AI-powered support: What to know before you buy

AI is revolutionizing customer support — but making it work takes more than just picking the right tool. Success starts with a solid foundation: a well-structured knowledge base, a clear implementation strategy, and a thoughtful evaluation process. Get these right, and AI unlocks a new set of possibilities. Get them wrong, and you risk adding friction instead of reducing it.

Support leaders are turning to AI to boost efficiency, eliminate repetitive tasks, and enhance both agent and customer experiences. When done right, AI can surface information faster, automate routine inquiries, and provide real-time guidance — freeing up teams to focus on complex, high-value interactions. But AI isn't a magic fix. Poor data quality, lack of system integration, and misaligned expectations can all stand in the way of real impact.

This guide covers everything you need to know before investing in AI for support, from optimizing your knowledge base to evaluating pricing models, proof-of-concept testing, and vendor selection. A strategic approach ensures AI isn't just another tool, but a true force multiplier for your team — today and in the long run.

Optimizing your knowledge base for AI

AI is only as good as the information it can access. A well-structured, up-to-date knowledge base is the back-bone of any AI-powered support system. Without it, AI tools risk delivering outdated, irrelevant, or inaccurate responses — frustrating customers and agents alike.

Want to go deeper? Learn how to build and maintain a knowledge base that powers AI with accuracy and efficiency.

Read the full guide ⬈

How to structure your knowledge base for AI success

To ensure AI surfaces the right answers, your content needs to be clear, structured, and easy to navigate. Here's how:

- **Use simple, direct language:** AI processes straight-forward language more effectively, improving response accuracy.

- **Break down complex topics:** Organize information into digestible sections to enhance readability.

- **Standardize formatting and tagging:** Use consistent headers, categories, and labels so AI can recognize patterns and retrieve relevant content.

Keeping your knowledge base up to date

Optimizing a knowledge base isn't a one-and-done project — it requires ongoing effort. Regular updates, feedback loops, and maintenance keep information fresh and relevant.

Red flags to watch for

- Outdated or conflicting information that leads to AI surfacing incorrect responses.

- Lack of structure, making it difficult for AI to retrieve the right answers.

- No system for continuous updates or agent feedback.

AI is only as strong as the information behind it. A well-structured, well-maintained knowledge base isn't just helpful — it's essential for AI to deliver meaningful results.

Breaking down AI support pricing models

AI-powered support tools come with a range of pricing models, and the details can have a big impact on cost, scalability, and overall ROI. While some pricing structures seem straightforward, hidden fees and vague value metrics can make it tricky to compare options.

Red flags to watch for

- Overcomplicated pricing structures that make cost comparisons difficult.

- Hidden fees that significantly increase total spend.

- Vendors who can't provide clear pricing models upfront.

- No accountability for AI quality — if a tool charges for every response, regardless of accuracy or usefulness, you may be paying for low-value interactions.

Common AI pricing models and how they impact cost

Here's a breakdown of the most common pricing models you'll encounter:

- **Per-seat pricing (for AI copilots):** Charges are based on the number of agents using the tool. Predictable but may not scale efficiently if usage fluctuates.

- **Usage-based pricing (for AI agents and copilots)** Costs depend on conversation volume, API calls, or interactions. Some vendors charge per session (each AI interaction), while others charge per resolution (only when an issue is fully resolved).

- **Flat-rate pricing (for AI agents and copilots):** A fixed cost for unlimited use. Easy to budget for but may not be cost-effective for smaller teams.

- **Feature-based pricing (for AI agents and copilots):** Higher tiers unlock more advanced AI capabilities. Important to evaluate whether those extra features are worth the price.

- **Outcome-based pricing (for AI agents and copilots):** Costs are directly tied to outcomes, like completing a phone call or resolving a chat conversation.

Curious about AI pricing? Get a detailed breakdown of common pricing models, hidden costs, and how to evaluate ROI.

Explore the full pricing guide ↗

How to evaluate pricing beyond the base cost

Before committing to a tool, get clarity by asking vendors:

- **What's included in the base price?** Are key features locked behind higher tiers?

- **Are there extra fees?** Look for costs related to integrations, implementation, data storage, security, overages, or premium features.

- **How does pricing scale as your team grows?** Does the cost increase linearly or exponentially?

- **What happens if usage spikes unexpectedly?** Are there penalties or auto-upgrades?

- **How does the tool ensure AI quality before charging for responses?** Can you set quality thresholds, review automations before they go live, and refine AI outputs to meet your team's standards?

〜

Beyond the sticker price, understanding total cost of ownership (TCO) is key. A lower upfront cost doesn't always mean a better long-term investment — what matters is whether the pricing aligns with your team's needs, quality standards, and usage patterns.

Validating AI with a proof of concept (PoC)

Before committing to an AI-powered support tool, a proof of concept (PoC) is the best way to test whether it delivers real value. Running a PoC lets teams evaluate AI in a real-world setting, measure its effectiveness, and identify any gaps — before making a long-term investment.

Setting up a meaningful PoC

A successful PoC isn't just about seeing AI in action — it's about proving whether it meets your business needs. To get real insights, define:

- **Clear success metrics:** What does success look like? This could mean improved resolution rates, reduced handle times, or increased agent adoption.

- **Realistic use cases:** Test AI in real support workflows, not just ideal scenarios. Check for sentiment accuracy, handling of misspellings, and translation capabilities.

- **A controlled test environment:** Run the PoC with a subset of tickets, interactions, or agents to measure impact at a manageable scale.

A well-executed PoC isn't just about validating an AI solution — it's about setting your team up for long-term success. When AI proves its value in a real-world setting, you'll have the insights and internal buy-in needed to scale with confidence.

Red flags to watch for

Not all AI solutions will make it past the PoC stage. Be cautious if:

- The tool requires extensive manual intervention to produce useful results.

- AI-generated responses are inaccurate, irrelevant, or difficult to refine.

- The vendor isn't willing to customize the PoC to reflect real-world conditions.

- You need vendor engineering resources to configure workflows or brand guidelines.

Ready to test AI before you buy? Our step-by-step guide walks you through running a PoC that ensures AI actually works for your team.

See the full PoC guide ⊠

Navigating the AI support RFP process

For larger or more complex AI support investments, a Request for Proposal (RFP) can help cut through the noise — ensuring you're evaluating vendors on more than just flashy demos and bold claims. But to run an effective RFP, you need clarity on what your team actually needs from AI.

When to issue an RFP

An RFP makes sense when:

- You need a structured way to compare multiple vendors.
- Your organization requires a formal procurement process.
- The AI solution will play a critical role in support operations, requiring security, compliance, or integration vetting.

Key sections to include in your RFP

- **Business needs and objectives:** What specific challenges are you trying to solve with AI? Whether it's reducing handle time, improving deflection rates, or enhancing agent productivity, define your goals upfront.

- **Required AI capabilities:** Outline must-have features like automation, integrations with existing systems, reporting and analytics, and customization options.

- **Security, compliance, and data privacy:** Ensure the vendor meets your organization's security and compliance standards, especially if handling sensitive customer data.

- **Scalability and support:** Can the solution grow with your team? What level of customer support does the vendor provide?

How to evaluate vendor responses

Look for clear, specific answers — not vague promises. Strong responses should include real-world examples, measurable outcomes, and transparency around pricing and implementation timelines.

Red flags to watch for

- Overpromising AI capabilities without proof of success.
- Lack of flexibility in pricing or contract terms.
- Unclear data privacy policies.
- No clear plan for onboarding and ongoing support.

A well-structured RFP does more than just help you compare vendors — it lays the groundwork for a solution that fits seamlessly into your support strategy. The more clearly you define your needs, the easier it will be to make a confident, future-proof decision.

Need to run an RFP? Learn how to structure your request, compare vendors, and avoid common pitfalls.

Check out our RFP guide ⧉

Key factors to consider before investing in AI support tools

Not all AI support tools are created equal, and choosing the right one goes beyond flashy features. The best AI solution aligns with your team's needs — not just today, but in the long run.

Seamless integration with existing tools and workflows

AI should work with your existing tech stack, not against it. Look for integrations with your help desk, CRM, and other essential tools to avoid manual workarounds that slow your team down.

Scalability: Will the AI tool grow with your team?

What works now might not work a year from now. Can the AI tool support multiple channels? Handle increasing ticket volumes? A scalable solution prevents costly migrations down the line.

Transparency and control: AI should work the way you do

Your brand reputation is on the line with every interaction — so don't leave it up to rigid AI templates. The right tool gives you full control over customization, letting you replicate your best agents and deliver AI-powered support that meets the same high standards as your human team.

Data privacy and security: No compromises

AI tools process sensitive customer data, so security isn't an afterthought — it's a requirement. Ask vendors about compliance certifications, data storage policies, and how they handle personally identifiable information (PII).

Reporting and analytics: AI shouldn't be a black box

To measure performance and prove ROI, you need strong reporting and analytics. Look for dashboards that provide actionable insights, not just surface-level metrics, so you can refine AI responses and optimize operations.

Customer support and implementation: A strong partner matters

Even the best AI tools require onboarding and ongoing optimization. Does the vendor offer hands-on support, training, or a dedicated success manager? A strong partnership can be the difference between an AI investment that thrives and one that stalls.

Red flags to watch for

- Limited or no integration support, requiring manual workarounds.
- A rigid pricing model that doesn't scale with your team's needs.
- Vague or incomplete security and data privacy policies.
- Weak reporting tools, making it difficult to track AI performance.
- Lack of dedicated customer support, leaving you to troubleshoot on your own.
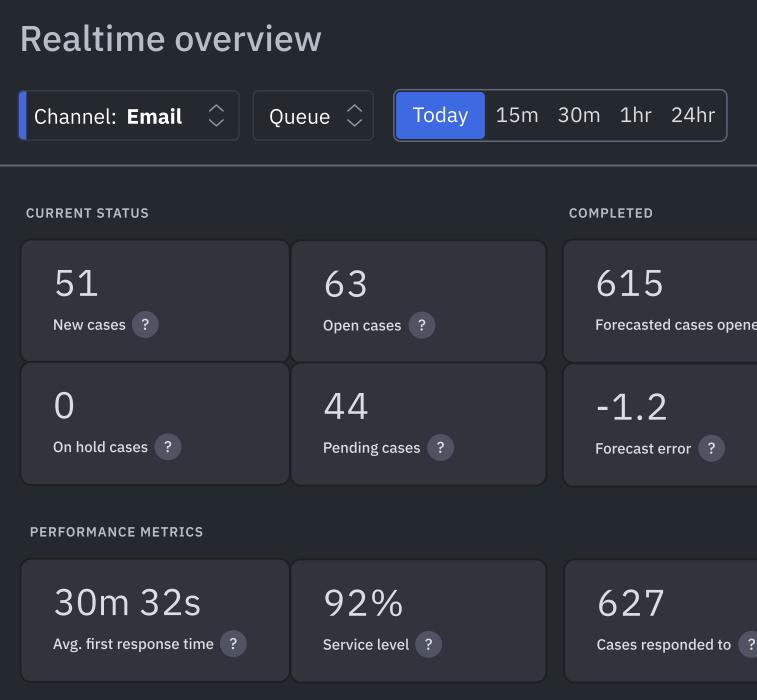
A well-vetted AI solution becomes a true asset to your team and not just another tool to manage. By focusing on integration, scalability, security, analytics, and support, you'll make an investment that drives long-term efficiency and impact.

See AI-powered support in action

Implementing AI in customer support can be transformative, but success hinges on choosing the right solution. Assembled Assist is designed to integrate seamlessly with your existing systems, ensuring scalability, robust data security, and insightful analytics. With features like omnichannel AI agents and an agent copilot, Assist empowers your team to deliver exceptional support across chat, email, and voice channels.

Ready to see how Assembled Assist can elevate your support operations? Explore our self-guided demo today.

Realtime overview

Channel: **Email**   Queue   **Today** 15m 30m 1hr 24hr

CURRENT STATUS

51
New cases ?

63
Open cases ?

0
On hold cases ?

44
Pending cases ?

COMPLETED

615
Forecasted cases opened ?

-1.2
Forecast error ?

PERFORMANCE METRICS

30m 32s
Avg. first response time ?

92%
Service level ?

627
Cases responded to ?