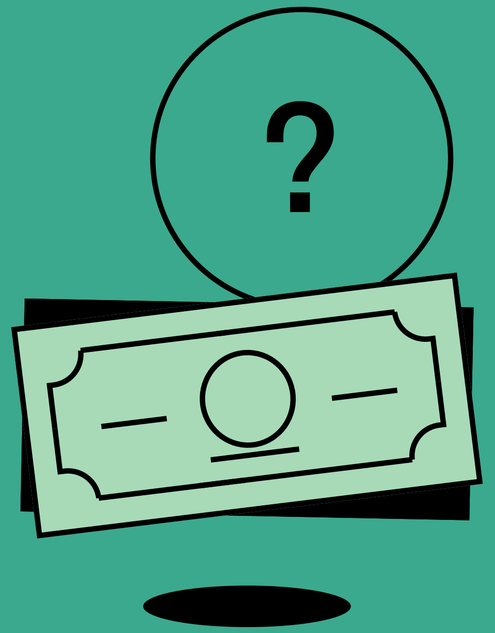
Assembled

How (most) AI support tools are priced, and the questions you need to ask before buying



You've put in the work to find the perfect AI tool for your support team — now comes the real challenge: making sense of the pricing structure and figuring out how it fits into your budget. AI tools promise to boost productivity and cut costs, but let's be real, their pricing can be a bit of a puzzle compared to traditional software.

Built with support leaders like you in mind, this guide breaks down the different AI pricing models and arms you with the right questions to ask. The goal? To help you cut through the noise and make sure your investment delivers the value your team deserves.

How AI support tool pricing differs from other CX software

AI support tools break away from the straightforward per-seat pricing model we're all familiar with. Older tools like contact platforms or workforce management software typically charge based on how many people are using or being scheduled in the system. This makes it simple to estimate and scale your software costs: more users mean higher costs. Why? Because the value of these tools is tied directly to the number of users engaging with them.

AI support tools, however, march to a different beat. Their pricing often focuses on outcomes or feature usage rather than the number of seats (though there are exceptions — we'll get to those). This shift happens because the value of automation-focused tools doesn't depend on how many people use them. In fact, an AI support tool often delivers more value when fewer people are needed to achieve results.

Think of it this way: AI tools are less like traditional software and more like a BPO (business process outsourcing) partner. Instead of charging for user access, they're compensated based on the amount of work they handle for your team. In essence, AI pricing reflects the outcomes it delivers, not just the number of users who log in.

This mindset can also help when you're thinking about how AI fits into your broader operations — but more on that later!

The spectrum of AI support tool pricing

When it comes to AI support tools, pricing isn’t one-size-fits-all. Instead, you’ll find a range of models designed to align with different use cases and outcomes. Here are the main ones you’re likely to encounter.

Per-user pricing (agent copilots)

AI tools like agent copilots, which assist your team in drafting responses or finding answers, often stick with the traditional per-user pricing model. It’s straightforward: the more agents you have, the higher your cost.

However, not all actions are created equal. Some are proactive, triggered automatically based on pre-set conditions, while others are reactive, initiated by the agent within the app. Because of these nuances, it’s worth clarifying with your vendor how (or if) pricing varies based on these different use cases.

Ask your vendor

- Is pricing tied to agent activity or just the total number of seats?
- Are there extra fees for usage spikes or high-volume interactions?

Usage-based pricing (AI agents)

AI tools that fully automate customer interactions typically charge based on the number of interactions or resolutions handled by the AI.

This model often works like a prepaid plan: you pay upfront for a set number of automated resolutions. While this approach can be flexible, it’s important to watch out for hidden costs like overage fees, which can sneak up on you if your usage exceeds the prepaid amount.

Ask your vendor

- How is a "resolution" defined?
- Are there overage fees, and how are they calculated?
- Do unused credits roll over to the next billing period?

Interlude

Understanding channel-specific AI pricing

When evaluating AI support software, don’t overlook how pricing can shift depending on the communication channel. Vendors often charge different rates for phone, chat, email, and API interactions because each channel comes with its own complexities and associated costs.



Voice AI agents

AI tools that manage phone interactions often come with higher per-interaction costs. Why? Handling voice requires advanced speech recognition, real-time natural language processing, and adherence to strict compliance standards. Unlike text-based channels, voice AI has to process and respond instantly, making it more resource-intensive.

Ask your vendor

- Are phone interactions priced higher than chat or email resolutions?
- Is there an additional cost for maintaining voice transcription logs?
- Does pricing vary based on call length or complexity?



Chat AI agents

Chat AI tools usually come with more predictable pricing than phone AI. These interactions are easier to automate since they don’t rely on voice recognition. However, costs can still vary based on the complexity of the conversation flow and the level of personalization required.

Ask your vendor

- How are multi-turn conversations counted in pricing?
- Are proactive chats (initiated by the AI) priced differently from reactive ones?
- Is pricing based on individual messages or resolved chats?



Email AI agents

AI tools for email interactions typically excel at automating ticket categorization, auto-responses, and prioritization. Pricing for these tools is often based on the volume of emails processed or the number of resolutions achieved, making it essential to understand how your usage aligns with the vendor’s model.

Ask your vendor

- Are there limits on the number of emails the AI can process?
- How is a "resolved" email defined?
- Is there a discount for handling bulk email volumes?

API Interactions

For companies integrating AI via APIs — whether to automate backend processes or programmatically handle support tickets — pricing is typically usage-based. This is often measured by the number of API calls or the volume of data processed, providing flexibility but requiring close attention to usage patterns.

Ask your vendor

- Are API interactions priced per request or by data volume?
- Are there usage limits, and what happens if you exceed them?
- How does the vendor ensure fair pricing for high-frequency API users?

Why usage definitions matter

When it comes to AI pricing, the term “resolution” can mean very different things to your business and your AI vendor. For some tools, a resolution might be any automated interaction — even if the case ultimately gets escalated to a human agent.

If you’re unsure about your projected usage, consider requesting a proof of concept (POC) period. This can give you a clearer picture of costs and usage patterns before committing to a long-term contract.

Ask your vendor

- Will a partial automation still count as a resolution?
- How are complex, multi-step cases handled in the pricing model?

Credit packages vs. pay-as-you-go

AI vendors typically offer two main pricing models: prepaid resolution packages and pay-as-you-go plans. Prepaid packages can give you predictable costs, which is great for planning. On the other hand, pay-as-you-go models might be a better fit if your support volumes vary throughout the year.

Ask your vendor

- Are there discounts for committing to larger volumes?
- What happens if I exceed my package — are overages charged at a different rate?
- Do credits expire, or can I roll them over?



Honeylove sees a 54% increase in productivity with Assembled Assist

Read the story

The secret costs of AI: Implementation and support

AI support tools might promise automation magic, but they're far from plug-and-play. Ongoing configuration, training, and support are essential to keep them running smoothly and optimized for your business needs. These hidden costs can add up, so it's important to account for them when evaluating pricing.

Implementation costs

Implementation can be a significant part of your AI investment. Some vendors bundle it into their pricing, while others charge it as a separate service. This process often involves fine-tuning the AI to fit your specific use cases, ensuring it's ready to deliver value from day one.

Ask your vendor

- What's included in the implementation fee?
- How long will it take to deploy the AI tool?
- Are there additional costs for complex use cases?

Ongoing support

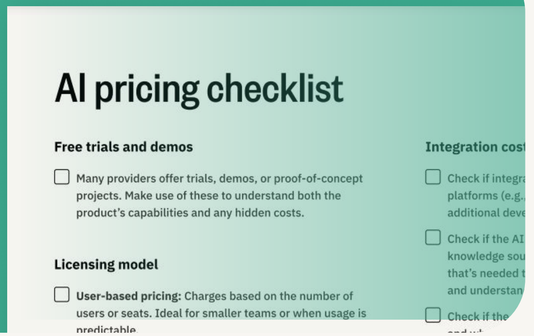
AI tools aren't a set-it-and-forget-it solution — they need regular updates and monitoring to stay effective. Understanding what kind of support is included in your contract (and what might cost extra) is crucial to avoid surprises down the line.

Ask your vendor

- How will we receive ongoing product support?
- Are new features included in our subscription, or are they add-ons?
- Will I need to pay for retraining the AI as my business evolves?

Download your copy of our AI pricing checklist

Get the checklist



Making the most of your AI investment

Investing in an AI support tool isn't just about spending money — it's about unlocking the potential to transform your customer experience operations. By taking the time to understand pricing models and asking the right questions, you'll set yourself up to choose a solution that delivers real, lasting value for your team and your business.

> Want to learn more about the real-world impact of AI in support? Check out case studies from companies like Thrasio and Honeylove, who are using AI tools to achieve transformative results.

