



Conversational AI Agent Development Process

Airlock AI specializes in developing Conversational AI Agents to help accelerate sales, create and drive leads, and guide customer support. When combined with expert automation services, AI can transform your business with digital intelligence on autopilot, saving you time, minimizing mistakes and increasing efficiency.

You can anticipate a one-time development fee, followed by a monthly recurring charge for monitoring and optimization. The final pricing will hinge on the specific scope of your project. Factors influencing this include the complexity of your AI, how many resources it needs to do the job you want, the number of integrations needed (such as CRM linkage) and the frequency of AI knowledge updates.

Development fees range from \$997-\$9997 or higher, with monthly recurring charges of \$197-\$997 or higher.

That's a broad range! Most businesses can expect development of \$1997-\$3997 and monthly charges of \$297-\$497.

- After an initial conversation with you **we'll build a prototype for free**. We want you to be able to envision outcomes.
- You're in charge of scope and you'll know what your investments will be before we get busy building your AI. No surprises.

Keep reading for how the entire process works, including what impacts pricing.

Conversational AI Agent development consists of several stages:

- Pre-Production
 - Prototype
- Production
 - Strategy and Scope
 - Design
 - Training
 - Development
 - Testing & Quality Assurance
 - Iterate & Adjust
 - Deployment & Integration
- Post-Production
 - Monitoring and optimization

Pre-Production

Prototype

Where we will, after an initial conversation with you, quickly create a visual flow and functional (though limited) AI to demonstrate basic capability, and allow us to talk through how things will work. The prototype is designed to let you make an informed decision about moving forward.

Required from you:

- An initial conversation of around an hour

Deliverable: Functional AI Prototype

Timeline: 1-2 days

Investment: No Cost

Production

Strategy and scope

Where we identify what we want your agent to do, when, why, and what are ideal outcomes, like schedule a conversation, visit a website page, share some information, etc.

- Our guided process makes this easy, helping us quickly move from “*what if we...*” to “*it will then...*”
- Often accomplishing the strategy involves integrating with additional tools, applications or platforms. Sending leads from the AI Agent to a calendar, or your CRM, for instance.
- We'll define how broad your AI's understanding and capabilities should be. This will influence how much training it needs on your business. A few website pages and PDFS is often all it needs. For more granular, specific abilities, AI will need detailed, numerous resources.
- **The broader the strategy and scope, the higher the investment**
- Where appropriate, we can suggest a phased approach to development.

Required from you

- Collaboration in generally one meeting of around an hour
- Email collaboration as we share development
- Sign-off

Deliverable: Scoping Document

Timeline: 3-5 days

Design

Where we build the logic and conversation flow of *if the user does this, the AI should do that.*

- With strategy and scope understood, we can quickly map out a non-functional or minimally-functional logic flow that, together, we can visualize and understand.
- **The more complexity required of the logic flow, the higher the investment**
- **At the end of this flow mapping, scope is set**
- **50% of initial invoice paid begin development**

Required from you

- Collaboration in generally one meeting of around an hour
- Email collaboration as we share development
- Sign-off
- 50% invoice payment

Deliverable: Visual Logic Flow

Timeline: 1-2 days

Training

Where we “teach” the AI your materials, brand, conversational tone and more.

- The AI is only as good as what it knows. We begin by collecting all the necessary information that will populate the AI's *Knowledge Base*. This includes FAQs, product details, pricing, company information, and any other relevant content that will help us satisfy the strategy. We train the AI with the collected data and knowledge base to ensure it understands user queries and provides accurate responses.
- Beyond informational data, we'll also train on any personality characteristics you have in mind. This can be brand guidelines, tone of communication, personal style or preference and more.
- Our easy submission process lets you simply upload what's necessary as a result of our Strategic decisions together.
- **The more resources required in the AI's Knowledge Base, the higher the investment.**

Required from you

- Resources to train your AI, like website urls, product or services pdfs, brand guides, etc.
- Which resources we need will be established according to Strategy and Design steps

Deliverable: Internal Only

Timeline: 1-2 days (depending on your ability to submit resources)

Development

Where we combine *Design* and *Training* for a quality user experience with the AI behaving as expected.

- AI technology is implemented, integrated, and configured to align with your specific requirements as identified in Strategy and mapped out in Design. This includes setting

up natural language processing (NLP), machine learning models, and other AI components.

- We'll also establish a cadence for how often knowledge for the AI is refreshed. For instance, if new blog posts to your site each week are relevant for the AI to know, or if monthly changes to your website are anticipated, etc.
- **The more complexity required of development, such as integrations or a large knowledge base, the higher the investment**

Required from you

- Login or API credentials for any integrations
- Patience and support as we build your Conversational AI
- Up to 2 check-ins for feedback as necessary
- Sign-off

Deliverable: Internal Only

Timeline: 1-2 weeks (sometimes longer for very large projects)

Testing and Quality Assurance

Where together we identify and rectify any knowledge or functional issues.

- Any outside integrations (for instance, your CRM) will be tested at this time too.

Required from you

- Patience and support as we build your Conversational AI
- A meeting of around 30 minutes as we share your Conversational AI
- Resources like time and people to test your Conversational AI

Deliverable: Conversational AI ready for testing

Timeline: 3-5 days (sometimes longer if integrations are complex)

Iterate and Adjust (Up to two rounds)

Based on our combined experiences and feedback, we iterate the Conversational AI for improved performance, up to two rounds.

- Any required additions to the Knowledge Base to accommodate improved performance are implemented

- Iterations are performed in-scope. If any feedback requires new scope, we'll be happy to discuss.

Required from you

- Patience and support as we iterate your Conversational AI
- Email collaboration as we share development
- Sign-off

Deliverable: Conversational AI ready for deployment

Timeline: 3-5 days (sometimes longer if iteration is complex or scope shifts)

Deployment and Integration

Where we prepare the AI to live on your website or other digital destination

- Typically the AI Agent will live on your website, especially if it serves to qualify leads or cultivate customers to sales, or if customer support is an imperative. If you have bespoke needs, like on an intranet or inside an app, we'll establish this understanding early in the Strategy discussion.
- Launch and monitoring, where we closely monitor its performance during initial stages.

Required from you

- Admin credentials to your website
- Final 50% of invoice paid

Deliverable: Conversational AI live on your site

Timeline: 1 day

Overall Timeline: 2-4 weeks (sometimes longer, depending on complexity)

Overall Investment to go from ideation to installation:

- Starting at \$997, up to \$9997 or higher
- Complexity, as indicated in all earlier steps, impacts pricing
- **Most businesses are between \$1997 and \$3997.** We'll establish this together early in the process. No surprises.

Post-Production

Monitoring and Optimization

Where the AI is continually fine-tuned to learn from real-world experiences, new data, technology or language model updates and more. With our ongoing support, your AI remains efficient, effective, and future-ready.

- Conversation transcripts are reviewed each day, each week or each month to identify new opportunities for training
- New training is conducted on the AI
- Analytics that suggest an opportunity outside of AI is shared accordingly with you. For instance, “We see that a lot of people are asking if this comes in red - maybe an opportunity there for you.”
- **The more frequent the review cadence and more complex the AI, the higher the monthly investment**

Timeline: Daily, Weekly or Monthly

Monthly Investment:

- Starting at \$197, up to \$997 or higher
- Complexity and cadence, as indicated, impact pricing
- **Most businesses are between \$297 and \$497/month.** We'll establish this together early in the process. No surprises.
- No long-term commitments. You can stop anytime with 60 days notice.
- Preferential pricing for semi-annual or annual payments.

Questions? Answers.

How does the Conversational AI personalize its interactions with different visitors? Does it adapt based on visitor behavior, location, or other factors?

Every interaction is unique. If Tony is engaging with your Agent on your website at the exact same time as Steve is, they're both having unique, original experiences. Personalization can be straight forward. If the AI obtains a name, for instance, it'll remember and use it throughout the chat. Personalization can also be achieved through thoughtful questions where the AI picks up on individual needs and issues. While the AI is trained to "speak for itself," we can be very specific on what it says and when. *At this point, ask a question about this thing. Then share an insight. Then ask another question.* And so on. Throughout this process, the AI is adept at taking the conversation's essence, mulling it over with what it's learned from its training materials, and then delivering a fitting response. This might be a direct answer, a helpful link, an invite to drop some contact details, or other actions that drive results.

How does the Conversational AI Agent integrate with existing sales and CRM systems, or HR systems? Is there a seamless handoff if a human needs to intervene?

Your AI can seamlessly integrate with a broad array of processes, including leading CRM and sales systems. Beyond that, stakeholders can be alerted of new conversational activities through email or text, and we can even append these details to a Google doc for effortless export.

As conversations unfold in real-time, our AI is skilled at discerning user intent and key events. We can then tag users based on these insights. This enables your reports to highlight elements the AI identifies as urgent needs, specific interests, potential timelines, or any other pivotal aspects integral to your sales journey.

Data & Privacy: How do you ensure data privacy and GDPR compliance? Does the AI store conversations or personal details of the visitors?

At present, our services are offered exclusively within the United States. The platforms and systems we utilize, including but not limited to ChatGPT, Claude, Amazon Web Services (AWS), and Google Cloud Platform (GCP), each have distinct privacy and security policies. We ensure these policies are made available to you before formalizing any engagement.

It's important to highlight that:



- Conversations facilitated by our system are not stored or repurposed for training beyond your specific use case and conversational system. We do not directly use any data for further training of any LLM models in general.
- Both encryption-at-rest and encryption-in-transit are rigorously implemented.

However, it's crucial to understand that we cannot be held liable for potential data or security breaches that might occur due to vulnerabilities in our software vendors' systems. Additionally, any communications produced by the AI, based on the materials you've provided, are not under our direct liability.

Can the AI operate in multiple languages?

Presently our capabilities are limited to US English.

Culture & Tone Adherence: How is the AI programmed to truly reflect a company's culture and communication style? What's the process to fine-tune this?

Just like with our knowledge base, any materials you provide on company culture will help train your AI. Any materials you possess that shed light on your organizational ethos can be used to fine-tune your AI. At its core, the AI is trained to resonate with your target audience. This could involve understanding their demographics, professions, or any other nuances central to your sales engagements. Guidance on tone like "casual yet professional" is taken to heart, and if you have specific jargon, acronyms, or unique brand expressions, those are integrated seamlessly. If you have a brand style or brand voice guide, we'll implement it. Additionally, we offer a dedicated 'Brand Style' service to fine-tune this communicative approach further.

Should the AI face questions outside its primary function, it's trained to not make guesses or venture beyond its knowledge base. It's trained to not speculate, opine or engage. Its training is for responses to stay strictly within the boundaries of its training. Hallucination is a term to describe AI's tendency to sometimes make up information or stray from its training. We have safeguards in place to, if not eliminate this phenomenon, minimize it significantly.

Before your conversational system is introduced, it undergoes testing, with your team playing a crucial role, to ensure it communicates in a manner consistent with your brand's standards.

Feedback Loop: If the AI doesn't know an answer to a question, can it be trained to know it next time?



Yes. When the AI encounters a question it can't answer, our system is primed for assessment and improvement. This often involves refining the AI's prompt configurations, supplementing its training materials, or reimagining parts of its conversational strategy. This is part of the crucial value of ongoing Monitoring and Optimization.

Customization & Branding: How customizable is the AI's interface? Can we adjust its appearance to match our branding?

At this time, icons and colors can be customized. Structural aspects like dimensions or placement on the page aren't customizable - for now, the AI conversational system will "live" inside a button on the lower right of a website page.

Continuous Improvement: How often is the AI updated or improved? Do you provide regular updates based on new findings or technologies?

Generative AI platforms, including ChatGPT, and the associated systems we utilize undergo regular updates, although the timeline is beyond our direct control. Rest assured, when significant updates arise that might influence the design or efficiency of your AI conversational system, we'll promptly inform you and offer relevant recommendations or guidance.

Insights Capture: What kind of insights does the AI capture from the conversations? Can these insights be exported or integrated into other platforms for further analysis?

The AI is designed to seamlessly capture essential data that drives your business conversations forward. Think of it as an interactive form: instead of static fields waiting to be filled, our system engages users in a dynamic conversation to gather the same data. This ensures a more user-friendly and engaging experience for your visitors. Moreover, the information collected is easily integrated into other platforms you may be using, ensuring continuous and efficient data flow for your business processes.

Lead Prioritization: How exactly does the AI evaluate and qualify leads? Can we set specific criteria for what makes a lead "high-quality" for our business?

Lead prioritization is entirely tailored to your specific needs. We collaborate closely with you to understand your existing sales practices and criteria. By grasping the nuances of your lead qualification, and by recognizing the vital questions that shape your sales scenarios, our AI ensures that every lead is ranked and managed in line with your unique business objectives.



Training & Setup: How much involvement is required from our end during the initial setup and training phase? Will we need to provide specific scripts or conversation flows, or is it more hands-off?

We have a simple process to retrieve essential knowledge assets from you when we start building your Conversational AI Agent, which are the materials we'll use to train and teach your AI about who you are and what you do. Beyond that, you'll have checkpoints to review and test your AI and make suggestions and refinements. But, you're never responsible for holding our hands or micro-managing. The entire development process is fun, interesting and easy.

ROI Measurement: Do you provide any analytics or reporting tools to help track conversions, lead quality, and other relevant metrics?

Analytics we capture include:

- Interactions: Total number of engagements users have had with your AI.
- Recognition rate: The percentage of messages understood by your AI, i.e. did not trigger the fallback intent.
- Unique Users: Unique users that have engaged with your AI.
- Unique Sessions: Unique user sessions with your AI.
- Top Intents: The most popular queries users ask your AI.

You can also compare different time periods to see how your AI is performing with each iteration of your assistant. You can then cross reference these analytics with your internal ROI measurement tools to determine how it's working. We also regularly use this data to improve your AI.

What might I expect to pay a dedicated lead-setter or support agent to monitor my website and answer questions in real-time?

For a dedicated lead-setter or support agent working 8 hours a day to monitor your website and answer questions in real-time, you can expect to pay between \$2,500 to \$4,000 per month. This estimate typically includes the cost of hiring a skilled professional who can manage customer interactions efficiently, using tools and systems to provide prompt and accurate responses.



Limitations: Are there any limitations or scenarios where the AI might not be the best fit for lead generation? For instance, industries or target audiences that may not respond well to AI interactions?

While scoping your project, if we identify areas where AI might be problematic—especially concerning sensitive topics—we'll recommend a human-centric approach. Though our platforms adhere to HIPAA and SOC-2 compliance standards, there are niches, such as within healthcare or finance, where human interaction is irreplaceable. That said, the value of an AI Expert is undeniable in most scenarios. For example, while we won't design a medical AI for diagnostics, we can craft one to route users to the appropriate medical professionals. Simply - we don't think AI is yet an appropriate tool for medical, financial, or sensitive advice or processing. We won't take on projects where we lack confidence in our ability to have a positive impact on you or your users.

All that said, we can create conversations that stay out of the specifics of sensitive advice. AI can easily answer pre-sale questions, for instance, as a part of lead generation.