

# *guide to building vs. buying an AI agent for customer service*

- p 3* total cost of ownership
- p 6* return on investment
- p 8* evaluation framework
- p 10* continuous improvement
- p 11* expertise
- p 12* when to opt for an in-house solution

## introduction



For businesses looking to enhance customer service, the appeal of integrating an AI agent into their operations is undeniable; AI agents can automatically resolve inquiries around the clock, offering personalized interactions that improve overall customer experience. And thanks to the availability of powerful foundation AI models, building generative AI applications like AI agents for customer service is more accessible than ever.

But just because you can, does that mean you should?

While it may seem like building an in-house AI agent is the best path — after all, who knows your business and your customers better than your own team — the reality is not that straightforward.

Before diving into the development of an in-house AI solution, it is crucial to weigh several key factors against the benefits of working with established AI agent vendors. By understanding these factors, you'll be better equipped to make an informed decision that aligns with your organization's needs and resources.

# 1. total cost of ownership

How much does it actually cost to build and maintain an AI agent for customer service? And how does that stack up against paying an annual fee to a third party vendor? Let's tally it up.



consideration	building in-house	partnering platform provider
<b>In-house team</b>  The human talent and resources you'll need to successfully launch and manage an AI agent for customer service.	Dedicated team of 6+ FTE with specialized skills.	One AI Manager to work with the partner platform and manage the AI agent.  <b>Related reading</b> → <a href="#">Understanding the privacy, data and security risks of AI agents</a>
<b>Development</b>  What it takes to build an AI agent from a technical perspective.	You'll need: <ul style="list-style-type: none"><li>Product Manager to lead the project</li><li>Product Designer to design the AI agent interface</li><li>Front-end Developer to craft what the customer experiences</li><li>Back-end Developer to create the logic and systems that make it all work</li><li>Integrations Engineer to build out integrations with business systems, such as your CRM, CSM, CCaaS, telephony, and/or email service providers (alternatively, you can use <a href="#">integration platform software</a>, in which case you'll have to add in that cost)</li></ul>	You'll need to: <ul style="list-style-type: none"><li>Select the right AI agent partner for your needs and assign a team member to become the AI Manager</li><li>Onboard the AI agent (this process will look different depending on which AI agent you hire)</li><li>Connect the AI agent to business systems using out-of-the-box integrations</li><li>Be ready to deploy in 6 - 9 days</li></ul>
<b>AI model costs</b>  There's a common misconception that "publicly available" means "free". The reality is, you still have to pay to use the AI models that will be powering your AI agent, including: <ul style="list-style-type: none"><li>Large language models (LLMs)</li><li>Speech-to-text</li><li>Text-to-speech</li><li>Translation</li></ul>	Typically for AI models, you have to pay per token. Depending on the model, a token can mean different things, for example: <ul style="list-style-type: none"><li>Number of words in a sentence (usually around 2 or 3 is one token)</li><li>Minutes in an audio file (usually one minute is a token)</li></ul> Each model will have its own pricing scheme, but your cost per conversation will essentially depend on how verbose your customers are.  This makes it difficult to forecast how much the AI models will actually cost month over month.	AI agent platform providers have partnerships with AI companies, and will have worked out these costs into their pricing models.  Crucially, these pricing models are tuned to dimensions of your business that you already forecast, so you can predict costs confidently and work them out into your own economics.  <b>Related reading</b> → <a href="#">Critical capabilities for your AI agent</a>

Table continues on the next page

# 1. total cost of ownership

consideration	building in-house	partnering platform provider
<p><b>Data warehousing and storage</b></p> <p>While AI models power the conversations you have with customers, they don't store them. You'll need data warehousing to access the conversation history for things like analytics, insights, and auditing.</p>	<p>Prices vary significantly based on requirements.</p>	<p>AI agent platform providers store all data related to the AI agent, and the platform fee includes this cost.</p> <p><b>Related reading</b> → <a href="#">Understanding the privacy, data, and security risks of AI agents</a></p>
<p><b>Ongoing maintenance and continuous improvement</b></p> <p>Building and launching the AI agent is not a one-and-done commitment. In order to remain competitive, you'll need to ensure the AI agent is constantly improving, much like you'd expect your human agents to improve at their jobs the longer their tenure is.</p>	<p>You'll need:</p> <ul style="list-style-type: none"><li>• AI Manager to manage and coach the AI agent</li><li>• Engineer for maintenance, troubleshooting, and additional integrations</li><li>• ML Specialist to improve the AI agent's capabilities as the underlying tech evolves</li><li>• Data Scientist to build out the Automated Resolution reporting dashboard and audit the accuracy of the AI agent's resolutions</li><li>• Security Manager to ensure that the AI agent meets privacy and security standards of operation</li></ul>	<p>You'll need:</p> <ul style="list-style-type: none"><li>• AI Manager to manage and coach the AI agent</li></ul> <p>The partner platform will handle everything else, including:</p> <ul style="list-style-type: none"><li>• Updating the underlying tech</li><li>• Developing new AI agent skills and capabilities</li><li>• Adding new management features</li><li>• Building out easy-to-use insights and analytics dashboards</li><li>• Handle any maintenance, troubleshooting, and additional integrations</li></ul> <p>You'll benefit from a higher frequency of updates that are well-researched and have higher ROI.</p>
<p><b>Scaling to new channels</b></p> <p>Customers want to connect with you on every channel available — and they want to be remembered when they do. Omnichannel service is necessary to maintain a competitive advantage.</p>	<p>Creating an omnichannel AI agent is not an easy task. Developing for channels like voice and email involves distinct challenges, which can complicate and prolong the development process for in-house teams.</p> <p>For example:</p> <ul style="list-style-type: none"><li>• In order to build your own solution for voice conversations leveraging large language models, you'd also need to build an orchestration layer to handle things like speech transcription, speech synthesis, and integration with telephony</li><li>• Seamlessly handling the nuance of voice conversations requires extra optimization, such as allowing your customers to interrupt, knowing when callers have finished speaking, or accurately capturing data such as order numbers, reference numbers, and more</li></ul>	<p>You can partner with an AI agent platform provider that works seamlessly across and is optimized for all channels, including messaging, voice, and email. Deploying the AI agent on new service channels becomes a matter of "when", not "if" or "how".</p> <p><b>Related reading</b> → <a href="#">Guide to finding the right AI Agent for omnichannel customer service</a></p>

## 1. total cost of ownership

# the cost of staying up to date

The rapid pace of AI innovation makes it increasingly difficult for businesses to keep up without continual reinvestment.

What's cutting-edge today can be obsolete in just a few months, forcing companies to constantly evaluate and implement new technologies or risk falling behind.

As we move closer to artificial general intelligence (AGI), AI systems will not only evolve in complexity but also demand more nuanced applications.

Businesses that have built their own AI agents will quickly discover the hidden costs of staying up to date.

This happened before with scripted conversational AI. The companies that built their own in-house solutions are now having to start from scratch, reinvesting in development, integrating new tools, and scaling for omnichannel service — all of which are essential to ensure customer satisfaction and prevent outdated tech from undermining the customer experience.

It's happening again already, with LLMs evolving to CoT models like GPT-o1. These open up new opportunities but require new skill sets and architectures to take advantage of — many of which haven't been invented yet.

And it doesn't even need to be big decisions like that. For example, if you want to introduce new integrations to automate different parts of your customer service strategy, your development and engineering teams will need to manually set up those integrations, test them, and deploy them.

So as you're tallying up the costs, don't forget to factor in the longevity of the solution, and what kind of resources you'll require to stay up to date.



# 2. return on investment

Now that we've broken down the costs, let's talk about the results. Assuming both options are working comparably and delivering a similar Automated Resolution rate, which one is giving you a better bang for your buck?

## time to launch

With an investment as high-value as building an AI agent, there are a number of time-consuming milestones between getting approval to work on a solution and rolling it out.

At the very least, you'll need to identify the key stakeholders, gather the requirements for the solution, hire the team or create a task force in your existing departments, develop the solution, test and iterate to improve it, train the teams on using it, and finally roll it out.

Think about how much time each of those steps take, and what kind of effort is involved in the details of those steps as well.

How much time to recruit the talent if you don't have it in house? How many people are involved in the process? Once the team starts working, how much time will they need to ramp on internal processes?

Think about it this way, if you spend 6 months building a solution in-house, that's 6 months where you're not automatically resolving support inquiries and losing out on savings while you build.

By the time the AI agent is live, you could have already broken even on the initial cost and started seeing positive returns with a partner platform.

How long does it take to launch an AI agent for customer service?

6 - 9 months

The average length of a development cycle for in-house builds

6 - 9 days

To launch with a partner platform

## 2. return on investment

# opportunity cost

Beyond the financial investment, companies must also consider the opportunity cost of building an AI solution for customer service in-house. While it's true that every company should leverage AI, the focus should be on areas that directly align with core business objectives and growth.

For example, a FinTech like Neo would gain more by developing AI applications that enhance portfolio performance, or a SaaS like Monday.com could better serve its mission by focusing on AI that makes workflows more efficient.

It's also important to evaluate what critical projects or teams will be impacted by reallocating resources to AI development. Will your recruitment team now have to shift focus towards finding specialized AI talent, or will you outsource to agencies, incurring high costs just to keep your internal teams focused on core functions?

Diverting resources to build a customer service AI agent can be a distraction from key business priorities.

"One of the things that we've seen is people don't understand generative AI as well as they think they do.

We did a survey in April [...] of buyers of generative AI. These were people at brands who were part of the buying process or part of the buying decision making process. One of the questions was: "if I have a generative AI system and I ask it the exact same question, it will always give me the exact same answer." And then agree or disagree with that statement.

Without getting into the minutia of the function of generative AI, it talks, it says things differently every time. That's the beauty and the pain of it. And 70% of the people in that survey said, yes, it will give the exact same answer every time.

And, it's possible to do that, there's some ways you can tune it to force that. But at that point you kind of lost the whole conversational ability of it. You basically just typed a prompt, and it's saying the prompt verbatim.

*People don't understand generative quite as well as they should."*

Max Ball  
Principal Analyst

FORRESTER®



# 3. evaluation framework

Before you deploy the AI agent to your customers, you need a testing environment to ensure that it's actually working. Otherwise, you risk exposing your customers to a potentially faulty AI agent and you have to deal with the consequences of any problems that might cause.

## testing AI systems

Evaluating whether an AI agent is ready to handle customers

Building an in-house AI agent

In addition to the AI agent itself, you'll have to build your own in-house testing systems.

Partnering with an AI agent platform provider

Benefit from the partner platform's testing tools.

It's easy enough to build a demo that looks good, but once you scale to a larger conversation volume or more diverse customer inquiries, you'll find many scenarios that the AI agent is not equipped to handle.

To avoid this, you need to have thoughtful, structured evaluation frameworks for your system in place.

Without them, it's easy to fall into a false sense of security and then wonder why your AI agent doesn't perform as well in the real world as it did in your own testing.

"From a company culture perspective, NinjaTrader aims to be innovative and push forward new tools and services. So AI was just something that was in our DNA to explore, and make sure that we were leveraging that to give our clients the best experience."

From a compliance perspective and a comfortability perspective, the ability to put in SOPs and have a structure in place to monitor [compliance-related conversations] was really important. Before we launched anything, we had to make sure that we had all that in place, and tested it in various situations, looked through the responses, and made sure that the evidence is documented internally. This helped us get comfortable about opening this up to our customers."

Walter Seldz  
President, Revenue & Client Services




### 3. evaluation framework

## insights and analytics

It should go without saying that deploying an AI agent is useless if the AI agent doesn't work — or worse, works badly. Afterall, the only reason you're thinking of this solution is because you want to provide your customers with a better experience, and the AI agent is a means to that end.

And to that end, you'll want to make sure that the AI agent is actually doing the job you hired it to do to the best of its ability. That means monitoring its results, auditing them when you need to, and analyzing the insights from the conversations to keep improving AI customer service.



#### How are you auditing and analyzing automated resolutions?

##### Building an in-house AI agent

Your team will need dedicated tooling to track the data that's most relevant to your customer service KPIs in related to the AI agent program, including:

- Containment rate (i.e. the percentage of conversations that were not handed off to a human agent)
- Customer conversations with the AI agent
- Customer data
- CSAT
- NPS

You'll also need to manually audit the conversations marked as "contained" to see whether they were actually resolved or the customer simply gave up on the conversation.

##### Partnering with an AI agent platform provider

You can use purpose-built dashboards to:

- Get an accurate Automated Resolution rate (i.e. the percentage of conversations that the AI agent is resolving), with the ability to drill into why the AI agent is classifying a conversation as resolved or unresolved
- Understand the actions your AI Agent took so you can coach it to continuously improve performance
- Measure your AI agent's success across channels
- Track the impact on key business metrics
- Bring your data into business intelligence tools to gain a deeper understanding of the customer experience

# 4. continuous improvement

A large component of the success of the AI agent is how well your team is managing and coaching it. And to do that, your AI Managers need tooling that enables them to identify gaps and drive higher automated resolutions. Let's take a look at some considerations for continuous improvement.



consideration	building in-house	partnering platform provider
<b>New skills</b>  As the underlying technology evolves, the AI agent should be able to learn new skills that make the customer experience even better. How will you be adding these new skills?	You'll need a dedicated team to take on the development of new skills.	Benefit from a continuous rollout of new skills.
<b>Feedback loops</b>  How does your team identify what part of the customer experience needs improvement?	<p>The feedback loop will look different depending on how you set up the analytics dashboards and who is in charge of reviewing the data.</p> <p>The team would analyze the uncontained conversations and cross-reference them with CSAT and customer data to uncover gaps and opportunities.</p> <p>The cadence of the feedback loop needs to be determined based on level of effort and time needed from the team members involved.</p>	<p>The AI agent surfaces conversation insights and improvement opportunities, pulling conversation and customer data from all channels it's deployed on as well as relevant connected business systems.</p> <p>The AI Manager can spend their time addressing the opportunities instead of looking for them.</p>
<b>Feedback implementation</b>  Who receives the feedback to improve the AI agent? And how soon until it is implemented?	<p>The AI Manager identifies a list of improvement opportunities based on the loop above, then they schedule these improvements with the Engineering team in charge of improving the AI agent.</p> <p>The team then triages this request and adds it to the relevant sprint. Depending on their triage criteria, it could take anywhere between 2 weeks and 2 months.</p>	<p>With the right coaching tools, the AI Manager can communicate the feedback directly to the AI agent using natural language. The AI agent implements any coaching feedback instantaneously. This means you make changes to the AI customer service experience as fast as your business evolves.</p> <p>The AI agent may also have the ability to quantify the impact of each opportunity on key metrics like CSAT or Automated Resolution rate, so the AI Manager can easily build a business case to invest in solving the most pressing problems your customers face.</p>

Related reading → [Guide to automated resolutions](#)

# 5. expertise

## AI for customer service

If you were building a large-scale LLM-powered solution that potentially all your customers will interact with, you'll need people who are experts in AI, customer service, and AI for customer service.

This is why it's important to work with partners who have been immersed in this world for years. Not only do they understand your pains and how to solve them, they also have working relationships with software and technology companies that give their AI agents an edge.

## AI agent program expertise

In addition to technology expertise, AI agent platform providers have worked with companies across various industries to launch and grow AI customer service programs. They have tried and tested structures and best practices that you can benefit from to truly transform customer service, beyond simply automating FAQ.

Enterprises can benefit from their partner platform's onboarding and education processes to train agents, leading to higher implementation success rates and better value for the business.

"If I just get Google or I get OpenAI or I get something, and I start building [an AI agent for customer service], I can build a self-service application, but there's just so many things that can go wrong. There's guardrails that need to be put in place. You need to watch out for hallucinations. You need to understand response times.

It's very important that you get something that takes generative AI and puts it into a context of AI for customer service. You need the right set of tools to do that. You don't want to just say, "IT go build me this stuff".

You need people who've been thinking about the guardrails, who've been thinking about the problems and living that stuff for a significant chunk of time.

You need a partner. Don't just say, "Hey, I'm going to take OpenAI, I'm just gonna build. This is going to be great." There's a lot of nuance. There's a lot of subtlety. There's a lot to think about. You need somebody who lives and breathes this stuff, who understands it."

Max Ball  
Principal Analyst

FORRESTER®



# when to opt for an in-house solution

While it may seem that there's a strong bias towards working with an AI agent platform partner, there are scenarios in which it actually makes more sense to build an in-house solution.


## You need a highly specialized AI agent

If you need an AI agent to take actions that are unique to your business, you're likely to get more value by building it yourself. While a partner platform can definitely build out specialized use cases, you would be bound by their development timelines.

For example:



Industry	Target customer	Specialized actions
SaaS	Developers	Analyze snippets of code for errors.
Travel	Vacationers	Build an itinerary based on travel dates, destination, and budget.




## You need complete control over the AI agent development roadmap

In the same vein as the points above, the biggest benefit of building an in-house solution is that you can customize it entirely to fit your business's exact needs. Perhaps your entire customer base prefers a specific communication channel, so you don't care about omnichannel capabilities and instead want to further optimize that one channel. Or perhaps you have multiple highly specialized customer intents that make up the majority of your conversation volume.

Whatever the case may be, building your own AI agent gives you the ability to prioritize the skills that make the biggest impact, and if the impact is big enough, it'll offset the costs of an in-house solution.

# meet your new AI agent for customer service

With ada, you get a platform that empowers you to deliver efficient, high-quality support at scale—continuously improving the speed and quality of results across all channels and languages, 24/7.



**5.5B**

customer interactions  
powered

**8X**

more productive  
than human agents

**83%**

up to 83% of  
conversations  
autonomously  
resolved

**99.9%+**

uptime of  
uninterrupted service

## accelerate your business with AI

[Book a demo](#)



**ada**