



GUIDE

How to implement human-centered AI agents in healthcare.



CONTENTS

Introduction

01. 02. 03. 04. 05.

Pre-build

Design

Build

Test and
launch

Post-launch

Conversational AI implementation in practice

Conclusion

INTRODUCTION

Most healthcare providers today are embarking on a digital transformation journey.

However, this journey is complex. Many hospitals and clinics still rely on aging phone systems and EHRs and must comply with strict privacy and security standards such as HIPAA, making it difficult to scale high-touch patient support.



Patient expectations

At the same time, patient expectations are higher than ever. Today's patients expect seamless, responsive service from their healthcare providers, just as they do in their everyday lives. AI agents have emerged as a solution to handle high call volumes and deliver faster, more personalized interactions at scale.



The reality for healthcare leaders

AI promises more personalized patient service and smarter insights. But for patient access managers and contact center directors, the day-to-day priority is simply keeping operations running. Staffing shortages, outdated systems, and patchwork technology leave little room for new projects. IT teams are already stretched between rolling out new tools and maintaining core systems. Without the right approach, AI can feel like added complexity rather than relief.



A clear roadmap to better patient access

This guide lays out a practical path for bringing conversational AI into your healthcare contact center. It covers every stage, from planning and design to testing and ongoing improvement. With this approach, you can scale patient services without a full infrastructure overhaul. Along the way, you'll improve patient experience, ease staff workload, stay compliant, and boost efficiency.



01.

Pre-build

Before you build your AI agent, it's important to take stock of where you are today. In this pre-build phase, you will assess your current patient engagement processes and assemble the right team and resources.

Completing the following five exercises will help you develop a conversational AI strategy aligned with your organization's current state and future vision. This groundwork keeps everyone focused on what matters most to your business.

EXERCISE 1

Define your change management strategy

Implementing any new technology in a healthcare organization will introduce change to operations, processes, and people. It's important to proactively manage that change.

Understand the “why” behind your conversational AI project and communicate it clearly

- Handle after hours call
- Decrease did-not-attends
- Increase patient satisfaction

Successful AI projects start with a clear direction. That means going beyond the tech and thinking about how AI supports your broader business goals.

If you can't connect AI-driven improvements to tangible business results, it will be hard to justify the upfront investment.

Support people through the transformation

Success goes beyond choosing the right technology. It starts with preparing your organization for the changes ahead. Focus on the impact on people, processes, and culture.

You need to create space and time for:

Learning: Help employees understand the technology you are investing in and how it will empower them.

Adaptability: Be clear about how AI will change roles, free up time, and allow employees to focus on more complex, high-value work.

Ongoing communication: Communicate openly throughout implementation to highlight what's working, where there's friction, and how to adjust.

Remain adaptable and deliver quick wins

To remain adaptable without losing focus, plan sprints, stay close to performance metrics and insights, and communicate effectively with all stakeholders.

To gain early momentum, focus on use cases that can deliver quick, tangible results without disrupting existing operations. Look for projects that are:

- Low in complexity but highly visible to stakeholders
- Backed by clear KPIs
- Able to demonstrate value quickly and build internal confidence

At PolyAI, healthcare customers see that up to 20% of calls are general FAQs, which can be automated in as little as two weeks.

By setting a clear vision, engaging people from the outset, and staying agile, you'll create a culture that embraces innovation instead of fearing it. This cultural readiness is just as important as the technology itself in achieving a successful conversational AI launch.

EXERCISE 2

Build your implementation team

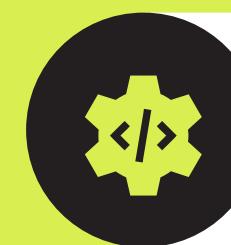
Implementing conversational AI doesn't have to be complicated. It should be a collaborative project where each team member plays a unique role in driving the overall success of the implementation.

We know time is limited in healthcare, so we keep the process focused and manageable for your team.

Your implementation team can be a mix of your staff and PolyAI experts, depending on what works best for you. Many healthcare organizations choose to offload parts of the setup to our team so theirs can stay focused on patients.



Patient engagement champion



Contact center technical lead



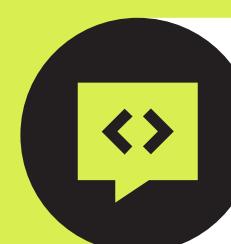
EHR/Systems integration specialist



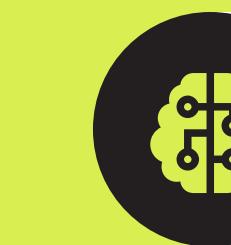
Data analyst



Project manager



Conversation designer



Machine learning lead



Executive sponsor



Patient engagement champion

This person knows the patient interactions you plan to automate. They can outline call types (e.g., scheduling, refills, billing), the steps staff follow, the systems used, common patient FAQs, and where improvements could be made. Based on frontline experience, they can also suggest improvements to the current call process.



Contact center technical lead

This individual manages your contact center technology stack (phone systems, IVR, etc.) They'll guide how calls are routed to and from the AI agent and ensure compatibility with existing infrastructure.



EHR/Systems integration specialist lead

This person provides access to backend systems and APIs. They enable backend system access (e.g., EHR, scheduling, billing) through secure APIs. They also ensure that data exchange is HIPAA-compliant and that the AI agent can retrieve or update patient information as needed.



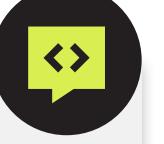
Data analyst

Quantifying the benefits of conversational AI is essential. A data analyst should gather current and historical metrics from your contact center. They will establish the baseline and later measure the impact of conversational AI (e.g., reduction in call handle time, increase in self-service rates).



Project manager

With multiple roles in your team, having someone responsible for keeping the project on track is essential. Your project manager is responsible for coordinating resources and organizing tasks.



Conversation designer

A conversation designer shapes how the AI interacts, making sure it's empathetic and natural. They create call flows and scripts that feel human to patients while following clinical communication best practices.



Machine learning lead

Leads the team of engineers who manage LLMs, training, safety guardrails, dialogue management, speech recognition, and spoken language understanding. They should have experience deploying large-scale AI agents in an enterprise setting.



Executive sponsor

If possible, bring a colleague from the leadership team into the project as soon as possible. Acting as the executive sponsor, this person can provide a high-level strategic view and help secure the necessary resources and buy-in from other stakeholders.





EXERCISE 3

Map out your current patient journey

Before you start mapping AI to your patient journey, your organization must ask: What do we want the patient journey to feel like? That vision is your north star. It keeps your AI efforts grounded in real outcomes, not just technology for technology's sake.

If your existing systems already have self-service capabilities, like digital appointment booking, it may not make sense to rip them out and start again. For example, conversational AI can be used to add more self-service capabilities alongside your IVR and to route calls from your AI agent to where self-service and digital resources already exist.

By mapping out your current experience and determining how conversational AI will fit in, you're taking a crucial first step toward understanding what "better" looks like. This exercise will also surface any dependencies (like required system access) and ensure you don't overlook parts of the patient journey when designing the new solution.



Gather the following resources to map out your existing experience.

What you'll provide		Why?
Call reports	The number of calls you receive per year and the breakdown of call types.	This process will highlight the high-volume, routine calls that drive the most value when handled by conversational AI.
Metrics	Your key contact center performance metrics such as average handle time, average speed to answer, abandonment, and agent utilization.	This data will help you identify areas for improvement. Once your AI agent is deployed, you can use this data to demonstrate its impact.
Agent training material	Documentation about the process agents use to complete calls and any systems used to enter and retrieve data.	Understanding this workflow is essential for designing an AI agent that seamlessly integrates with existing processes.
Contact center architecture	Diagrams for telephony & digital chat architecture and network connectivity.	You can identify potential integration points and dependencies by visually mapping out the connections between different systems.
API documentation	Documentation of all available APIs and input & output data formats.	This enables your voice team to understand your available APIs and how to integrate the conversational AI platform with your existing infrastructure and applications.
Security questionnaire	Questionnaire to identify any specific security concerns.	Ensure that every customer interaction meets your company's security and compliance requirements.
Analytics and transcripts	Call recordings or transcripts of customer calls and digital interactions.	To highlight customer behavior and preference patterns, understand specific needs and pain points, and identify common issues or recurring customer problems.
Screenshots and call recordings	Screenshots of agent desktops and systems used.	Screenshots provide insight into the workflow followed by agents during calls. By examining all aspects of agent interactions captured in screenshots, your voice AI vendor can identify key touchpoints and decision-making processes that shape the call-handling process.



EXERCISE 4

Gather baseline data and set clear objectives

Start by reviewing your current data and defining clear goals for introducing conversational AI. Know what success looks like and be realistic about what AI can achieve at each stage.

Choose the right starting point

Pick a call type that needs immediate attention, such as high-volume, routine calls that take up too much agent time. Decide what a meaningful improvement in your metrics would be, focusing on achievable and measurable outcomes. For example, you might begin with prescription refill calls, aiming for a specific containment rate before tackling more complex interactions.

Scale thoughtfully and stay compliant

The size of your initial deployment matters. Too large, and it may struggle to launch; too small, and it won't make an impact or justify the investment. Set targets that are specific, realistic, and aligned with regulatory and quality standards. Reducing handle time is helpful, but never at the expense of patient safety, required verifications, or quality of care.





Here are some goals you may want to consider. Remember to be specific and define the before-state as well as the after-state.



Improve appointment slot utilization by X%

Proactively reach out to patients to confirm, reschedule, or cancel appointments, helping your clinic fill open slots and make better use of provider time.



Reduce patient no-shows by Y%

Automate reminders and rescheduling to ensure patients show up or adjust their appointments in advance, keeping schedules on track.



Increase resolution of billing and collection reminders by Z%

Handle payment reminders and balance inquiries with sensitivity and consistency, helping more patients resolve their bills without human follow-up.





EXERCISE 5

Technology audit: Understanding your technical architecture journey

At this stage, the goal is to map and document the high-level system architecture required for calls or digital interactions to reach your conversational AI solution, how calls will be transferred, and identify API integration points for data.

In a healthcare environment, you'll also need to consider any additional security and compliance requirements for handling patient data in this architecture.



You can break your technical discovery into 3 key areas:

Contact center platform	Voice delivery	EHR and other system integrations
<p>What platform are you currently using? Does it support features like agent “screen pops” to enhance efficiency?</p> <p>Where does your solution sit? Is it on-site or in the cloud? And who’s managing it? If your call center is tightly integrated with your EHR, such as Epic’s call scheduler, note that as well.</p> <p>A contact center platform isn’t required to implement conversational AI, but if one is needed, flexible options are available to support your goals.</p>	<p>Decide whether you prefer using PSTN or SIP for calls. Can voice be delivered via SIP over the internet?</p> <p>Check if you can send custom header information in the SIP header and transfer the information back to your contact center. SIP headers will enable you to send messages in a screen pop alongside the call transfer to assist your agents with a more seamless handoff. If using SIP or any internet voice, ensure it’s encrypted (mutual TLS) and compliant with HIPAA security for voice data in transit.</p>	<p>You can automate many calls without touching sensitive patient data, but the real impact comes when your AI agent can access systems like the EHR. To handle tasks such as checking appointments, verifying patient identity, viewing balances, or updating records, you’ll need APIs. Plan these integrations carefully and ensure data is protected with encryption and strict access controls.</p> <p>REST and JSON are just the start. In healthcare, true interoperability depends on APIs built on standards like FHIR and HL7, which translate raw data into clinically usable information. Whether you’re integrating with an EHR or automating patient workflows, healthcare-specific APIs are essential.</p>





Your pre-build checklist

Your pre-build checklist	
Define your change management strategy	Clarify and communicate your project's 'why'. Engage stakeholders (agents, managers, clinicians, IT, executives) early and foster a culture open to innovation. Plan for training and communication to encourage adoption, and remain flexible to incorporate feedback as you go. <input type="checkbox"/>
Build your implementation team	Assemble experts across patient access operations, telephony/IT, integrations (EHR/CRM), data analysis, project management, design, and machine learning. Identify which roles will be filled by internal staff and which by the conversational AI vendor <input type="checkbox"/>
Map out your current patient experience	Document current call flows and patient experience. Understand how your IVR (if you have one) works today and identify opportunities for conversational AI to enhance or replace parts of it without disrupting what already works. Gather scripts, FAQs, and metrics on current performance. <input type="checkbox"/>
Gather baseline data and set objectives	Establish your key metrics (e.g., abandonment rate, wait time, containment, AHT, patient satisfaction). Collect current baseline numbers and set realistic target improvements for each. Ensure these objectives prioritize patient experience and align with your organization's goals. <input type="checkbox"/>
Complete a technology audit	Map out how the AI agent will technically integrate. Determine your contact center platform details and any needed upgrades, decide on PSTN vs SIP for call audio, and identify necessary API integrations for accessing patient info. Pay special attention to compliance (HIPAA) in your architecture, including secure data exchange and system access controls. <input type="checkbox"/>





02.

Design

With technical architecture understood, the next phase is designing a voice experience that truly engages and helps your patients. Designing for voice automation in healthcare requires careful attention to conversation flow, empathy, and clarity.



Designing for automation

Conversation design sounds simple in theory. Staff talk to patients every day, surely we know how a conversation should go?

While common sense is important in conversation design, it's not enough to create engaging systems. That's because we don't always speak to automated systems in the same way we speak to other people.

And with no visual interface to help callers through complex transactions, design know-how is key.

Make sure you work with a conversation designer who has a proven track record of designing conversations for enterprise AI agents at scale.

In healthcare, empathy and clarity are paramount, so your design should be reviewed by those who understand patient needs (for example, patient advocates or experienced call center nurses) to ensure it sets the bar for a good "bedside manner" over the phone.

**Empathy isn't optional;
it's expected**

67%

say, in patient intake, it's very or extremely important for AI to:

- Recognize emotional cues (distress, hesitation, urgency)
- Respond appropriately in tone and outcome.



Here are five key tenets of voice design to guide your process:



First impressions matter

The first thing your AI agent says plays a huge role in how (and if) a patient will interact with the system. When designing an AI agent, pay special attention to the first ‘turn’ of the conversation to earn the caller’s trust. Consider what voice characteristics would make patients feel comfortable and trust the system. Do you want a warm, soothing tone for a hospital context? Should it sound very professional or more conversational?



Empathy is essential

Word choice, intonation, volume, and the use of silence work together to relay meaning. Expressions of empathy are one of the hardest things to get right in conversational design.



Reduce cognitive load on the caller

Overly complex instructions make a conversation hard to follow. Since patients can’t see information, the AI agent shouldn’t present too much at once or ask the caller to remember things from earlier in the call. Good voice design anticipates various ways a patient might respond and is prepared to handle interruptions, questions, or confusion gracefully.



Plan for seamless handoffs to humans

Some calls need a human agent. Set clear criteria for when to transfer like—urgency, emotional distress, or patient frustration. Make sure your AI agent can pass context to a human agent and explain the handoff to the caller to keep the experience smooth and reassuring.



Use clear, jargon-free language

Patients come from all backgrounds, and not everyone speaks “healthcare.” Use everyday language, explain any necessary medical terms, and avoid acronyms or system-specific phrases. Clarity builds trust and helps patients feel confident using an AI agent.



Crafting seamless handoff protocols

Sometimes calls will need to be passed from the AI agent back to your team.

While technically, this is a simple SIP transfer, you'll want to consider the handoff protocol to ensure maximum efficiency, without diminishing patient experience.

When designing handoff protocols, you'll want to consider:



Urgency

Does the caller have an emergency that requires prioritization?



Sensitivity

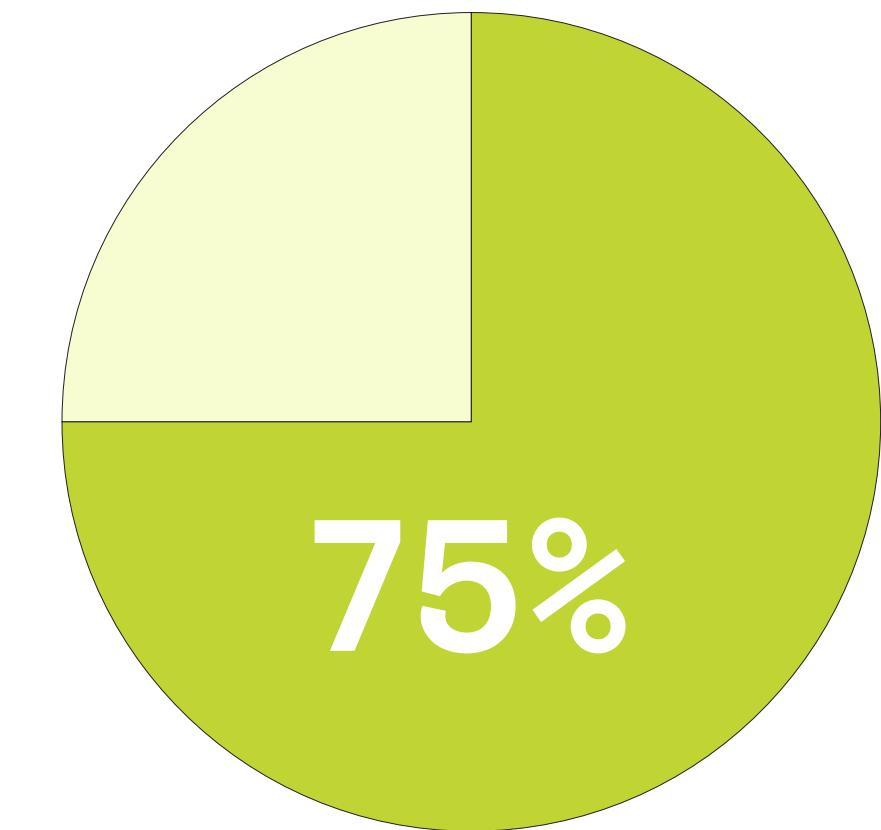
Is the caller in a vulnerable state or situation that would benefit from the human touch?



Patient retention risk

How likely is the caller to hang up and take their business elsewhere?

Timely handoff is key to trust and efficiency



of patients expect to speak with a human clinician quickly after an AI-led intake.

Your voice design checklist	
Pick a voice	Review and select a voice that fits your health system's brand and patient needs. The voice should be clear, compassionate, and representative of the tone you want for patient interactions. <input type="checkbox"/>
Design for automation	Create engaging, user-friendly conversations. Emphasize a warm greeting and empathetic tone. Keep prompts and questions simple, and use natural, human-like dialogue. Adjust the AI agent's tone based on context (cheerful vs. calming as needed) and minimize the information burden on callers by breaking down tasks. <input type="checkbox"/>
Create a robust handoff protocol	Define the rules for transferring callers to a human agent or clinician. Consider factors like medical urgency, emotional distress, and potential patient frustration or attrition if their needs aren't met quickly. Ensure the system can pass along the call context to the human and that the handoff process is clearly explained to the patient to maintain trust. <input type="checkbox"/>



03.

Build

An engaging conversational AI solution will require a specialized tech stack including speech recognition, spoken language understanding (SLU), natural language understanding (NLU), voice synthesis, dialogue management and AI guardrails.

To understand the process of building an AI agent, it's helpful to consider conversation as three constituent parts.



1. Listening
Understanding what callers are saying, however they speak.



2. Reasoning
Guiding callers through business logic to complete transactions.



3. Speaking
Giving appropriate responses in an engaging tone of voice.

PART 1:

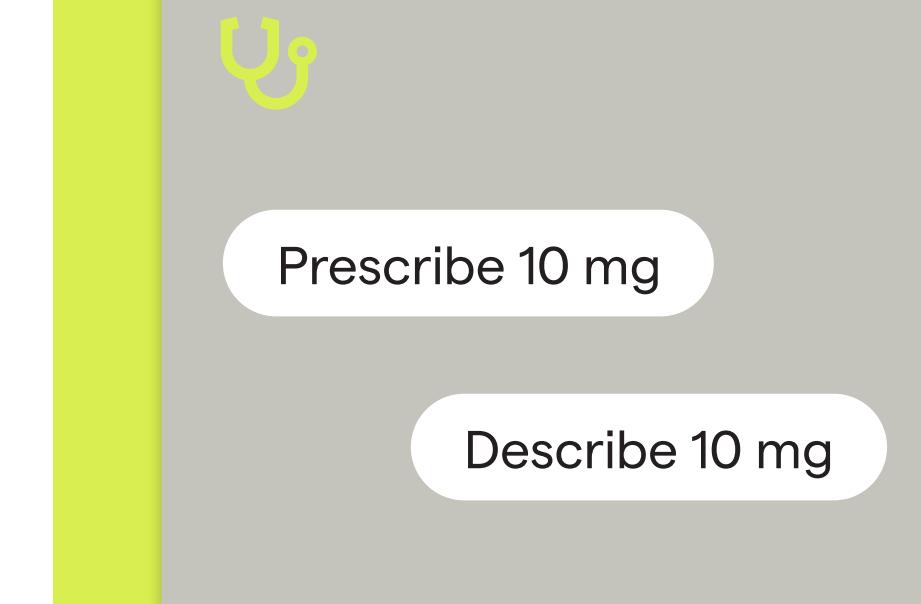
Listening

During any conversation, the ability to listen is impacted by multiple external factors. Is the environment loud? Does the person appear to be listening to you? Are there interruptions? Patients might use colloquial or medical terms that are tricky to recognize.

These factors similarly make automating human-machine conversations more challenging and can lead to speech recognition errors that make it difficult for AI to capture spoken language accurately.

A great listening stack includes:

- **Automatic speech recognition (ASR)** - These systems transcribe spoken language into text that can be digested by LLMs. They are often fine-tuned for specific accents, languages, and use cases, so you may need to use several models concurrently for accurate understanding.
- **Spoken Language Understanding (SLU) models** - Even the best ASR will result in gaps and errors in transcriptions. SLU models recover important information from incorrect speech transcriptions, using context and customer information to infer the correct input



Identification & Verification (IDNV)

It is an indus
Hernández

Contextual Recognition

Can the appointment be at hate o'clock?
eight

Alphanumeric Parsing

Apple Charlie Echo double eight zero for seven
A C E 8 8 0 4 7

Ensemble of Recognizers and N-best List

I need sun cream to help me with my eczema.
I need some cream to help me with my eczema.
I need some cream to help me with my eggs ma.

Timing is everything: Balancing latency and interruptions.

One of the most crucial elements in building an AI agent is smooth, timely interactions.

Delays between user input and system response can frustrate users if the agent takes too long to respond. But interruptions to fast or incomplete responses present a problem to harmonious human-machine communication.

There needs to be a balance where the system responds quickly and accurately without cutting off or confusing the user. Unfortunately, there's no hard and fast rule on when you should allow patients to interrupt. It depends completely on what you deem important to your business and what your patient deems important.



Should you allow patients to interrupt?

Here are some useful questions to ask when considering if and when patients should be able to interrupt an AI agent:

- How important is it that the patient listens to everything? Do you want to ensure the patient does not interrupt the AI reading out terms and conditions?
- What is the potential impact of a patient interrupting? If you're reading out a list, and they skip later options, is this likely to result in an error further along the conversation?
- What is the impact on patient experience? Will allowing or preventing the patient from interrupting have a negative or positive impact on their experience?
- How do you want to balance patient experience and business processes? If patients are likely to skip the boring stuff, how do you strike a balance between your priorities and theirs?

PART 2:

Reasoning

Once the speaker's words have been transcribed, the AI agent needs to understand the context behind what the caller is saying and how to respond in a way that continues to move the conversation toward an appropriate resolution.

A great reasoning tech stack will include:

- **Large Language Models (LLMs)** - These machine learning models can extract meaning from words and sentences and define the next steps the system should make in the context of the conversations.
General-purpose LLMs aren't optimized for structured, goal-directed dialogue. Owning and fine-tuning a model gives you more control over latency, behavior, and deployment flexibility, allowing you to optimize every deployment for a given use case.
- **Safety guardrails** - In healthcare, trust is everything. AI agents can handle a wide range of topics, but without the right safeguards, they risk saying the wrong thing to vulnerable patients. What's appropriate in a retail context may be deeply harmful in a clinical one.
You should be able to configure safety filters to block harmful language—both from patients and the AI agent—in a way that aligns with clinical standards and brand values. The use of retrieval-augmented generation (RAG) keeps responses grounded in your approved medical knowledge, helping prevent hallucinated or misleading answers.
These safety guardrails help ensure your AI agent stays compliant, respectful, and aligned with the sensitive nature of patient care.



PART 3:

Speaking

Once the AI agent has listened and understood the caller's intent and the appropriate response, it must then turn that response into speech.

Even with the best technology, a robotic, unnatural voice provides a subpar experience and discourages callers from engaging, eliminating the benefits of voice automation.

A great speaking tech stack will combine voice cloning technology with state-of-the-art synthesis and to create an experience that sounds like talking to a real person. The AI agent's spoken responses should be as engaging and appropriate as a well-trained human agent. This ensures patients stay comfortable interacting with the AI, which drives higher usage and containment.



Integrations

Once your AI agent is built, it needs to be integrated into your existing tech stack to enable calls to pass back and forth to your contact center.

Call center integration - SIP

A simple SIP or PSTN connection is all that's required to route calls between your voice assistant and your team. This is virtually the same for every AI agent and your IT team will handle this with support from your conversational AI vendor.

A typical high-level approach is shown in the [image on page 26](#):

1. The call enters your existing infrastructure from the PSTN / carrier
2. Calls are forwarded to your solution via SIP INVITE. NOTE: PSTN forwarding is an alternative here. However, SIP is preferable. Using a combination of AI techniques, your solution handles the call. Where applicable, the process will use your APIs and services to maximize automated resolution.
3. Where it isn't possible to resolve the patient call, it is returned to the patient infrastructure, via SIP REFER, SIP INVITE, or PSTN.

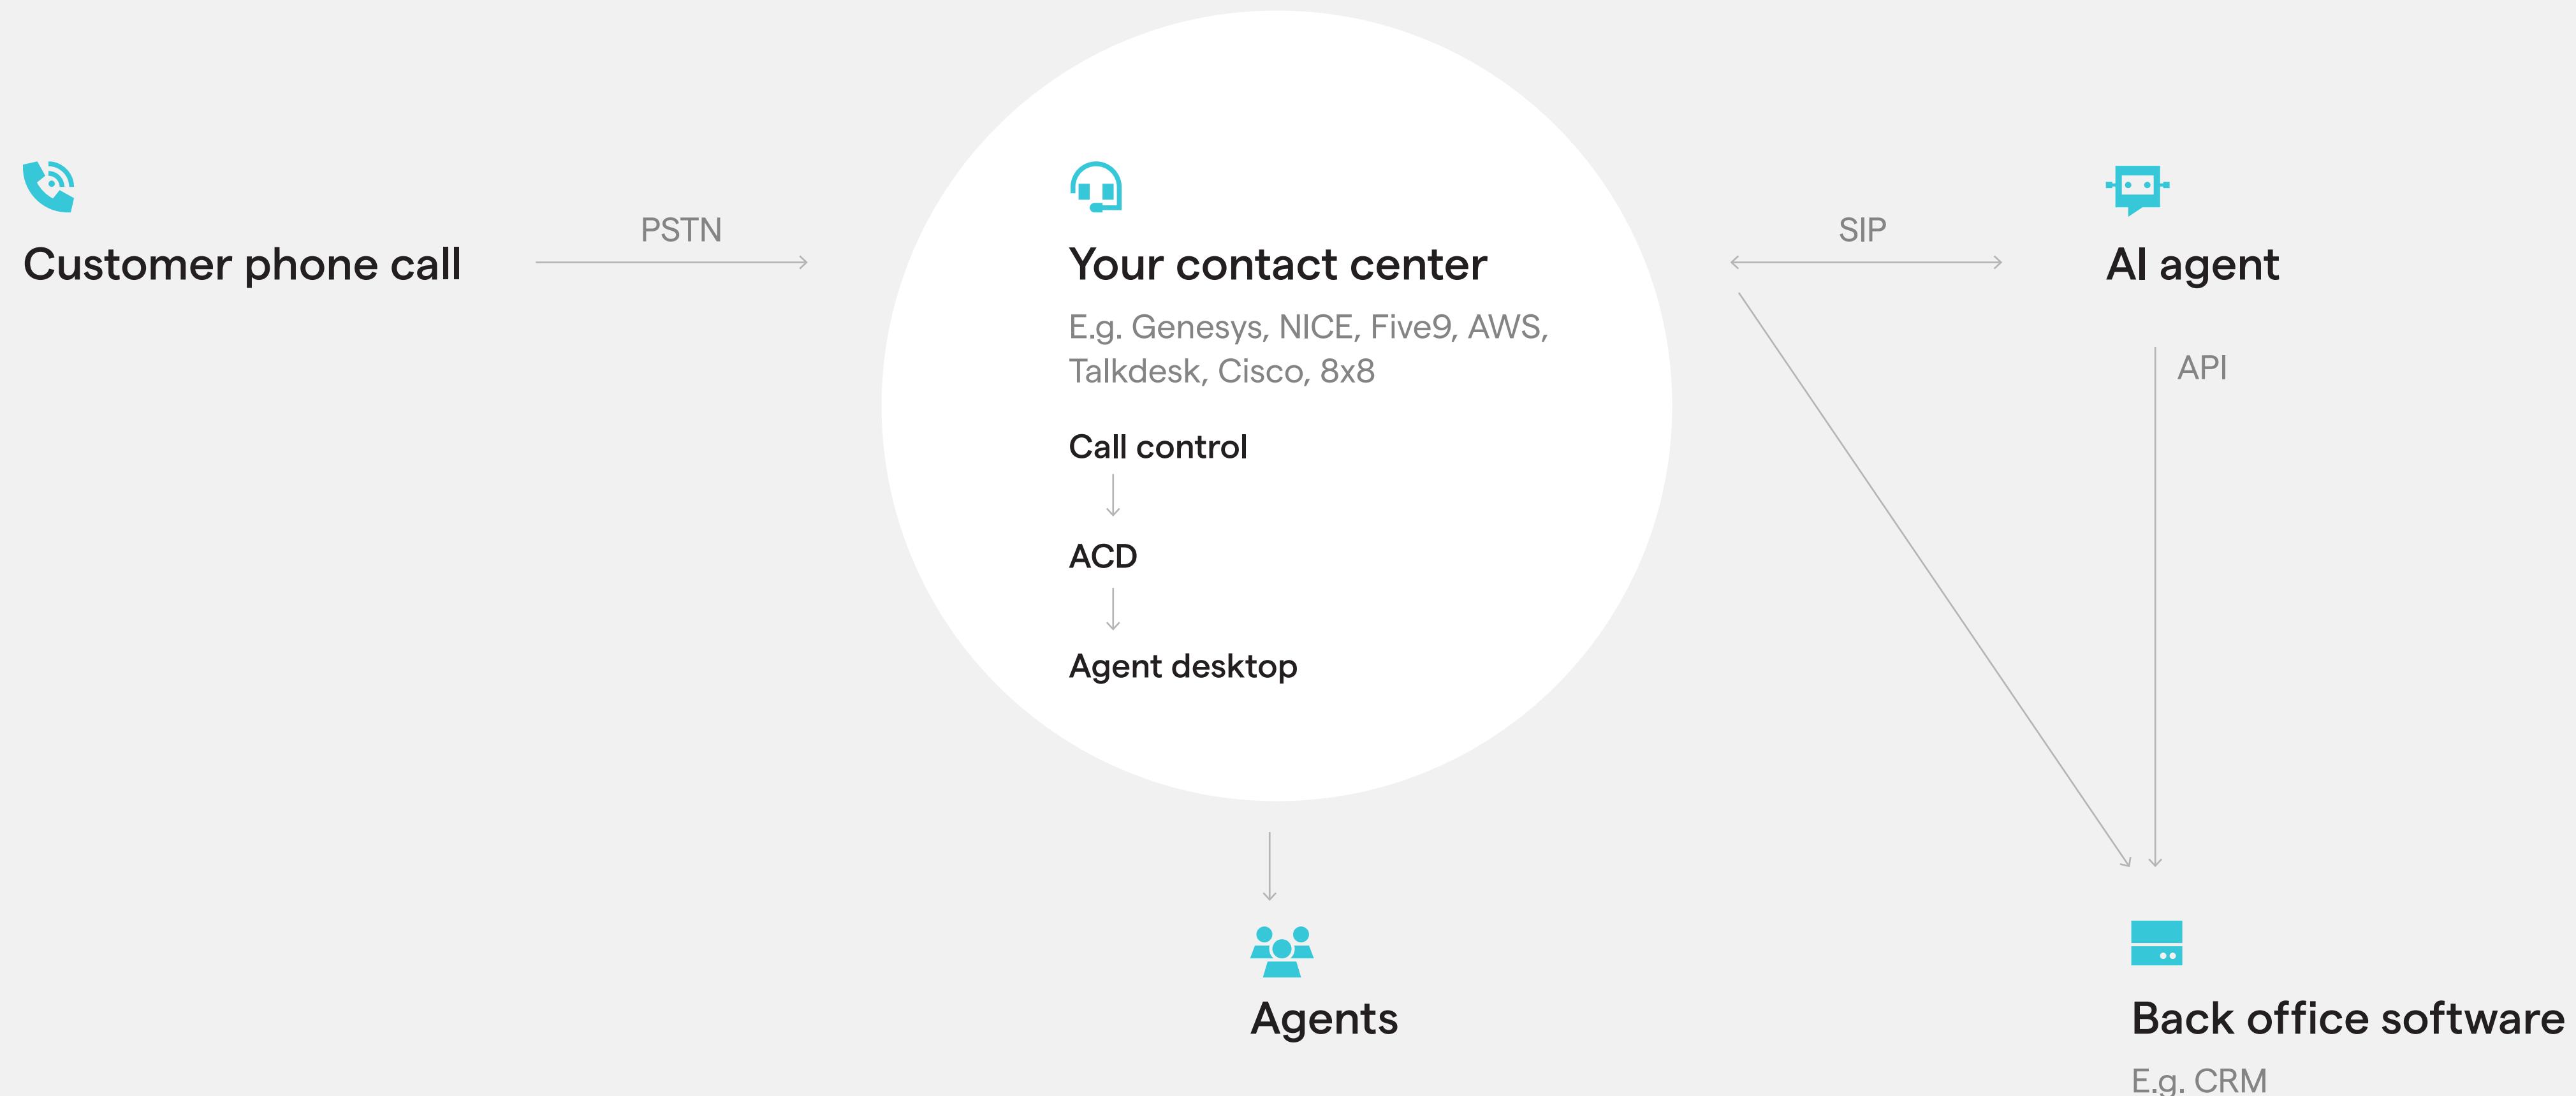
4. The call is routed to the relevant agent group/queue based on the information gathered on the automated portion of the call. A screen pop with a summary is sent to the agent (e.g., "Caller wants to reschedule cardiology appointment; chest pain flagged – transferred to triage nurse") to ensure a smooth handoff.

Remember, as you build:

- Test components individually and together.
- Use real patient call examples to check speech recognition.
- Verify APIs with test data to ensure your AI agent can pull up and report information correctly.
- Run full call flows yourself to confirm your AI agent handles calls and routes to humans when needed.



Here's a high-level example of what your system architecture might look like.



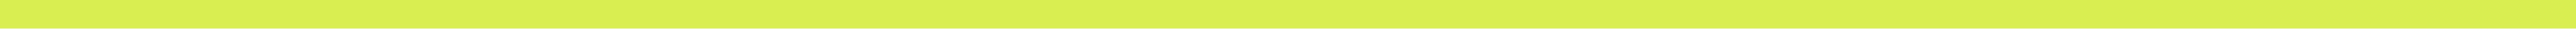


Your build checklist		<input checked="" type="checkbox"/>
Train machine learning models	Fine-tune for an accurate understanding of specific accents, languages, and use cases.	<input type="checkbox"/>
Develop your tech stack for engaging conversational AI	Build systems for speech recognition, spoken language understanding (SLU), natural language understanding (NLU), voice synthesis, dialogue management, and AI guardrails.	<input type="checkbox"/>
Integrate your AI agent into your existing tech stack	Integrate the AI agent into your tech stack via a SIP or PSTN connection to enable seamless call routing between the AI agent and your contact center, with support from your IT team and conversational AI vendor.	<input type="checkbox"/>
Integrate the screen pop function	Ensure the AI agent provides the agent a screen pop with relevant caller information during handoff.	<input type="checkbox"/>
Ensure compliance and security	Follow healthcare security best practices. Use encrypted connections (e.g., HTTPS, TLS) for any data transit. Verify that PII is handled according to HIPAA.	<input type="checkbox"/>





04.



Test and launch

Now your AI agent is almost ready to deploy.
But first, testing!

Testing

You'll want to run a number of rigorous tests before going live with your patients.



1. Quality assurance - To test various user interactions, commands, and responses to ensure your AI agent delivers the expected experience and identifies any issues. Like different ways to ask for appointments or prescription refills, and check that the agent responds appropriately.



2. Load testing - Simulate high call volumes to ensure the conversational AI platform and your telephony integration can handle peak loads, such as Monday morning rush, or a vaccine clinic announcement causing a spike.



3. Team demos - Have your team call in and try the system for themselves. Your agents are great for this, as they know what patients are most likely to say! When staff feel included in testing, they become advocates for the new system rather than skeptics.



4. Workflow integration testing - Test full end-to-end scenarios to make sure every system connection works as expected. See [example on page 30](#). These tests help catch issues like missing data or slow transfers. In healthcare, where calls often involve multiple steps, it's critical that nothing gets lost along the way.

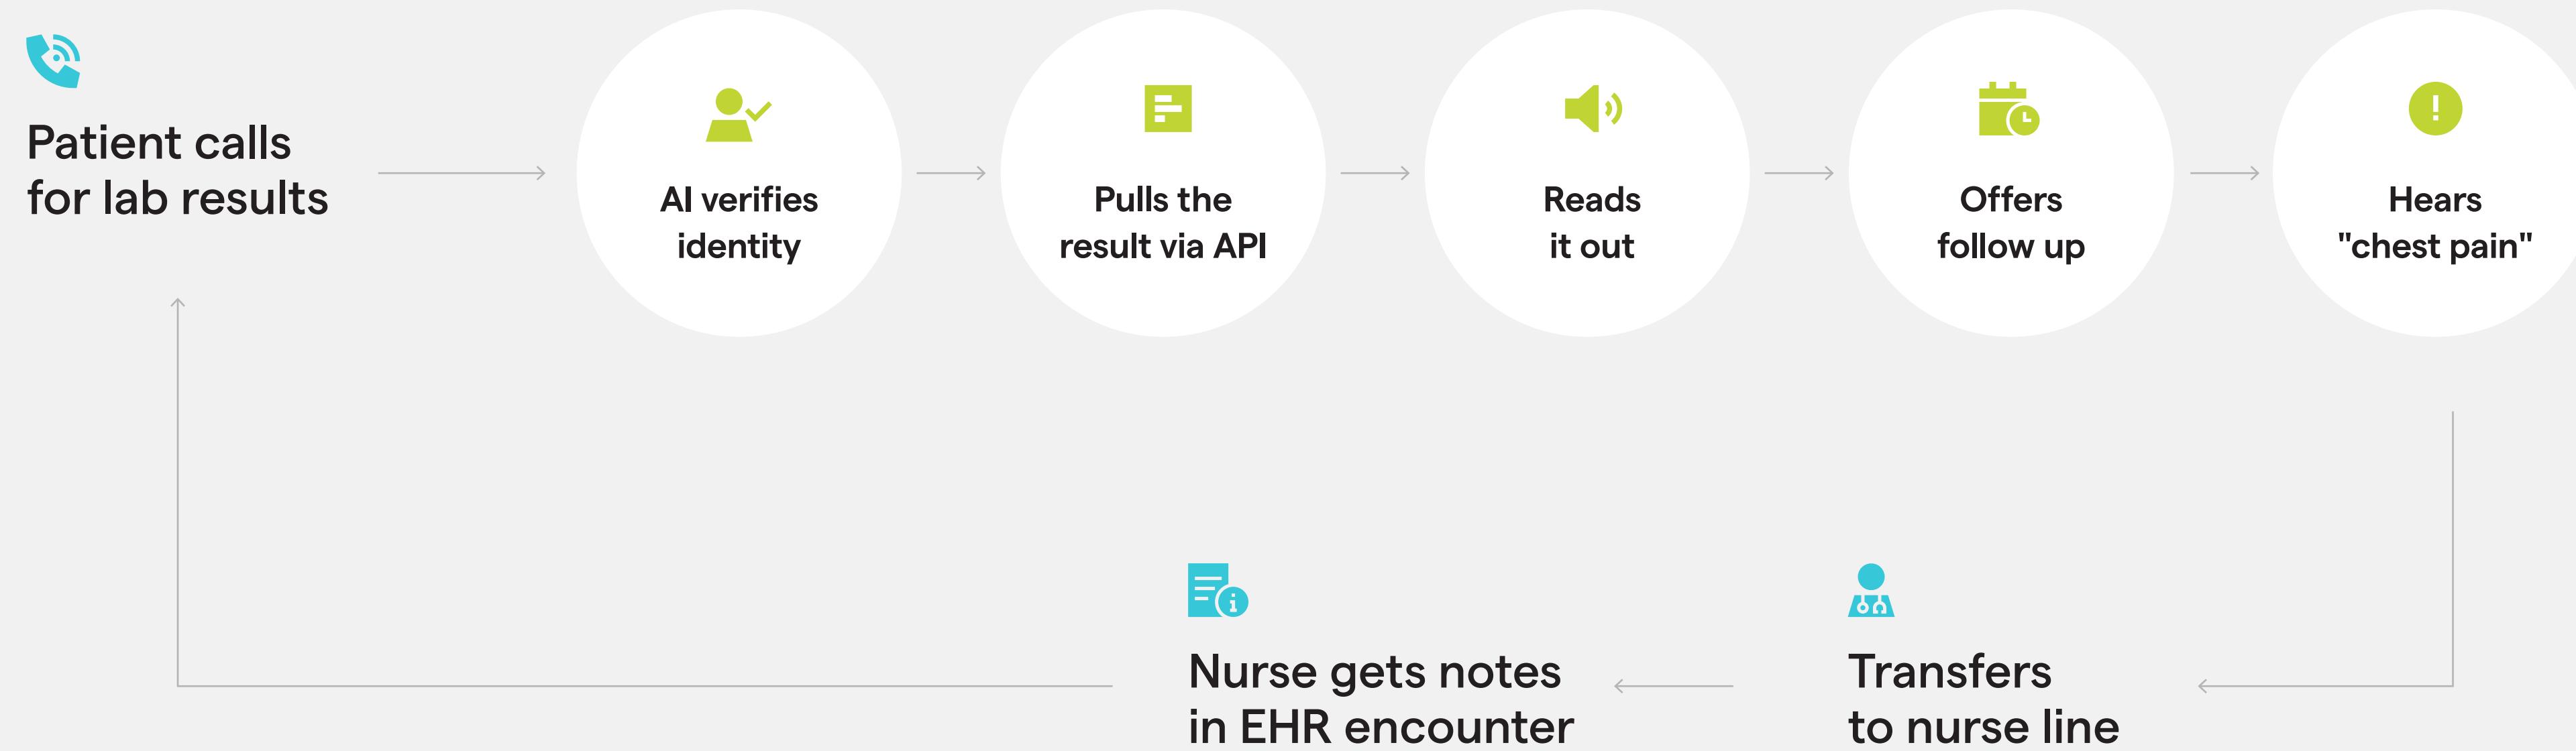
Launch

Some healthcare organizations opt to go live with a percentage of calls before rolling out the AI Agent to everyone, such as after-hours calls, one department, or 10% of mainline traffic. If you go this route, plan your ramp-up and keep staff informed so they're ready for any AI handoffs.

Some take the dive to reap the rewards faster. In healthcare, a full launch is typically done when the use case is low-risk and the organization is eager to address an issue immediately.

However you approach your launch plan, your vendor should have a support team on hand to monitor early calls, make sure everything is running smoothly, and jump in to make amendments if needed. If you encounter unsuccessful calls due to unexpected patient inputs, use them as an opportunity to work with your vendor to develop new use cases that may not already be supported.

Workflow integration testing example



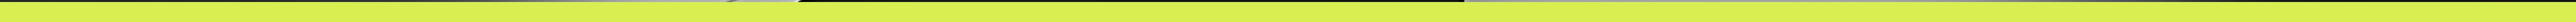


Your test and launch checklist		<input checked="" type="checkbox"/>
Quality assurance	Test various user interactions, commands, and responses.	<input type="checkbox"/>
Team testing	Have your team, especially agents, test the system to simulate patient interactions. Ensure the AI agent can handle high volumes of calls.	<input type="checkbox"/>
Optimizing accuracy and efficiency	Monitor the performance in a live environment, identifying areas for optimization and fine-tuning to enhance accuracy and efficiency.	<input type="checkbox"/>
Launch	Decide whether to go live with a percentage of calls or full deployment.	<input type="checkbox"/>





05.



Post-launch

In the weeks following the launch, you should closely monitor performance and make adjustments to improve based on your business goals.

Continuous improvement

Every patient call is different, and it isn't until you've launched your AI agent that you'll get to see how callers react.

By closely monitoring early calls and making adjustments where needed, you'll improve patient experience rapidly.

Key metrics such as AHT, call volume, and containment should be reviewed regularly with stakeholders, and tweaks and enhancements can be made to the ASR, machine learning, and dialogue design to ensure the best results for your patients.



Making updates

If there's one thing to expect in patient service, it's the unexpected! As your AI agent beds in, you'll want to add new functionality, edit prompts and make further tweaks to improve experience and efficiency.



Further automation

Once your AI agent gets to work on your chosen use case, it will instantly begin gathering data that gives you a clear indication of your most popular calls and helps determine the next-best call types for further automation. This data will be extremely useful when it comes to encouraging buy-in for further investment.



Listen to calls

Listen to recordings or read transcripts (with privacy in mind). You'll hear how patients actually use the system; what they say, where they get stuck, and what works well. It's one of the best ways to uncover unexpected behaviors and improve the experience.



Managing your knowledge base across multiple product lines, departments, or compliance requirements can get complicated quickly.

Look for a platform with version control that makes tracking changes and maintaining accuracy easier. This will allow your teams to collaborate on updates and ensure all changes are clear and well-documented.

You'll have a well-structured knowledge base that ensures your AI agent retrieves accurate information, performs tasks reliably, and delivers a consistent user experience across multiple environments.





Conversation review

AI agents are powering more flexible and natural automated conversations than ever, but off-the-shelf LLMs don't always offer the control enterprises need. When you're choosing a conversational AI vendor, make sure you can review and manage what the system is doing in real time, including the ability to:

- Access transcripts and listen to call recordings
- See what topics the AI agent is referencing during the conversation
- Review function calls to understand how patient actions trigger specific responses

This level of visibility makes it easier to troubleshoot, improve, and stay compliant, so you can make fast updates with confidence and keep improving the patient experience without delays.



Knowledge base updates

Your AI agent should be as informed as your best-trained staff member. When patients get incorrect information, it often comes down to staff working with outdated or incomplete guidance. The same goes for AI agents. If your AI doesn't have access to the latest clinical protocols, location-specific policies, or scheduling rules, it can't provide safe or accurate answers.

Your vendor should make it easy to keep your AI agent current through call reviews or the ability to jump directly to the relevant knowledge base entries and update information in real time.

The result is no more outdated instructions or missed context—just consistent, reliable support that reflects your most up-to-date care standards.



Scaling and further optimization

Successful scaling starts with low-risk, high-impact use cases that can quickly demonstrate value. Use these wins to develop standardized playbooks that can be applied to new use cases. Document everything: what worked, what didn't, and what needs to be adjusted for broader implementation.

Establish strong feedback loops, not just from analytics dashboards, but from your agents, patients, and teams. These insights become essential for refining and expanding your strategy.



Your post-launch checklist

Your post-launch checklist	
Regular monitoring	Monitor performance, analyze caller interactions, and adjust the AI agent accordingly. <input type="checkbox"/>
Make updates	Edit prompts and make tweaks to enhance patient experience and efficiency. <input type="checkbox"/>
Review metrics	Monitor AHT, call volume, and containment regularly and regularly share these metrics with stakeholders to show impact and catch any negative trends. <input type="checkbox"/>
Review conversations	Access transcripts and listen to call recordings, see what topics the AI agent is referencing during the conversation, review function calls to understand how patient actions trigger specific responses. <input type="checkbox"/>
Identify further automation opportunities	Use the data to pinpoint popular call types or pain points that could be improved with further automation. Engage stakeholders to prioritize the next phase. Keep a roadmap of conversational AI enhancements and align it with your broader patient experience strategy. <input type="checkbox"/>



CASE STUDY

Howard Brown Health provides personalized patient experiences with PolyAI

Howard Brown Health partnered with PolyAI to build an AI agent that integrated with their existing RingCentral (later Dialpad) telephony system and Epic EHR. Without overhauling infrastructure, they delivered fast, multilingual support to their diverse patient base. Here's how it works:

1. Every call is forwarded securely via SIP to PolyAI. The voice assistant ("Alex") greets callers immediately in English, Spanish, or Polish.

2. Patients request services like appointment scheduling or prescription refills. The AI authenticates callers and uses backend integrations (e.g., Epic MyChart, pharmacy systems) to fulfill requests.
3. Sensitive issues or emergencies are escalated instantly to live agents or nurses, with a screen pop summarizing the conversation so far to avoid repetition. Alex also recognizes distress cues and prioritizes appropriate routing for patients in crisis.

30%

of calls fully contained by the voice assistant

72%

reduction in handle time for agent-assisted calls

4%

increase in patient satisfaction scores

0

hold times for automated calls - every patient is greeted instantly

24/7

coverage, including during public health emergencies and call surges



Conclusion

Implementing conversational AI in healthcare involves a journey from initial assessment and team-building through careful design and development to thorough testing, launch, and ongoing optimization.

By embracing conversational AI, healthcare providers can modernize their patient experience. The journey requires collaboration across departments and a clear vision, but the end result is a responsive, efficient contact center that enhances patient care and organizational performance.

In a world of rising patient expectations and ever-changing healthcare demands, conversational AI is fast becoming a vital part of the solution. It enables providers to deliver empathetic, timely service at scale. With the right implementation roadmap, your organization can be at the forefront of this positive change.

[Learn more](#)

