



The Ultimate Guide to Choosing an AI Agent for Your Website



What are the best shoes for running a marathon?




BUY NOW



BUY NOW



Is a family suite available in April?



BOOK NOW

Yes, it has an ocean view.



Custom Agent

Resort Activity Agent

Restaurant Finder Agent

Concierge Agent

Create Agent

Name: Now Playing Movies Agent

Description: When someone asks about movies that are now playing

Method: GET

Book 2 seats for Back to the Future.



Booked!



Back to the Future
December 4, 2025 @ 7 pm

Seat 121 Section GW



I ordered a blue sweater from your site but I received a green one instead.

Attached is a prepaid return label and I've ordered the blue sweater which will be sent to you via express shipping.



Order #15T9M4

Get Label

Thanks!



Table of Contents

- 2 | AI Agent: Benefits and Definitions
- 3 | Generative AI Overview
- 4 | From Clunky Chatbots to Intelligent AI Agents
- 5 | Why Having an AI Agent Matters Now
- 7 | Use Cases of an AI Agent Across Verticals
- 9 | What to Look for in a Copilot Platform
- 13 | What to Look for from a Team’s Perspective
 - What Developers Want
 - What Marketers Want
 - What Content Creators Want
 - What Digital Strategists Want
- 15 | Evaluating a Copilot Platform Solution
- 15 | Conclusion
- 15 | About ai12z



You've heard a lot about artificial intelligence (AI) by now – some good, some less so. You may have been tasked with figuring out a way to use it for your organization – or been told you shouldn't.

Either way, AI is here to stay and organizations that use it will become more competitive. Those that don't, well, they can expect to be left behind. So, if you are looking for guidance on an effective way to use AI in your organization, this eBook is for you.

Introduction

There are a lot of new generative AI (GenAI) tools out there – especially for writing. Tools that help you write better emails or write your press releases or product descriptions. And tools that let you send out emails and bring in potential leads. All those tools are great.

But we're going to talk about a different approach to gain value from AI. This eBook looks at how AI can impact your customers' digital experience every day so that you can deliver better business outcomes on your website, whether your goal is to increase sales, downloads, donations, or memberships.

This guide is intended for marketing, digital strategy, growth, or customer success leaders. It provides information to get you started on your AI journey, and to help you and your organization choose an AI agent for your website.

The eBook details why you need an AI agent, how you can take advantage of one on your website and any application you run, what product features you should look for, and how to evaluate the platform (known as a copilot platform) that underpins an AI agent.

AI AGENT

Benefits and Definitions

Every company has a website. People go online every day and land on websites to find information or to take actions such as making a purchase or a reservation. But these experiences still aren't the greatest – website visitors face endless browsing, scrolling, and searching for information to find what they're looking for or to complete a task.

These experiences can become exponentially better through AI.

So, if you are wondering “Where should my company start with AI?” we suggest starting with your website.

Just as every company now has a website, inevitably every company will eventually have an AI agent for its users. These will help marketers turn browsers into buyers and prospects into customers, and ultimately move individuals further along their customer journey, faster.



An AI agent can be deployed on your website or via your mobile app. Most of the time, your site visitors have a goal when coming to your website. An AI agent will help them find what they need faster and likely increase your conversion rates.

Generative AI makes AI agents possible, so before going any further, let's look at the basics of what generative AI is and define terms like LLM, RAG, copilots, tokens, and more.

An AI agent may look like a traditional chatbot, but it is much more powerful. An AI agent can:

- Answer your site visitors' questions
- Automate workflows
- Provide recommendations
- Access real time data from third-party systems by using agents (specialized tools to extend the agent's capabilities beyond basic question answering)
- Present calls-to-action to the user
- Make a reservation (such as for a restaurant, activity, tour, etc.) or complete a purchase

Terminology often evolves as quickly as technology. An AI agent can also be known as an AI assistant or as a digital assistant or as a copilot, so you may have seen those terms used interchangeably.

For this eBook, we use AI agent for the entity that interacts directly with the website visitor. You can also imagine it as a website concierge for your website visitors. We use the phrase copilot platform to refer to the software platform that powers the agent. (Note that when we say, “copilot platform” we are not referring to Microsoft copilot, which is that company's proprietary “assistant.”)

Essential Terms

- Artificial intelligence (AI)
- Generative artificial intelligence (GenAI)
- Machine learning (ML)
- Deep learning (DL)
- Artificial neural network
- Discriminative model
- Generative model
- Large language model (LLM)
- Retrieval-augmented generation (RAG)
- ReAct (Reasoning + Acting)
- Reasoning engine
- Transformer architecture
- Token

OVERVIEW

Generative AI

Artificial intelligence (AI) has been around for some time. AI is a branch of computer science where computer systems show behavior and perform tasks like decision-making that could be interpreted as human intelligence.

Machine learning (ML) is a subset of AI. It's a program or system that uses data to train a model that is designed to learn on its own. The model looks at patterns to learn new things and performs tasks without being explicitly programmed. Most people are familiar with ML because it's used in everyday experiences like music recommendations.

Deep learning is a subset of ML. It uses artificial neural networks that mimic the human brain to process even more complex patterns than traditional ML.

Artificial neural networks are composed of interconnected nodes known as neurons. These neurons process large amounts of data, learn complex patterns, make predictions, and learn how to perform tasks. Think of robots or self-driving cars that use machine learning's neural networks to detect pedestrians crossing a street.

Deep learning can be categorized as discriminative models or generative models.

Discriminative models are usually trained on labeled data. They learn the relationship between the features of the data points and their labels, focusing on distinguishing between classes. For example, a discriminative model might learn to recognize which patterns (like a shape or texture) in an image correspond to a "dog" or a "cat." In other words, discriminative models are primarily used to classify or differentiate between existing data points.

Generative models generate new data that is similar to the data they were trained on. For example, a generative language model can predict the next word in a sentence based on the context. Given the prompt, "The sky is __," it might predict "blue." However, with the prompt, "The weather has thunder and lightning," the model might generate, "The sky is stormy." Or for "Peanut butter and __," it would likely predict "jelly."

Generative AI generates new content (whether text, images, audio, or video) as its primary output, based on what it has learned from the existing data.

Generative AI relies on a **large language model (LLM)**. This is a neural network model that is trained on trillions of pieces of content including web pages, articles, books, and code. (Training means learning from content.) An LLM understands natural language and code and learns associations between words. It is then trained to predict the next word in a sentence, generating it based on previous words it has seen before in that context.

And it's trained to solve common language problems such as:


- answering questions
- completing sentences
- summarizing documents
- classifying text
- translating text into different languages.

There are two popular frameworks for generative AI:

1. RAG (retrieval-augmented generation): answers questions based on your organization's own content, such as website pages or PDF documents stored and processed in a vector database. RAG retrieves relevant content and uses an LLM to generate coherent and contextually accurate responses.

2. ReAct (Reasoning + Acting): the LLM will reason about a task and then act (using different tools or agents) before arriving at an answer. RAG is just one of the many possible agents it can use. In this setup, the model can query your content through RAG, as well as interact with a range of other external services (like Google Maps, email clients, weather APIs, and more) as it reasons step-by-step to arrive at the best answer.

We will discuss these frameworks more in the section "What to Look for in a Copilot Platform."



An LLM is a reasoning engine.

Think of it as a human brain: it understands context and meaning, and it can reason by putting together a plan to respond to a user's query. When a user interacts with the AI agent, the LLM determines whether it needs to call different integrations to retrieve data from other systems. For example, imagine your company is a printer manufacturer. A user visits your website or mobile app and asks the AI agent, "How do I replace the ink?" For the LLM to generate an answer, it may first ask, "What printer do you have?" This is a simple use case of RAG, where the answer is found in the user manual for the specified printer, which the agent can access. Now, consider a scenario where the user asks about the status of a purchase order for printers. This involves ReAct, where the AI agent acts as a concierge: it identifies the relevant third-party system to retrieve real-time data about the purchase order and provides the user with the status. Similar logic applies to other use cases, such as asking about the weather, hotel bookings, or dinner reservations.

Well-known LLMs include OpenAI GPT, Google Gemini, Meta Llama, Anthropic Claude, Amazon Nova, and Mistral AI.

The performance of the LLM typically improves when the model has more parameters; this increases its capacity to learn and store patterns, thus allowing it to generate more accurate results. A larger model has a lot more neurons, which means more associations between words and groups of words. More vertically focused LLMs are also being developed that are suitable for industry-related questions.

Generative AI typically involves the **transformer architecture**, which is known for language understanding. In this architecture, the transformer models relationships between all words in a sentence.

A token is a part of a word and is the atomic unit that LLMs work in. For English, one token is about three-quarters of a word, or you can think of this as 100 tokens equals 75 words. LLMs generate text one token at a time. They look at the current context, a fixed number of tokens, and predict the next token.

Essentially, generative language models are:

1. Neural network models
2. First trained to predict the next word in a sentence
3. Optimized to perform tasks well when given instructions

Generative AI can be applicable across many use cases, including music composition, code generation, content creation, language translation, agents/conversational agents, and summarization.

From Clunky Chatbots to Smart AI Agents

You've likely used an agent on a website before. And you've likely been frustrated by the experience.

That's because traditional agents don't understand context and meaning. In other words, they don't understand what we're asking. And so, we're put in a queue and left to wait for a live representative to get our questions answered.

For any developers who've worked on building a agent, it probably took months to build. This is because traditional bot frameworks require frequent manual updates, extensive training data, and programming for numerous question variations. Their reliance on decision trees makes them struggle with unexpected inputs. Going live with a agent has typically been complex, requiring significant coding effort. Last, but certainly not least, these frameworks are difficult to integrate with third-party systems.

But the power of GenAI changes all of that, and it's now easy for organizations to go live with an AI agent.



An intent model cannot reason – it is programmed with about 20 variants, and it does a fair job of handling them, but can still fall short. You can spend a lot of time debugging queries that seem to be worded exactly like a variant, only to find that a minor difference, like a user typing 'an' instead of 'a,' caused the issue.

Why Having an AI Agent Matters Now

Remember the early days of the World Wide Web?

Businesses started to have an online presence.

And then Google Search and SEO became important.

And so did e-commerce, where retailers could start selling their products online.

Then came mobile. The list goes on.

We're in our next technological revolution with generative AI.

Now, every company will have an AI agent for their users – just as every company has a website. Why? Because AI agents offer the following advantages:



24/7 support with instant responses to questions

An AI agent ensures your company is always available, whether someone is shopping at midnight or researching during business hours. In addition, your company can scale customer support to handle high volumes during peak times and manage thousands of inquiries without compromising response quality. By ingesting your website content, the call center PDF, and any other product manuals, guides, or documentation, an AI agent can answer questions quickly, reducing customer wait times and cutting operational costs. You can always have live service escalation if users need to speak directly with a representative.



Increased sales and conversion rates

An AI agent can act as a digital sales agent/concierge/advisor, etc. guiding prospects through their buying journey with product recommendations and upselling opportunities. In addition, many site visitors will leave a website because they can't find what they're looking for, but an AI agent can help visitors find the right information before they abandon the page. Plus, it helps prospects make faster purchasing decisions, leading them to buy sooner.



Direct paths to action

Inserting the right call-to-action button at the right time allows your users to immediately take the desired action ("Buy Now," "Book a Demo," "View Related Products," "Book a Reservation," "Sign Up for Event," etc.) instead of having to navigate through multiple steps or pages. This will likely increase your conversion rates and drive business outcomes, in addition to enhancing engagement and retention.



Personalized experience

A prospect can tell an AI agent what exactly they're looking for, or you can connect the agent to your CRM to provide responses based on their preferences or past purchases. This eliminates the traditional research project that website visitors are forced into today; instead, the agent becomes a personal concierge.




Internal efficiency

By handling routine queries, AI agents allow your team to focus on more strategic and creative work rather than on routine tasks.

USE CASES FOR

AI Agents Across Verticals

Every company, across every vertical, can go live with an AI agent. Let's look at some industries to see how these digital assistants can benefit those users when they reach a website.



AI AGENT USE CASES



Media & Entertainment

People come to a theater website to see what’s playing. They may have specific criteria in mind, such as the type of show, whether it’s family-friendly, a certain price range, or specific dates and times.

Typically, a website visitor spends time browsing multiple pages and navigating filters to find information. But an AI agent can serve this information directly to prospective theatergoers, helping them find the right show, select seats through a 3D seat map, and purchase tickets. It can also answer general questions, such as: “How do I get to the venue from downtown?” or “What food options are available?” or “What is Hamilton about?”



Manufacturing

On a manufacturing website, users may browse products, check their order status, or seek answers to service-related questions.

An AI agent can enhance their experience by enabling them to compare two products side-by-side to evaluate features and differences, check the status of their order, or troubleshoot product issues.

Internally, employees at a manufacturing company can use an AI agent to quickly access information from training materials, standard operating procedures (SOPs), or inventory management systems, and provide customer service and sales teams with access to product information, allowing them to address inquiries quickly.



Nonprofits & Associations

Nonprofit and association websites want to educate visitors, provide resources, and increase their donations.

An AI agent can answer any questions their visitors may have, guiding them to key resources such as program details, upcoming events, or volunteer opportunities. It can also simplify the giving process, walking donors through the steps and even providing updates on the impact of their contributions.



Retail

Retail stores can use an AI agent on their website or mobile app to act as a digital shopping advisor. The agent can answer customer questions about products, provide personalized recommendations based on specific criteria, and help shoppers compare products. It can also check product availability and allow customers to make purchases directly with a “Buy Now” button, making the shopping experience easier and faster.



Education

Prospective students go to a university’s website because they’re interested in finding out more information about the academic programs and/or campus, booking a tour, or checking the status of their application. An AI agent can answer their questions and provide images and videos, allow them to book a tour, and inform them that their application is in review.

A current university student could use an AI agent to get course recommendations for the next semester. Instead of making an appointment with a student advisor days later and all the way across campus, that student can talk to the AI agent. The agent connects to the student CRM, the course directory, and the class calendar, so it understands what courses are required for the student’s major and what courses they’ve taken already, and it can provide a recommended course list and schedule. It can even accommodate a preference such as wanting to take Art History as an elective.



Travel & Hospitality

Vacationers often research not only their destination but also the hotel or resort where they plan to stay.

An AI agent on a hotel website will serve as a digital concierge for travelers. They can ask it questions like:

- What activities does your hotel offer? It answers with a list of activities and prompts the guest to sign up by providing forms they can fill out to reserve a spot.
- What restaurants are nearby? It provides a list of restaurants, images, the distance from the hotel, and the opportunity to reserve a table.
- What room do you recommend for a family of four? Or, what room do you recommend for a person with a walker? It will make appropriate recommendations and enable the future guests to book the room.
- Can I book a spa treatment? It serves up the form for the user to select the treatment type, date, and time.

What to Look for in a Copilot Platform

Many agents or conversational platforms say they use AI, but they still require a lot of code to get going. Or they're only able to handle RAG, meaning they can do question answering but can't work with third-party tools, GraphQL, or agents like Google Maps, email, SMS, or a CRM to provide more relevant and actionable answers.

Here's what you should look for in a copilot platform solution to power your AI agent:

1. RAG

(retrieval-augmented generation)

RAG is a technique where a prompt provided to your AI agent retrieves your organization's content from a vector database and combines it with a large language model (LLM) to handle question answering (QA).

Here's how it works:

Your content is uploaded into a vector database via a CMS connector or site scraper. In the database, the content is converted into vectors (numerical representations, also called embeddings) that encode the meaning of words, images, and other data. These vectors enable the database to understand and compare content based on context and meaning. The vector database retrieves and surfaces data that closely matches the query, even across different languages.

When someone asks a question, that query will be turned into embeddings. Then the RAG system searches the vector database for the most relevant content embeddings that match the question. The system might review 20 documents, and re-rank them to select the three to six most relevant ones to answer the question.

The content retrieved from the database is then used by a generative model to create a coherent and informative answer. Hence the name, retrieval-augmented generation.

2. Image AI

Users always get text in response to their agent interactions, but why not images? After all, a picture is worth a thousand words.

Imagine an individual planning a trip to London. They're browsing a travel website for information about the city. With an AI agent on that tourism website or mobile app, they can ask questions and get replies with not just text, but also images.

The images are retrieved from the tourism website – or, in other use cases, from marketing brochures or product manuals. In short, the AI agent gets images from the content that the company has given the copilot platform to ingest or scrape.

And great image AI can go beyond what is just in the ALT text. It can analyze the image using an AI vision model, enhance the image description, and create an image caption that shows up below the image.

When a site visitor asks a question, the copilot platform will examine all the content that's been ingested in order to provide an answer with contextually relevant images. Leveraging image AI ensures that the visuals match the answers provided.



Most copilot systems do not have image AI, so be sure to check with the vendors you're talking to.

3. Multilingual

Your customers are global. Having an AI agent that understands and responds in their language of choice is important.

An AI agent can respond in the same language as the user's query, providing a localized experience. For example, if they ask a question in French, the AI agent can respond in French.

Also look for a system that detects the user's language from their browser settings so that the control (aka the AI agent) will show up in the same language as their browser – for instance, any text in the agent would display in their language.



4. Integrations

Integrations allow your AI agent to perform tasks beyond simple question answering by connecting with third-party systems to retrieve real-time data and execute specific operations. This means the agent can handle scenarios such as completing a purchase, booking a reservation, or filling out a form, which creates a more engaging and productive user experience.

An AI reasoning engine within an agent identifies which integrations are available and what parameters can be passed to them. Based on the user's request, the engine can initiate a workflow that uses multiple integrations to accomplish specific tasks.

A great copilot platform includes out-of-the-box integrations your AI agent can call to start workflows or complete tasks, such as Google Maps, CRM, Salesforce case management, email, SMS, stock, weather, Google Search, and more.

Organizations should be able to easily create custom integrations for their specific business needs. Their backend systems (like marketing automation platforms, CRMs, and other technologies) can be connected to an AI agent via a JSON REST API or GraphQL.

What's a JSON REST API?

- JSON stands for JavaScript Object Notation. It's a format in which data is exchanged.
- REST API stands for Representational State Transfer Application Programming Interface. This is a style of web service design that uses HTTP requests to interact with resources.
- Each resource typically has its own endpoint, and actions are specified by HTTP methods (GET, POST, PUT, DELETE).
- A JSON REST API is a REST-style web service that uses JSON to transmit data between the client and the server.

What's GraphQL?

GraphQL is a modern way of calling and fetching data. It allows you to easily filter what you want to return to the AI agent. For instance, if a user asks a hotel website's AI agent about movies playing nearby, GraphQL could filter specific information to return, such as the name, rating, and available times, while excluding details like movie description and release date.

How are they accessed?

Both a JSON REST API and GraphQL use URLs to locate their endpoints.

What's an endpoint?

- An endpoint is the specific web address your AI agent contacts to fetch or send data.
- In a JSON REST API, each endpoint is tied to a specific resource and returns data in JSON format.
- In GraphQL, there is usually a single endpoint that can process flexible queries, allowing you to specify exactly what data to retrieve.

Integrations can also be HTML templates, which are widgets. One example is a mortgage calculator. Another example is an integration that calls an endpoint through a REST API or GraphQL and returns the data in a Google Chart. There are thousands of out-of-the-box widgets you can use, and you can also create your own using JavaScript and HTML.

Integrations can call other LLMs as well, which could be valuable when you need a specialized model like one from Hugging Face. The ability to call Hugging Face models and APIs allows you to leverage a rich ecosystem of pre-trained models, tools, and resources without having to train large models from scratch. It can accelerate development, reduce costs, and improve the quality of your AI solutions.

5. Forms

An AI agent should have forms to capture leads or allow site visitors to book a reservation or tour. Non-technical users should be able to create these forms easily.

Look for a copilot platform that offers multipage forms with formatting options, dynamic time slots, responsive design, multimedia elements (images and videos), inline help, various input types (text, dropdowns), and the ability to connect with backend systems.

Example of a multipage form:

When a user books an activity via an AI agent, the first page of the form might ask for their name and email address, the second page lets them select a time slot, and the third page allows them to choose add-ons, like a golf cart if they booked a tee time.

6. Comparisons

Your site visitors have products in mind, but they don't know which to choose. To help them decide, they can ask an AI agent to compare products for them.

This is ideal for comparing multiple items. The agent will gather data for each item and then provide the comparison to your customer or prospect.

Note that this only works with the ReAct Framework. With RAG, the AI agent would need to find a document that already compares the items. With ReAct, it will call the RAG system multiple times to get information for each product, then the Reasoning LLM will analyze each product and create the comparison.

7. Scrolling list or carousel slider view

When querying a backend system with multiple items, such as nearby restaurants, it's important to present the data in a user-friendly format. For example, you might have an integration that retrieves all the restaurants nearby, and you want to bypass the LLM to send the results directly to the client. These results can then be displayed in either a scrolling list or a carousel slider view.

Choose a copilot solution that supports scrolling lists and a carousel slider view for displaying items visually. A carousel view in an AI agent works perfectly when images are available, making it easy for users to browse options.

Imagine you're displaying a list of theater performances with dates. A carousel view allows users to easily view what's playing. When a user clicks "Available Performances," they'll see a carousel they can scroll through to explore performances and select specific dates and times for a particular show. Users can also click "Back" to return to the main list of options.

Custom carousels allow you to call specific endpoints, such as MLS listings in real estate. To maintain your brand identity, the carousel can be passed through a custom template, ensuring your data is displayed in a consistent and organized way.

8. Structured data

Structured data is information organized to make it easy to search and process. It's like the data you find in tables in a database or arranged in an Excel spreadsheet. Digital tools can quickly find and use relevant data, which is especially helpful for tasks such as managing product inventories in retail, navigating course catalogs at universities, or optimizing supply chains in manufacturing.

A copilot platform that supports ReAct can enhance traditional RAG models with integrations that specialize in structured data. By allowing the LLM reasoning engine to pass parameters to an integration specifically designed to handle structured data, the system can quickly filter and retrieve relevant information based on user-defined criteria. Whether the data resides in databases, PIM systems, or CMS platforms, the integrations can interact with these backend systems to fetch and filter data. A copilot platform that can handle structured data is able to streamline data retrieval and boost the relevance of results.

Let's look at an e-commerce example:

a company that sells medical devices, including a wide variety of blood pressure monitors.

When a customer asks the AI agent which blood pressure monitors fit specific criteria – such as price range, features, or brand – the integration quickly filters the structured data to retrieve relevant products for the customer. This can be done via GraphQL or a custom integration that can query a system with filters based on the parameters passed to the LLM.

9. Voice

Some users prefer to use voice rather than typing when interacting with an AI agent. Look for a copilot platform that offers a feature where users can ask questions via voice in their language of choice and hear the chatbot's voice response.

10. Call-to-action (CTA) buttons

You want to be able to present call-to-action (CTA) buttons to the user to prompt them to take action.

Examples of CTAs could be:

- "Buy" (an item)
- "Book" (an activity)
- "Reserve" (a table)
- "Apply" (for an apartment)
- "Schedule" (a demo or an appointment)
- "Contact" (fill out a form with contact information)
- "Watch Trailer" (for a movie)
- "Donate" (to a cause)

This will guide users to the next step in their customer journey and help increase conversion rates for your team. These CTAs should trigger actions or responses by the reasoning engine, such as providing additional information or initiating tasks. For example, a CTA like "Learn more about golf" could prompt the agent to provide details about golf packages with another CTA to "Sign up for golf."

CTAs should also support web events, such as JavaScript events, enabling web developers to define actions on the event page. For instance, a CTA like "Add to Cart" could be programmed with an item ID that integrates seamlessly with the website's shopping cart.

11. E-commerce

An e-commerce feature is essential. An AI agent on your website can be a personal concierge for your visitors and even turn browsers into buyers.

It can answer their questions about the products you offer and allow them to compare options right in the agent interface.

Interacting with the agent, they can provide their preferences and price range and view a list of options. If they make a decision, they can click on the "Buy Now" button to purchase right from the agent or put the product(s) in their shopping cart.

12. Brand guidelines

Ideally, your copilot platform will let you set the agent's voice and tone and follow your brand's identity. In this case, you can upload your brand guidelines into the copilot platform. You can also brand your AI agent with your icon and colors.

13. Avoid hallucinations

You want to be able to instruct your copilot to only use the information from your organization's content, rather than the knowledge it has from an LLM. That way it will not generate replies with false information (known as hallucinations). Instead, you can tell your AI agent to respond to queries that aren't related to your brand with "I don't know" or "I don't have information about that. Please ask me questions related to [insert your company's name]."

14. Security

Organizations are concerned about their data, and rightly so. Look for a copilot solution that will not use your data to train the large language model. That means that any content you upload, or any chat data (customer questions and copilot answers) will not go into the LLM to train that model.

In addition, you want your AI agent to detect sensitive or inappropriate queries, such as those related to hate, explicit content, or harmful topics, to ensure responses remain ethical, professional, and aligned with brand values. This safeguards the user experience and can prevent harmful or offensive interactions.

Ensure that your vendor is building their technology to be SOC 2 compliant. SOC 2 is a compliance framework used to evaluate and validate an organization's information security practices. It's widely used in North America, particularly in the SaaS industry.

If working in healthcare, understand whether the vendor is willing to sign a BAA (Business Associate Agreement). In United States healthcare, if a vendor handles protected health information (PHI) on behalf of a covered entity (such as a hospital, clinic, or insurance provider), then that vendor is considered a "business associate" and must sign a BAA. Under the Health Insurance Portability and Accountability Act (HIPAA), covered entities are required to enter a BAA with any business associate to ensure that the PHI is properly safeguarded. A BAA clearly delineates each party's responsibilities regarding the privacy and security of PHI, stipulates permitted uses and disclosures, and outlines the consequences if the vendor fails to comply with HIPAA's requirements.

15. Fine-tuning an LLM

Your copilot platform should allow you to fine-tune the LLM. LLMs are trained on a wide range of content. While this enables the LLM to have general reasoning and understand multiple languages, the LLM can still lack domain-specific proficiency. That's where fine-tuning can come into play. Most of the time, it's not required for an AI agent, because you can simply add examples to the system prompt of your copilot platform. But there may be times where you want to include a significant number of examples. Fine-tuning allows you to adapt the model to specific tasks, domains, or styles, and will improve relevance and accuracy.

Fine-tuning on curated, domain-relevant data, such as technical documentation, medical articles, or legal briefs will help the model internalize the precise vocabulary, jargon, and nuances of your field.

Fine-tuning will also save you money and time by providing chosen examples rather than collecting massive amounts of new training data and building a model from scratch. Instead, you can adjust and refine the model on smaller, focused datasets, reducing the time-to-market for specialized applications.

16. Support development, staging, and production environments with sync

Just as websites support development, staging, and production environments, your copilot should do the same. A sync solution for your AI agent makes this possible. It enables you to refine content and features without disrupting live deployments. Export and import capabilities allow you to bundle your copilot's configurations and assets, automatically replicating them across environments without the need to do so manually.

Sync functionality allows you to test updates, gather feedback, and ensure stability before going live. You can iterate on new features or content while maintaining consistent behavior across environments. Sync capabilities also simplify deployment timelines and make it easy to create new environments for testing, development, or demonstrations without affecting the main production environment.

17. Analytics

It's important to track how your copilot is performing to improve user experience and outcomes. Look for a solution that provides detailed engagement metrics, such as the total number of questions asked, session durations, and peak inquiry times.

Thumbs up or down ratings can help you assess response quality and user satisfaction. Analytics should also highlight unanswered questions to identify areas for improvement. An "I don't know" response from an AI agent may indicate that relevant content is missing.


Tracking what users are searching for and asking can reveal what they're most interested in and highlight any content gaps. Addressing these gaps can ultimately improve your SEO and better align your content with user needs. Additionally, monitor the performance of calls-to-action buttons, such as click-through rates and conversions. Integration with tools like Google Analytics allows for deeper analysis of trends and user behavior.

A copilot with built-in quality assurance testing helps validate its functionality and the performance of its agents. It allows you to create and manage test scenarios, run tests to evaluate your copilot's responses and agent workflows, and review performance metrics to ensure everything is working as intended.

With these insights, you can continually analyze and optimize your copilot's performance to better serve your users and achieve your business goals.

What to Look for from a Team's Perspective

Each team will have their own wish list when going live with a copilot platform solution. It is important to have a solution that meets the needs of each of these groups.



What developers want

Your developers will want a copilot platform that has enterprise scalability, flexibility, and the right development tools, as well as the ability to connect to other key business systems.

Low-code and developer productivity

Many systems require a lot of code when going live with an AI agent. Developers have likely experienced this if they've worked with Microsoft Bot Framework, Amazon Lex, or Google Dialogflow. Look instead for a low-code solution for your development/engineering team. Developers are more efficient when working in an environment they're at ease with. The platform should be intuitive to use, giving them the ability to accomplish more while writing less code.

Structured data

Your AI agent does not have to be limited to handling unstructured content. Its integrations can connect with structured data systems, such as product information management (PIM) systems, or use GraphQL agents to filter and refine responses from existing data sources. Find a solution that works with structured data to deliver accurate and relevant responses.

Flexibility to build

While you want to minimize the required effort, you still want to give developers the flexibility to easily build custom integrations and connect them with your backend systems. This lets the agent provide a personalized experience to the user.

Whether syncing with a CRM, Google Maps, or backend systems via REST APIs, these integrations can then enable the AI agent to automate workflows and complete tasks for the user.

Documentation

Look for a solution that has thorough, user-friendly documentation that includes both text and training videos. You also want to work with a team that provides support either through its own professional services or through its partner ecosystem.

What marketers want

Marketers want the website to provide engaging customer experiences in order to achieve greater business results online – leads, sales, downloads, memberships, and donations.

They are generally non-technical users who are looking for a copilot solution where they can go live with a Q&A agent in less than a day, and one which requires no code.

Branded AI agent

Marketers need to be able to input their brand guidelines and tell the AI agent how to behave. They can be specific about the tone of voice they want the agent to embody, reflecting your organization's brand identity.

Content is still king

Content matters. It's what makes up your website, your blogs, your brochures, and your user manuals. You have and will always need great content. Marketers can easily take the content they've already created – whether from their website or from articles or marketing materials – and upload that into a copilot system. The platform will then use that content to answer your customers' questions. They can even instruct the system to exclude certain pieces of content (such as blogs published before a specific year) or prioritize answering questions with information from specific web pages or documentation instead of blogs.

If analytics reveal that existing content doesn't address certain questions, marketers can create question-and-answer pairs, i.e., a question and its corresponding answer. They don't have to worry about creating different variations of that question as the LLM reasoning engine is smart enough to know. For example, if someone asks, "What's the address of your office in the UK?" the LLM would know that's the same question as "What's your England headquarters address?" or "What's the address of your London office?"

Create a relevant experience and increase conversion rates

Marketing has become a revenue generator, driving growth and leading customers further down the funnel. Presenting relevant content and answering a user's specific questions can turn the user from a prospect into a customer. A user can input their preferences or be specific about what they're looking for, and the AI agent can provide recommendations and prompt the user to take the next step with a call-to-action. These can be sign-up forms, reservations, or buy now options.

Analyze and optimize

Marketers also want to be able to analyze and optimize results. They can review the questions being asked of the AI agent and see which are or are not being answered, or which answers received a thumbs up or a thumbs down.

Seeing which questions received an "I don't know" answer or a thumbs down will help marketers be able to improve their SEO results. They'll learn what content their site visitors are searching for, and if they don't have that content, they can create it. They can also review the tone of the agent when reviewing answers to questions.

What content creators want

Content creators will need to think about content not just for their website, marketing brochures, blogs and press releases, emails and newsletters, videos, podcasts, ads, eBooks, and whitepapers, but now also for their AI agent. Think of an AI agent as a new digital channel that will require its own specific content.

Easy content creation and multilingual capability

Provide content creators with a copilot platform that allows them to easily create new content for this new digital channel. In addition, they'll want a platform that can talk to its users in multiple languages for their global audience.

What digital strategists want

Digital strategists are looking for a copilot platform that can enhance user experiences, drive business outcomes, and easily integrate into their existing digital ecosystem.

Personalized experience

They want to present users with a personalized experience, answering their exact questions and providing them with calls-to-action to move them further along their customer journey.

An AI agent should have suggestions for products or services that fit their needs. And it should be able to execute complex tasks.

Integration with backend systems


The digital strategy group will look for a system that can integrate with their CMS, CRM, e-commerce platform, and other existing systems, and have the APIs and webhooks to easily connect to third-party systems to access that real-time data.

Flexibility and scalability

Look for a solution that supports multiple use cases, such as customer support, guiding customers and prospects to take the next step, and even internal use cases. You also want a solution that has a scalable infrastructure so it can handle growing traffic.

Business outcomes

Digital strategists will want to drive more online results with a copilot platform that engages users, answers their questions, and integrates CTA buttons like "Buy Now" or "Sign up to become a member" to move them further along their customer journey.



Evaluating a Copilot Platform Solution

In addition to the features already listed, you want to ensure the platform you choose fits well into your current infrastructure and will interoperate with your existing sales and marketing applications.



Cost and scalability

A lot of systems are pricing their solutions at a dollar or two per transaction! That is far too high. Copilots in enterprises with many site visitors will interact a lot with those individuals. Look for a system that charges cents per interaction, not dollars. In other words, look for a solution with a platform that is meant to scale.



Integration with different LLMs

Look for a copilot solution that can work with different large language models, from OpenAI's GPT series to Meta Llama to Anthropic Claude and more.



Free Trial

Getting a free trial allows you to try out the solution hands-on. Look for a vendor who will give you a 1 week or upwards of a month for a free trial.meant to scale.



Training and support

You'll want an organization that provides training and support options to ensure the success of your copilot(s) implementation. Such an organization is likely to offer discovery, getting started, or professional services packages that you can choose from.



Working with an agency

If your internal resources are limited, consider collaborating with a digital or web agency. Find one who will work with you to understand your business goals and deploy the copilot solution for you.

Conclusion

The best copilot platform drives online results by handling question answering and moving customers further along in their journey. It enables users to take meaningful actions that increase conversion rates, whether it's booking appointments or completing purchases, or navigating multistep workflows. With its reasoning engine and ability to call integrations that interact with external systems, the platform becomes a true concierge and digital assistant for your users.

Everyone wants a personalized experience, and that's exactly what a copilot platform delivers. Instead of users struggling to find information or complete tasks on a website or mobile app, an AI agent provides immediate support. This leads to more satisfied customers and increased sales, bookings, memberships, and success for your business.

About ai12z

ai12z powers digital assistants on websites and mobile apps, helping site visitors quickly find information, complete tasks, and make purchasing decisions. With ai12z's AI agent, marketers can engage and acquire new customers, driving conversions while moving users further along their journey.

Learn more at <https://ai12z.com>

To see a live demo of ai12z's copilot platform, please click here:

REQUEST A DEMO