ASAPP

Generative AI agents: Key questions to ask every solution provider

Every vendor who sells a generative AI agent for contact centers makes the same big claims about what you can achieve with their product – smarter automation, increased productivity, and satisfied customers. That language makes all the solutions sound pretty much the same, which makes a fair comparison more difficult than it ought to be.

If you want to get past the vague language, take control of the conversation by asking these key questions. The answers will help you spot the differences between solutions and vendors so you can make the right choice for your business.

What *exactly* does your AI agent do?

Some AI agents simply automate specific processes or serve up information and other guidance to human agents, while others can operate independently to talk to customers, assess their needs and take action to resolve their issues. Ask these questions to distinguish between them.

- Can your genAI agent handle customer interactions from start to finish on its own? Or does it simply automate certain processes?
- *How* do your agents use generative AI?
- What channels does your AI agent support?

Look for a solution that uses the full range of generative AI's capabilities to power an AI agent that can work independently to fully automate some interactions across multiple channels, including voice. This type of agent can listen to the customer, understand their intent, and take action to resolve the issue.

Is there more to your solution than a LLM + RAG?

Retrieval augmented generation (RAG) grounds generative AI agents on an authoritative source, such as your knowledge base. That helps the solution produce more accurate and relevant responses. It's a dramatic improvement that's invited some to ask whether RAG and a foundational model is all you need. The simple answer is no. Ask these questions to get a fuller picture of what else a vendor has built into their solution.

- Which models (LLMs) does your solution use? And why?
- Besides a LLM and RAG, what other technologies does your solution include? And how is it structured?
- Will I get locked into using a specific LLM forever? Or is your solution flexible enough to allow changes as models evolve?

Look for a solution that uses and orchestrates a wide variety of models, and a vendor that can explain why some models might be preferred for certain tasks and use cases. In addition to the LLM and RAG, the solution should include robust security controls and safety measures to protect against malicious inputs and harmful outputs. The vendor should also offer flexibility in which models are chosen and should allow you to swap models later if another would improve performance.

How will your solution protect our data (and our customers' data)?

Security is always a top concern, and generative AI adds some new risks into the mix, such as prompt injection, which could allow a bad actor to manipulate the AI into leaking sensitive data, granting access to restricted systems, or saying something it shouldn't. Any AI vendor worth considering should have strong, clear answers to these security questions.

- How do you ensure that the AI agent cannot be exploited by a bad actor to gain unauthorized access to data or systems?
- How do you ensure that the AI agent cannot retrieve data it is not authorized to use?
- How does your solution maintain data privacy during customer interactions?

Look for a solution that can detect when someone is trying to exploit the system by asking it to do something it should not. It should also have strong security boundaries that limit the AI agent's access to data (yours and your customers'). Security and authentication in the API layer are especially critical for protecting data. And all personal identifiable information (PII) should be redacted before data is stored.

How do you keep your AI agent from ticking off my customers or damaging my brand?

We've all heard stories of bots that spouted offensive language, agreed to sell pricey products for a pittance, or encouraged people to do unsafe things. Solution providers worth considering should have robust safety mechanisms built in to ensure that the AI agent stays on task, produces accurate information, and operates ethically. Get the details on how a vendor approaches AI safety with these questions.

- How do you mitigate and manage hallucinations?
- How do you prevent the AI agent from sharing misinformation with our customers?
- How do you prevent jailbreaking?

Look for a solution that grounds the AI agent on information specific to your business, such as your knowledge base, and includes automated QA mechanisms that evaluate output to catch harmful or inaccurate responses before they are communicated to your customer. The solution should also incorporate a variety of guardrails to protect against people who want to exploit the AI agent (jailbreaking). These measures should include prompt filtering, content filtering, models to detect harmful language, and mechanisms to keep the AI agent within scope.

How hard will the solution be to use and maintain?

Conditions in a contact center can change quickly. Product updates, new service policies, modified workflows, revised knowledge base content, and even shifts in customer behavior can require your agents to adapt – including your AI agents. Ask these questions to find out how well a solution empowers your team to handle simple tasks on their own, without waiting on technical resources.

- What kinds of changes and updates can our contact center team make to the solution without pulling in developers or other technical resources?
- What will it take to train our supervisors and other CX team members to work with this solution?

Look for a vendor who has invested in user experience research to ensure that their solution's interfaces and workflows are easy to use. The solution should have an intuitive console that empowers non-technical business users with no-code tools to manage changes and updates on their own.

How will we know what the AI is doing – and why?

When a human agent performs exceptionally well – or makes a mistake – you can ask them to explain their reasoning. That's often the first step in improving performance and ensuring they're aligned with your business goals. It's equally important to understand how an AI agent is making decisions. Use these questions to learn how a solution offers insight into the AI's reasoning and decision-making.

- How will we know what specific tools and data the AI agent is using for each customer interaction?

- In what ways do you surface information about how the AI agent is reasoning and making decisions?

Look for a vendor who provides a high degree of transparency and explainability in their solution. The AI agent should generate an audit trail that lists all systems, data, and other information sources it has accessed with each interaction. In addition, this record should also include an easily understood explanation of the AI agent's reasoning and decision-making at each step.

How does your solution keep a human in the loop?

Solution providers acknowledge the importance of keeping a human in the loop. But that doesn't mean they all agree on what that human should be doing or how the solution should accommodate and enable human involvement. These questions will help you assess how thoroughly the vendor has planned for a human in the loop, and how well their solution will support a cooperative relationship between the AI and your team.

- What role(s) do the humans in the loop play? Are they involved primarily during deployment and training, or are they also involved during customer interactions?
- When and how does your genAI agent hand off an interaction to a human agent?
- Can the AI agent ask the human agent for the input it needs to resolve the customer's issue without handing over the interaction to the human?
- What kind of concurrency can we expect with a human in the loop?

Look for a solution with an intuitive interface and workflow that allows your human agent to provide guidance to the AI agent when it gets stuck, make decisions and authorize actions the AI agent is prohibited from doing on their own, and step in to speak with the customer directly as needed. The AI agent should be able to request guidance and then resume handling the interaction. The solution should be flexible enough to easily accommodate your policies for when the AI agent should ask its human coworker for help.

Why should we trust *your* team?

Trust depends on a number of factors, but it starts with expertise. What you really need to know is whether a vendor has the expertise to deliver a reliable solution now – and continue improving it for the future. These questions will help you determine which solution providers are best equipped to keep up with the pace of innovation.

- What components of your solution were developed in-house vs. acquired from third-parties?
- What kind of validation can you share from third-parties?
- Can you point me to your team's research publications and patents?

Look for a vendor with a strong track record of in-house development and AI innovation. That experience is a good indicator of the vendor's likelihood of continuing to expand their products' capabilities as AI technologies evolve. Patents, published research, and third-party validation from industry experts and top-tier analysts underscore the vendor's expertise.

This list of questions is not exhaustive. There's a lot more you could – and should – ask. But it's a good start for rooting out the details you'll need to make a fair comparison of generative AI agents.

Want to ask ASAPP some questions about GenerativeAgent?

Get in touch. We've got answers.

SCHEDULE A DEMO

ASAPP

ASAPP is an artificial intelligence cloud provider committed to solving how enterprises and their customers engage. Inspired by large, complex, and data-rich problems, ASAPP creates state-of-the-art AI technology that covers all facets of the contact center. Leading businesses rely on ASAPP's AI Cloud applications and services to multiply agent productivity, operationalize real-time intelligence, and delight every customer.