PolyAI

Scaling AI in the enterprise:

Leading with agentic AI.

INTRODUCTION

Contact centers are embarking on a digital transformation journey.

AI adoption is firmly on the C-suite's agenda.

According to a **recent report**:

80%

of CX and contact center leaders say their C-suite is pushing for faster AI adoption

1/3

say their top goal over the next 5–10 years is to implement an AI solution that delivers clear ROI, turning pressure from the top into a career opportunity.

Agentic AI is a growing part of that future. **Gartner predicts** that by 2028, a third of enterprise software will include agentic AI, up from just 1% in 2024.

But adoption is already showing signs of strain. An MIT study found that:

95%

of companies' AI pilots are failing.

However, the same report found that **how companies adopt AI is crucial**. Purchasing AI tools from specialized vendors and building partnerships succeed about 67% of the time, while internal builds succeed only one-third as often.

The slow progress rate of agentic AI isn't just early-stage tech growing pains; it's largely because companies struggle with adoption strategies. Many internal pilots fail not due to the technology itself, but because organizations lack the right partnerships, expertise, and **change management frameworks** to deploy AI effectively. Initiatives are often rushed to keep up with market pressure, leaving little room to integrate AI in ways that deliver measurable impact.

Enterprises are already seeing real results with Agentic AI. This guide is about how to join them, not with demos or vague automation promises, but with something you can put into production and scale.

Rethinking how enterprise AI works

The next wave of enterprise AI is rooted in customer experience and systems that understand, act, and improve over time. This is the concept of the agentic enterprise: where businesses are built around intelligent agents that can act dynamically, learn continuously, and coordinate across functions.

In an agentic enterprise, AI-powered teams can be directed with precision. Unlike humans, they don't forget or get distracted. They can also be reconfigured instantly to meet new priorities and scale effortlessly as business needs evolve.

> A greater degree of automation gives you more visibility into your operations. As opposed to managing thousands of agents across locations and even vendors, you can essentially get a single pane of glass when you deploy AI agents. You're no longer guessing what's being said on the phone or digging through transcripts to analyze conversations between employees and customers.

Nikola Mrkšić
Founder and CEO of PolyAI.

Meet your newest team member: Agentic AI

Understanding agentic AI is the first step to successful implementation. Think of it like getting to know the personality and strengths of a new team member before bringing them on board.

Agentic AI is a type of artificial intelligence that can act on its own to get things done. Instead of waiting for constant human input, these agents can:

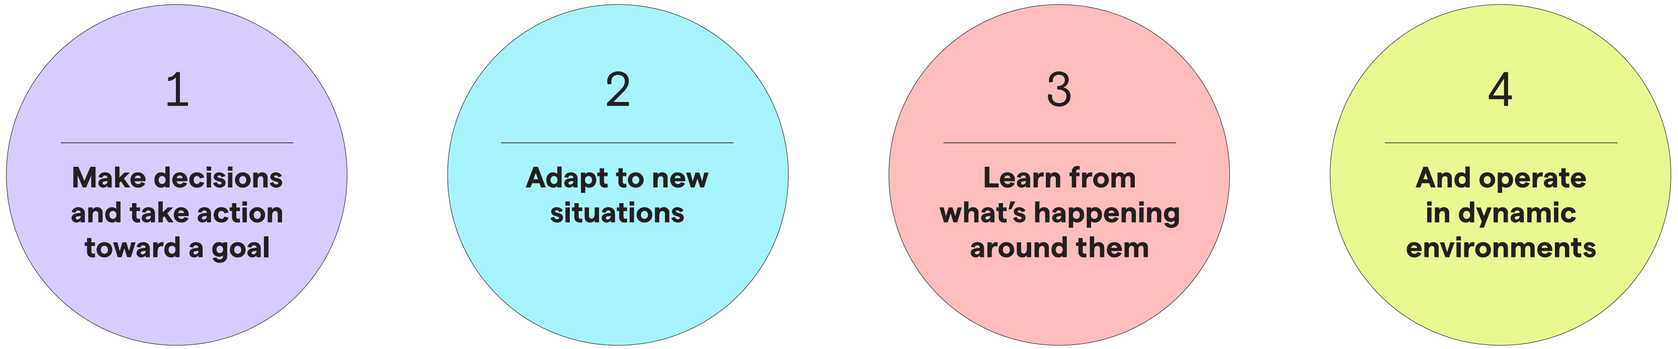

Whether it's handling repetitive tasks or navigating complex decisions, the key trait of agentic AI is its ability to work independently, guided by what it's been programmed to do or what it has learned over time.

What does an agentic workflow mean?

An agentic workflow is a process where AI autonomously manages tasks, coordinates actions, and adapts in real time without relying on human input at every step.

The 4 types of agentic AI

Not all agentic AI operates the same way. At a high level, agentic AI falls into four major categories: reactive, deliberative, learning, and multi-agent systems.

Each represents a different stage in autonomy, adaptability, and decision-making complexity. Here are the key characteristics of each:

Reactive agent

Reactive agents respond instantly to changing conditions without planning ahead or storing past experiences or internal models. They operate in the present, i.e., adjust their actions based on immediate inputs. While this makes them fast and reliable, it also limits their ability to improve over time.

Deliberative agent

Deliberative agents analyze, plan, and weigh different possibilities before making a decision. So, instead of reacting impulsively, they predict future outcomes and select the best course of action.

Learning agent

Learning agents evolve based on new data and past experiences. Unlike reactive agents that repeat the same behavior, learning agents continuously refine their decision-making, adjusting their responses to become more effective.

Learning AI is critical for long-term efficiency in customer service. An AI agent that handles thousands of calls per day should improve at recognizing customer intent, improving responses, and reducing friction in interactions.

Multi-agent systems

Instead of a single AI agent working alone, multiple autonomous agents collaborate to solve complex tasks.

In customer service, an AI agent might handle the initial customer inquiry, while a separate AI system retrieves account information, and another processes refunds or updates orders.

Instead of relying on a single AI model, multi-agent systems distribute tasks dynamically, improving response times and efficiency.

How does agentic AI work?

When agentic AI is implemented effectively, it’s like a highly capable employee. Not one that needs constant supervision, but one who can listen, think, act, and learn, all while handling dynamic situations on the fly.

Here’s how agentic AI processes tasks and problem-solves:

Perceive

AI agents gather and process data from various sources depending on the task they have been assigned.

Reason

A large language model acts as a reasoning engine to coordinate solutions for the task.

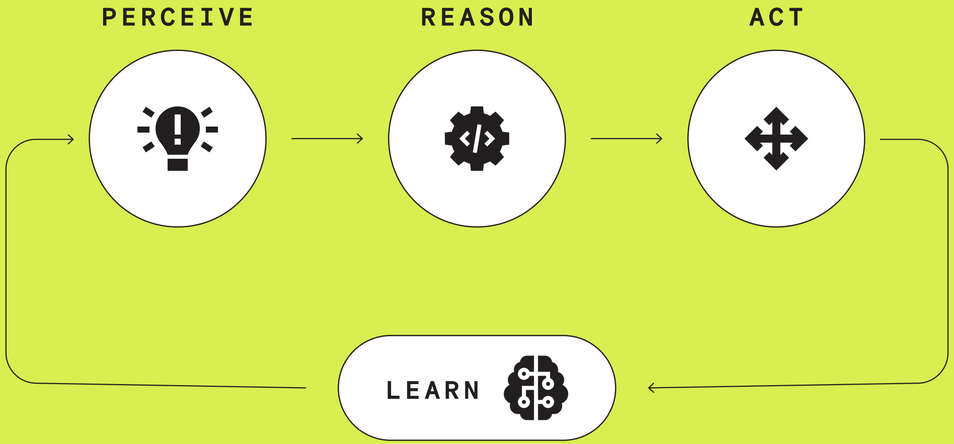

Act

The agent then acts based on its reasoning result, which may either involve calling on more APIs or acting on the response.

Learn

Agentic AI should continue to improve through feedback and continuous learning.

Understand the risks of agentic AI and take control

Putting customer interactions and business processes in the hands of agentic AI systems takes trust from both your business and customers.

The more autonomy AI has, the harder it can be to predict, control, or correct its behavior. If an agent is optimizing for efficiency, what happens when it chooses speed over accuracy or fairness?

Without clear safeguards, agents can make mistakes with no obvious way for teams to intervene.

But these risks can be managed with the right technical foundations.

SAFEGUARDS

How to build safeguards into AI agents

If a human employee wasn’t sure how to answer a customer, you’d expect them to double-check their response. The same principle applies to AI.

1 Retrieval-augmented generation (RAG)

RAG enables AI agents to cross-reference knowledge from a generative model with a knowledge base. This means an AI agent checks its generated responses against information your organization has confirmed as factual before responding.

This prevents inaccurate, irrelevant, and inappropriate responses and keeps customer conversations within established limits.

2 Documenting actions in your knowledge base

Many AI agents struggle to reliably take action because they have access to too many tools or APIs and little to no direction on which tool to use and when.

Breaking down your knowledge base allows you to ascribe specific actions to various topics. For example, you can build a specific part of your knowledge base related to updating account information.

Retrieval-augmented generation (RAG)

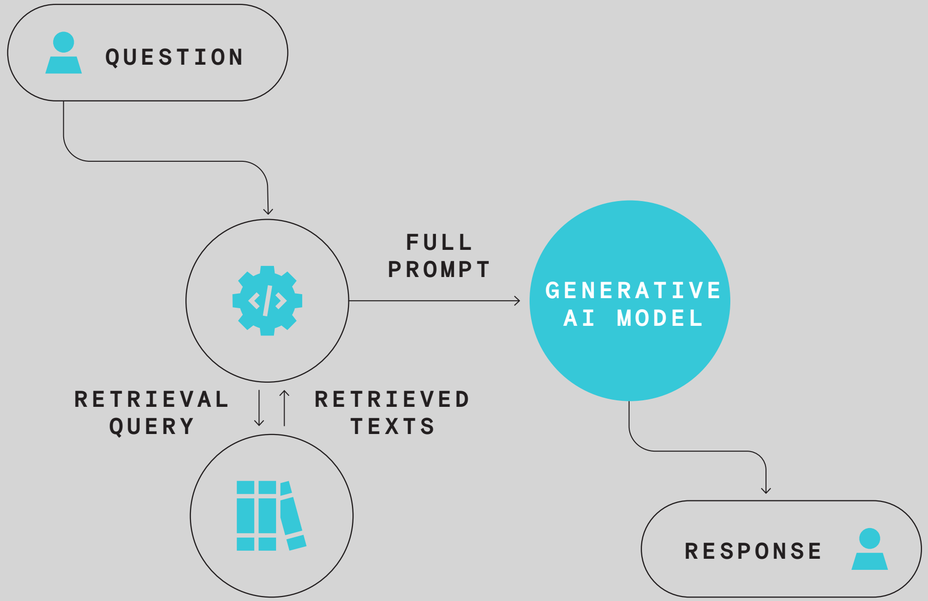

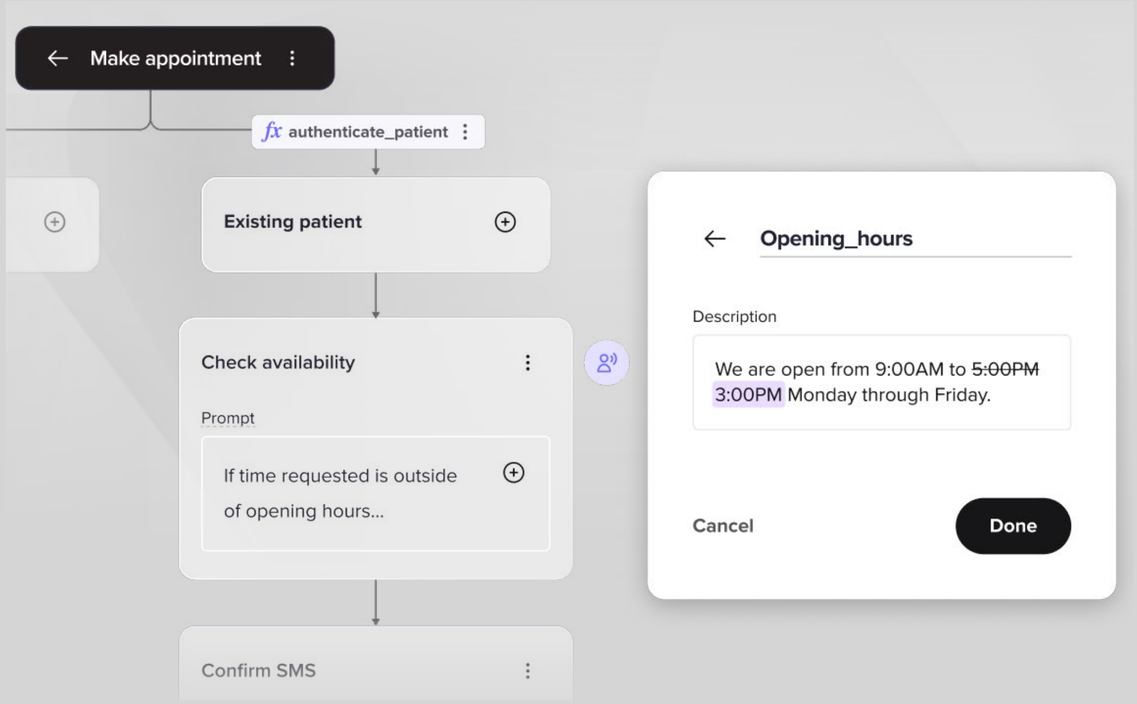

3 Building flows for key transactions

LLMs are smart enough that you don't need to design specific flows to answer FAQs. But transactions like taking bookings or sending payments need to follow a specific set of steps.

While you can write simple prompts that enable LLMs to hold a fairly normal conversation, it's safer and more reliable to design flows that show LLMs how to move through a conversation, including what actions to take at every step.

4 Building in checkpoints for specific tools

As a final safeguard, it's a good idea to build checkpoints for certain types of transactions that require specific tool use or API calls. These checkpoints remind the AI agent to ensure that certain actions have been taken based on call type, specific utterances, or certain types of transactions.

Agent-to-agent (A2A) collaboration enables real task execution

Most AI agents operate in isolation. They can handle a single task within a narrow context, but struggle with multi-step processes like pulling data, verifying identity, updating a record, and notifying the customer.

With A2A, agents can delegate work to other agents. For example, one agent might triage a request, another handles verification, and a third completes the task. Each brings a specific capability, but they function together as a system.

A2A lets you build more robust, coordinated workflows that reflect how real teams operate.

Multi-call prompting (MCP) gives agents the ability to think across steps

Most prompting approaches are: one prompt in, one response out. That's fine for basic tasks, but not for complex, multi-step tasks.

Multi-call prompting allows agents to reason iteratively. They can pause, gather new information, reassess, and revise their output before taking action, a lot like a human agent checking multiple systems before responding.

For example, a customer wants a multi-leg flight. The agent checks flights, pauses to gather seat, layover, and pricing information, asks clarifying questions, and then finalizes the booking with confirmation.

This approach makes agents more adaptable, more accurate, and better suited for real-time use cases.

Why both MCP and A2A matter

To scale agentic AI in production, you need both depth and coordination:

- MCP gives agents the reasoning power to handle ambiguity and multi-step logic
- A2A gives you a way to break work into parts and distribute it effectively

Together, they make AI agents more capable, more reliable, and more scalable. Take retail refunds, for example. Imagine you receive a damaged package and want to return or exchange it.

Agentic AI can handle the end-to-end process by asking the customer to describe the issue, sending them a secure link to upload photos or videos, and passing those to an AI agent that reviews and summarizes the damage. From there, the refund can be approved automatically without the customer waiting.

It's a faster, smarter way to handle common requests while reducing manual review.

Here it is in action:

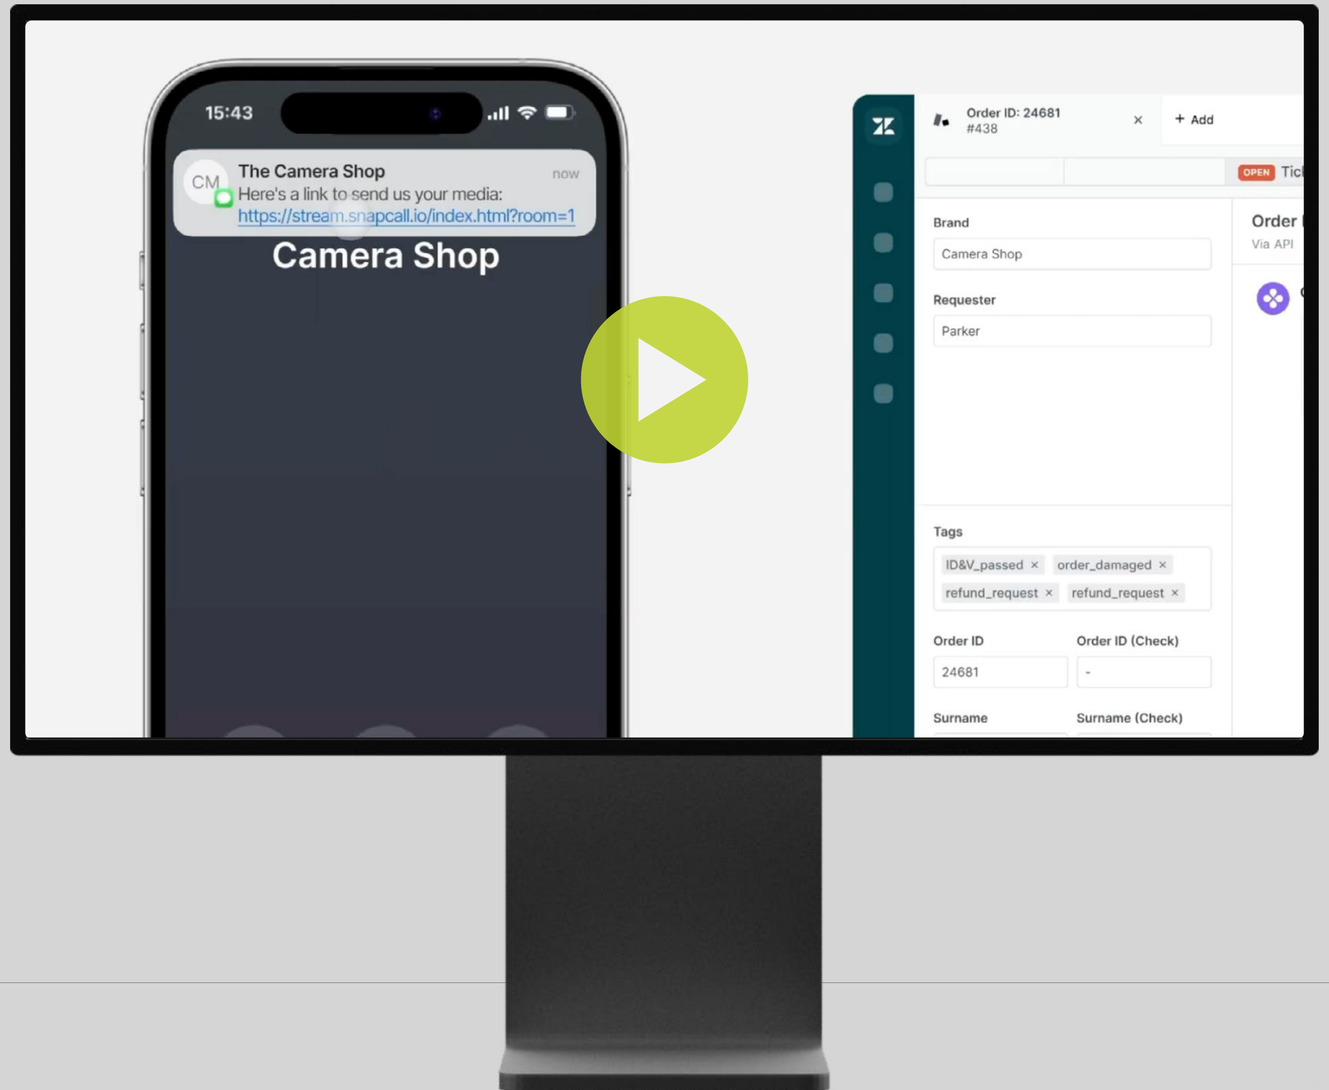

Building a foundation for agentic AI to scale

Scaling automation isn’t just about volume; it’s about doing it without breaking trust, processes, or customer experiences.

That starts with a solid platform. AI agents need to meet enterprise standards for uptime, compliance, and data privacy from day one. It also means deep integration with core systems like CRMs and call routing tools, so agents can access real-time data, personalize interactions, and support faster decisions.

Successful scaling starts with low-risk, high-impact use cases that can quickly demonstrate value. The contact center is a natural entry point. With its high volume of repetitive interactions, it offers clear, measurable opportunities for improvement without overhauling the entire business model.

Why the contact center is a strong starting point for AI transformation

Contact center operations have historically focused on lowering call volumes, optimizing agent utilization, and improving routing accuracy. Meanwhile, CX teams have worked to build loyalty and grow lifetime value.

These goals aren't in conflict. AI can help achieve both.

With AI agents, enterprises can reduce handle times, improve resolution rates, and lower costs while delivering a more enjoyable customer experience.

The future of customer service is agentic AI

Imagine a contact center where managers aren't stuck handing out scripted answers to every simple question. Instead, their primary work can focus on solving complex issues as they happen.

With AI agents in the contact center, your organization can reduce handle times, improve resolution rates, and lower costs while delivering a more enjoyable customer experience. One where customers have the freedom to talk naturally and receive humanlike responses that build trust and confidence that their issues will be resolved.

When AI is designed around real customer needs, it resolves issues without escalating unnecessarily, freeing up agents for high-value conversations. This balance of efficiency and experience is where modern CX begins.

Agentic AI can join existing contact center workflow ecosystems and breeze through repetitive tasks for agents. When AI is fully integrated, the contact center starts to function more like a command center, which means

- Information flows back and forth smoothly
- Teams have better visibility and control
- Agentic AI can take action for the agent, handling tasks behind the scenes to help connect the dots and reduce friction across systems.

Now that customers can and will engage with automated systems and enterprises can safely and reliably connect sources of customer data, AI agents can transform the way businesses communicate with their customers.

Drive relevant insights for the entire business

When contact centers can easily pull structured data from customer conversations, they can make better decisions for both agents and the business. For example, if a manager has access to clear, organized data, it's easier and faster to report on what's happening and why. Understanding the root of issues - like why demand isn't being met - can help shift the contact center from a cost center to a revenue driver.

Agentic AI doesn't just impact customer experience; it can also benefit your wider organization. If the C-Suite sees the quantifiable success of AI deployment within their contact center function, it drives adoption within other teams. Data from marketing, sales, invoicing, or shipping can then start feeding into the contact center, building a true 360-degree customer profile.

> With PolyAI's Analyst Agents, we can ask questions in plain language and get answers right away. This gives us a clear view of what's happening across our customer interactions and helps us make smarter, faster decisions to improve our service.
>
> 
>
> **Brian Jeppesen**
> Director, Contact Centers at Golden Nugget

Meet your agentic AI team: QA, Analyst, and Builder Agents

With new agentic AI roles, PolyAI's Agent Studio platform expands from front-line automation to unlock intelligence and growth.

PolyAI Agents continue to speak directly with customers, while QA Agents measure quality, Analyst Agents uncover insights, and Builder Agents keep the system evolving. At the core of the three new roles are PolyAI's agentic AI capabilities, which allow these agents to reason, adapt, and take action:

QA Agents

Automatically score every customer call on six quality factors, providing a continuous, standardized view of quality. These agents help CX leaders spot trends and target improvements faster.

Analyst Agents

Turn raw conversation data into clear answers, surfacing patterns and opportunities in plain language. These agents help CX leaders understand what's working and act on it quickly.

Builder Agents

Guide onboarding, development, and ongoing maintenance of PolyAI Agents. These agents help CX leaders put their best foot forward with customers and evolve without slowing operations.

Building trust and value for long-term AI success

Successful agentic AI implementation is about building trusted partnerships, working with experts who understand your business, and integrating AI thoughtfully into your organization. The organizations that win with agentic AI will not be the ones that adopt fastest. They will be the ones that prove value early, build trust deliberately, and scale with a clear strategy.

Ready to start your CX transformation?

PolyAI's Agent Studio is the world's only voice-first omnichannel platform for conversational AI.

Request a demo to find out more about how PolyAI can help you resolve up to 90% of calls, access first-party customer data, optimize phone support in an instant, and deliver effortless CX at scale.

Request a demo