** Revenue gain
- **Sample Metrics:** Increase in average order value (contributing an uplift in overall conversion rate); positive feedback on the ease of finding suitable products.

88. Store information and Inventory

Description and Value: Customers often call stores asking for hours, locations, or if a specific item is in stock at a nearby store. A generative AI agent can address these automatically: provide today's store hours, directions via text link, and check inventory systems to say "Yes, the 55 TV is available at our Elm Street location." This saves store associates from constantly fielding calls and gives customers rapid answers.

- **Deployment Time:** 2–4 weeks
- **Value Drivers:** Efficiency gain, cost reduction
- **Sample Metrics:** ~80% of store inquiry calls handled by the AI; associates reclaim significant time to assist in-store customers; customer inquiries answered in seconds, improving experience (reflected in fewer complaints about store phone lines).

89. Loyalty Points Balance & Redemption

Description and Value: A loyalty program generative AI agent allows customers to query their loyalty/rewards points and even redeem them. For example, via chat: "How many points do I have?" followed by "Redeem \$10 coupon with my points." The AI confirms and applies the reward to their account or provides a code. Instant gratification and clarity around loyalty rewards encourage program engagement. It also answers related questions (tier status, how to earn points). This reduces calls and keeps loyal customers happy.

- **Deployment Time:** 4–6 weeks
- **Value Drivers:** CSAT improvement
- **Sample Metrics:** Self-service redemption up (more customers redeem points, engagement +40%); loyalty-related calls down significantly; CSAT among loyalty members up, noting ease of use.

90. Order Modification

Description and Value: If a customer needs to change an online order (shipping address, quantity, or cancel an item shortly after ordering), the 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
- **Sample Metrics:** AHT for order change and cancellation calls reduced ~20%; percentage of modifications done automatically (prior to shipping) high, leading to fewer manual interventions; improved goodwill – customers appreciate timely intervention (fewer "I couldn't change it in time" complaints).

91. Conversational Shopping Assistants

Description and Value: 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 example: "I'm looking for a gift under \$50 for a 10-year-old who loves science." The AI suggests products, explains features, and can even offer bundled deals.

- **Deployment Time:** 2+ months
- **Value Drivers:** Revenue gain, cost reduction
- **Sample Metrics:** 10–20% decrease in abandoned carts; Average Order Value (AOV): Impact of AI-driven upselling/cross-selling (can increase 5–20%); Containment rate: % of sessions resolved without human escalation (70–90%)

92. Warranty Claim

Description and Value: 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 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
- **Sample Metrics:** ~40% of warranty claims initiated via self-service; reduction in claim processing time by days (immediate initiation vs. waiting for email replies); lower support cost per claim processed.

93. Payment Issue

Description and Value: If a customer calls about a payment problem (e.g., they were double-charged, or their coupon code didn't apply), the AI can quickly access the transaction, verify the error (comparing system records), and help the customer understand their options (process a refund of the extra charge, or manually apply a discount and refund difference). The AI can handle the transaction, for a quick and satisfactory resolution.

- **Deployment Time:** 4–6 weeks
- **Value Drivers:** Efficiency gain
- **Sample Metrics:** Resolution time for billing issues reduced (often can be done within one call vs. investigations); decrease in escalations to billing department by 40%; improved customer perception of fairness/competence as errors are corrected on the spot.

94. Gift Card Inquiries

Description and Value: 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 (expiration, where it can be used). Automating these inquiries saves the retailer considerable call volume during holiday seasons when gift card use is high.

- **Deployment Time:** 2–4 weeks
- **Value Drivers:** Efficiency gain
- **Sample Metrics:** Containment ~90% for gift card inquiries; peak-season support load reduced.

95. Price Match Inquiry

Description and Value: Some retailers offer a price match guarantee. A customer might contact support saying "I found this item cheaper at Competitor X." The AI agent can handle the negotiation: it will ask for the competitor price and a link, verify it (through a quick web lookup or database of competitors), and if it meets the criteria, automatically apply a price adjustment or provide a promo code for the difference. If not eligible, it politely explains 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
- **Sample Metrics:** Time to resolve price match requests cut significantly (from days of review to real-time in call/chat); increased successful price matches processed (customers more likely to request since it's easy, but also more likely to buy because they got their match); reduced need for managerial approvals on straightforward matches.

96. In-Store Appointment Scheduler

Description and Value: For retailers that offer in-store services (personal shopping consultations, electronics repair, makeup tutorials, etc.), customers can book appointments through a generative AI agent that checks the store's appointment calendar, offers time slots, and books the customer appointment. It can also reschedule or cancel existing appointments upon request. This spares store staff from handling scheduling phone calls and provides customers the convenience of self-serve booking at any time.

- **Deployment Time:** 4–6 weeks
- **Value Drivers:** Efficiency gain
- **Sample Metrics:** Majority of appointments (e.g., 60-70%) booked online by the generative AI agent rather than phone, reducing calls to humans; appointment attendance and no-show rates improved with immediate confirmations and reminders from the system; store staff labor saved – fewer interruptions for scheduling tasks.

97. Post-Purchase Support

Description and Value: 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
- **Sample Metrics:** 60-80% containment; 5-15% reduction in returns attributed to setup/usage issues; 30–50% faster time to resolve warranty claims

98. Delivery Rescheduling

Description and Value: For scheduled deliveries (furniture, appliances) or signature-required packages, customers sometimes need to change the delivery date or time window. A generative AI agent can handle this 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
- **Sample Metrics:** Successful self-initiated reschedules (reducing failed deliveries by 30%); fewer calls back and forth with delivery coordinators; customer satisfaction with the retailer's delivery process is improved.

99. Seasonal Volume Containment for Contact Centers

Description and Value: 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
- **Sample Metrics:** ~70% containment of peak-period interactions; 40% lower seasonal staffing requirements; maintained or improved CSAT during high-demand windows.

100. Product Recalls and Incident Response Automation

Description and Value: 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
- **Sample Metrics:** ~70% of recall-related inquiries resolved without agent handoff; 40% reduction in response time during incident surges; improved compliance through standardized responses and documentation handling.


Conclusion

Each of these use cases demonstrates how a generative AI agent can automate customer-facing interactions in 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, businesses can modernize their customer service while tracking metrics to ensure each deployment delivers real value.

Sources

1. McKinsey & Company. (2023). Automation in claims processing: Reducing turnaround time. ResearchGate.
2. McKinsey, "Extracting value from AI in banking: Rewiring the enterprise," 2024.
3. DECTA, "Banking Automation in Digital Banking," 2025: "Automation reduces banking costs by up to 20–30%. Modern fraud detection systems significantly improve banking security."
4. Healthcare Finance News, "Improving Patient Access Through the Contact Center," (2022).
5. Healthcare IT News, "Healthcare Contact Center Automation Survey Results," 2024.
6. Accenture, "Digital Health Tech Vision," 2023. "Automated self-service tools can improve first contact resolution and enhance operational efficiency in healthcare call centers."
7. Google Cloud Blog, "Customer Engagement Suite: Stronger results, and new AI features," 2025.
8. Consultancy.me, "AI and Gen AI to boost worker productivity and disrupt industries," citing Bain & Company, 2023.
9. Accenture. (2025, April 24). AI Transformation for U.S. Telcos: Insights. Accenture.
10. Forrester, "Contact Centers Win At Leveraging AI For Operations," June 2025.
11. TMForum, "Digital Transformation and Intelligent Automation Benchmark Reports," 2024/25.
12. Gartner, Inc. (2024, August 19). Survey finds only 14% of customer service issues are fully resolved in self-service. Gartner Newsroom.
13. McKinsey & Company. (2025). Winning the cost battle: Success factors in digital transformations for energy retailers—noting 30–50% cost-to-serve reductions and up to 70% automation potential.

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.