ASAPP

The contact center buyer's guide to generative AI agents

The promise of a generative AI agent in your contact center is clear – increased capacity, a better customer experience, and lower cost to serve. But the actual value you realize depends heavily on performance, safety, and flexibility. That makes choosing the right AI agent for your business critical.

We want to be candid here. We want ASAPP to be the answer. We would love for you to choose GenerativeAgent® for your customer service – but only if that's the right choice for your business. With that in mind, here's our best advice on what to consider as you decide.

What does a robust solution include?

A safe and reliable AI agent is far more than a large language model (LLM) grounded in your internal knowledge sources. It's a sophisticated system that integrates multiple LLMs, layered input and output safety mechanisms, automated QA for responses, secure connections to external systems, transparent documentation of the agent's behavior, and no-code tools that allow business teams to manage and optimize performance with agility.

This complexity is essential to achieve the safety, accuracy, and performance that enterprise operations demand. It's also what separates a functional prototype from a mature, production-ready solution that delivers consistent results at scale.

A safe and reliable AI agent should include:

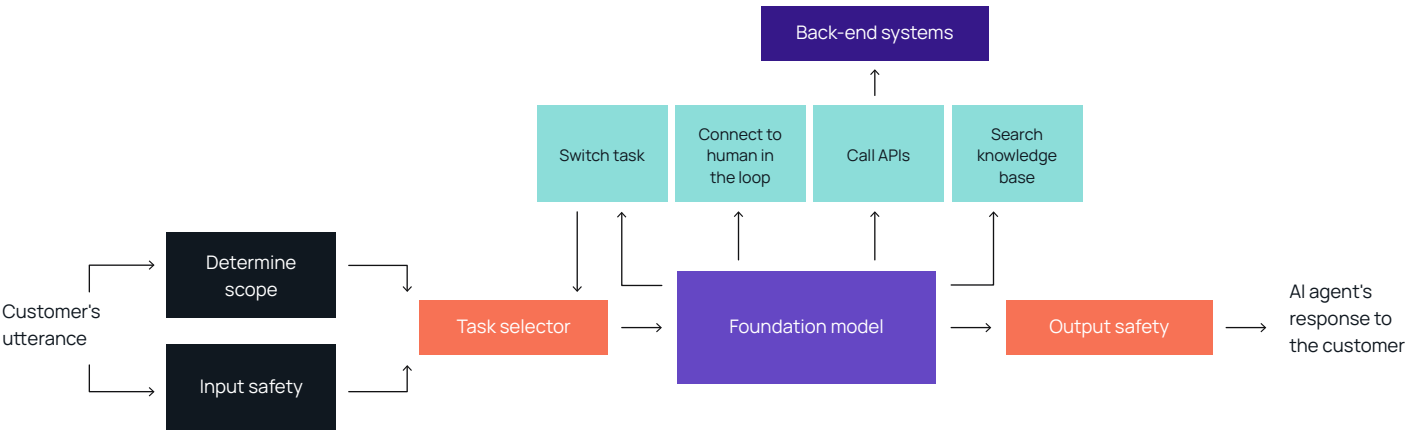

ASAPP

Multiple LLMs

No one model will be best suited for every task, so it's important to incorporate multiple models into your AI agent solution, using each one for what it does best. One model might be tasked with primary reasoning, while another serves as a judge to assess the accuracy and relevance of responses before they are sent out to your customer. This real-time evaluation is critical because LLMs tend to favor their own content, increasing the likelihood of inaccurate or irrelevant responses.

AI models are evolving quickly, so flexibility is essential. The best model for a task today may not be the best model for that task in the future.

How ASAPP delivers

GenerativeAgent incorporates multiple LLMs, each dedicated to a different task. The architecture of GenerativeAgent is designed so that we can swap out the backend LLM as needed whenever a more advanced or suitable model becomes available. Because ASAPP is focused on AI research and development, we stay on top of model differences, and we manage all testing and QA for seamless maintenance, even when a model is swapped. That reduces your development costs, minimizes your risk, and produces consistently high performance.

Secure integration with external systems

As your AI agent works to resolve customer inquiries effectively, it will need access to other systems like CRMs and billing platforms, to retrieve data and take action on behalf of the customer. The AI agent will rely on APIs to access those systems. In many cases relevant APIs likely already exist in your technology ecosystem; however, they may not deliver data in a format your AI agent can process directly. Your technical teams will either need to modify or wrap those APIs, or develop another method for translating their output so your AI agent can use it.

How ASAPP delivers

ASAPP has developed an API transformation layer that simplifies the retrieval and passing of information to GenerativeAgent. This ensures high performance while at the same time filtering out data that should not be exposed to an LLM. **As a result, there's no need to rewrite your existing APIs.** This reduces costs for both initial deployment and later expansion to additional use cases.

ASAPP

Model orchestration

LLMs are powerful but there are still serious limits to what they can accomplish on their own. They sometimes struggle with multi-step processes and applying contextual information in a meaningful way. Using multiple models, each for a specific purpose, can help overcome these limitations – but only with nuanced orchestration. Model orchestration manages the complex workflow across components and processes, including the interaction between the various LLMs, API calls, data retrievals, prompt management, and more. **This orchestration layer is actually the backbone of any enterprise-ready AI agent solution.**

How ASAPP delivers

GenerativeAgent is built on a sophisticated approach to model orchestration. In addition to a foundational LLM, the system uses more than ten models that each perform discrete tasks. The orchestration layer manages their complex, simultaneous interactions to deliver superior accuracy and safety.

Data security

Because your AI agent will access other systems using APIs, security and authentication in the API layer are critical to ensure that the AI agent can access data only for the customer it's actively interacting with and only for the task it's currently performing, and that it cannot retrieve data it is not authorized to use.

During customer interactions, personal identifiable information (PII) is sometimes part of the conversation and is often necessary to resolve the customer's issue. To maintain data privacy, this PII should always be redacted before the data is stored. If you use third-party components in your solution, you'll need to be able to guarantee that if any data is shared with a third party, it's redacted before leaving your infrastructure.

How ASAPP delivers

Authentication in ASAPP's API layer only permits the retrieval of data for the authenticated user. This approach ensures that GenerativeAgent cannot make API calls for other users or systems, which enhances security and data privacy. ASAPP redacts all data before storage, including when working with third-party providers (e.g. LLM model providers).

ASAPP

Input safety

Effective AI agents need multi-layered defenses against malicious or manipulative inputs. These defenses must detect and block exploits such as prompt injection, abuse attempts, or out-of-scope requests that could harm your brand. A layered approach—covering prompt filtering, harmful language detection, and safeguards to keep the AI within defined boundaries—is crucial.

How ASAPP delivers

GenerativeAgent is a multi-layered solution that employs automated input safety evaluators to filter malicious inputs and guard against a wide range of jailbreaking techniques. These mechanisms include prompt filtering, content filtering, models to detect harmful language, and mechanisms to keep the AI within scope.

Output safety

Even with strong grounding, LLMs can hallucinate, producing inaccurate or misleading responses. AI agents must include safety checks that review outputs before sharing them with customers, along with clear escalation paths for when AI-generated responses fall short. Whether the AI retries or transfers to a human agent, these processes are essential to maintain trust and service quality.

How ASAPP delivers

GenerativeAgent addresses and significantly reduces the risk of AI hallucinations through several measures, including the use of guardrails and safety agents that analyze traffic and outputs to identify and prevent the generation of inaccurate or misleading information. In addition, both automated quality assurance testing and manual red teaming focus on identifying hallucinations. A human-in-the-loop escalation process provides the opportunity for your customer service team to intervene and rectify any AI-generated inaccuracies. We also ensure that responses are firmly grounded in your data, policies, and knowledge base.

For actions that alter customer data, GenerativeAgent supports a confirmation workflow outside of the LLM that requires explicit user confirmation before the action is executed.

ASAPP

Human-in-the-loop

Few enterprises are ready to turn over too much of their customer service to AI without clearly defined processes for human intervention and oversight. In other words, there's broad agreement that it's important to keep a human in the loop. **The question is, what do those processes look like, and what exactly do these humans in the loop do?** You have a range of options to consider:

- Treat the humans as an escalation point to take over an interaction when the AI agent fails, much like they do with traditional deterministic bots
- Require human oversight and approval for every action the AI agent takes, which severely limits the impact the AI agent can make on your contact center's capacity
- Define a middle ground between complete AI independence and total human oversight by creating a solution that can work collaboratively with human agents.

Whatever path you choose, your solution will need to enable the human-AI relationship you're envisioning.

See why Forrester says the ASAPP human-in-the-loop capability enables the transition to AI-led customer service by capturing the tacit knowledge of human agents.

How ASAPP delivers

GenerativeAgent supports multiple approaches for integrating humans into the workflow. It can often avoid a hard escalation by asking a human-in-the-loop agent for the help it needs to continue resolving a customer's issue on its own. It's smart enough to know when it needs help, and it knows how to ask for what it needs. Through real-time collaboration with a human in the loop, GenerativeAgent can continue serving the customer without handing off the interaction. Because it was designed and built with this collaboration in mind, GenerativeAgent enables the human-in-the-loop agent with user-friendly tools in an intuitive interface for an efficient and streamlined workflow.

You can also configure GenerativeAgent to ask humans for help at certain critical points, or even more frequently.

This advanced capability creates new possibilities for how you can incorporate GenerativeAgent into your customer service workflows. With this innovative approach to the human in the loop, you can expand automation opportunities safely, improve the performance of GenerativeAgent through self-learning and continuous optimization, and increase your contact center capacity without shortchanging the customer experience.

ASAPP

Intuitive controls for business users

Customer service operations must adapt constantly. AI agents need no-code tools that empower non-technical business users to update processes, refine tone, and launch new use cases—without waiting on development resources. This agility prevents bottlenecks and ensures the AI stays aligned with evolving business needs.

How ASAPP delivers

GenerativeAgent was built for business users to operationalize quickly with fewer labor hours and at lower operational cost. No-code tooling enables non-technical business users to handle regular optimization and adjustments to policies on their own. Over time, as you expand your use of GenerativeAgent, you can add intents without heavy reliance on IT.

Production monitoring

Generative AI behaves probabilistically, which means its outputs can shift unexpectedly over time. Continuous, real-time monitoring is essential to catch issues early, understand performance trends, and intervene before small problems escalate.

How ASAPP delivers

ASAPP provides robust production monitoring of 100% of GenerativeAgent responses, and has an ongoing quality improvement process (including human annotation) to identify and reduce problems that happen at scale. And, if something does go wrong, ASAPP can alert you early so that you can get ahead of any potential issues while they are still small.

Production testing

Generative AI generates new information and content based on patterns, often producing varied responses that can be difficult to predict. **This variability makes testing an AI agent exceptionally difficult, but it is a crucial component in any enterprise-grade solution.** Your AI agent should include a robust pre-production testing process that combines multiple techniques, including simulation-based testing from production data.

How ASAPP delivers

GenerativeAgent comes with easy-to-use testing and simulation tools to help you ensure that it performs as you expect, even at scale. These tools include mock APIs, mock data, and simulation testing interfaces.

Keeping pace with AI innovation

As quickly as the technology is evolving, keeping up with every model improvement and new technique for applying AI in the real world can seem overwhelming.

How ASAPP delivers

ASAPP is an AI company with dozens of leading experts in AI, data science, and solution development. Our commitment to AI research is evident in the 60+ patents we've been granted so far, with more pending. We've published dozens of research papers on AI, and we are continuing to push the boundaries of AI solution development. That means GenerativeAgent is backed by a team that is totally focused on AI innovation.

Conclusion

The nature of generative AI, the speed with which LLMs are evolving, and the lingering limitations of the technology make this effort quite different from other software and infrastructure purchases. We'd love the chance to show you how GenerativeAgent can fit into your customer service operation and deliver on your long-term business goals.

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. 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.