

# Magic Quadrant for Conversational AI Platforms

13 August 2025 - ID G00826020 - 49 min read

By Gabriele Rigon, Justin Tung, [and 4 more](#)

The conversational AI platform market is evolving, with use cases expanding to include conversational AI agents and tools increasingly leveraging generative AI. This report guides application leaders in selecting conversational AI platforms for complex automation and multimodal interactions.

## Strategic Planning Assumption

By year-end 2027, conversational AI applications will automate approximately 70% of customer support interactions within enterprises, up from approximately 50% in 2025.

## Market Definition/Description

Gartner defines conversational AI platforms (CAIPs) as SaaS products that primarily enable the development of applications simulating human conversation across multiple channels and media. CAIPs leverage composite AI, including generative AI (GenAI) and natural language technologies. Conversations can use a mix of modalities such as text, voice and visual content. To support the building of conversational applications, platforms provide extensive coding options, from pro-code to no-code. Application areas include chatbots, virtual assistants (VAs) and conversational AI (CAI) agents.

CAIPs are used to create, deploy and manage AI-driven conversational interfaces. These platforms enable businesses to develop VAs and conversational AI Agents that facilitate both customer-facing and internal interactions through pro-code/low-code/no-code tools.

CAIPs empower businesses to centralize and democratize the development and management of multiple, diverse CAI initiatives, leading to more cohesive and efficient operations. The blend of capabilities provided by CAIPs is distinctive compared to those offered by other CAI solutions, such as targeted extensions for CAI found in other enterprise applications (e.g., CRM systems, contact center platforms) or stand-alone GenAI-native apps. In comparison, CAIPs are a better fit for strategic and scalable enterprise-grade CAI adoption.

Typical use cases for CAIPs are:

- **Customer interaction automation:**

- **Scope:** Focuses on enabling customers to interact with the brand across multiple channels and modalities, integrating with multiple back-end systems
- **Focus:** Empowers customers to resolve their inquiries and perform transactions without human intervention. Such experience requires the conversational application to understand, reason or triangulate, fetch information or transact
- **Example:** Self-service VAs for customer service and support

- **Employee assistance:**

- **Scope:** Supports both customer service representatives (CSRs) and, in general, employees across various job functions
- **Focus:** Improves productivity and decision making by providing real-time insights and automating routine tasks
- **Examples:** CSR assist solutions, HR or ITSM VAs

- **Sales and marketing automation:**

- **Scope:** Enhances sales and marketing efforts through goal-driven interactions and lead generation activities
- **Focus:** Increases conversions and campaign effectiveness by enhancing personalized and proactive customer engagement
- **Examples:** AI-driven sales assistants, proactive marketing chatbots

- **Enterprisewide conversational AI agents:**

- **Scope:** Emerging use cases where conversational AI agents operate autonomously or semiautonomously across enterprise applications, integrating with multiple systems
- **Focus:** Achieves goal-oriented tasks through proactive orchestration of data and systems
- **Examples:** Conversational AI agents for enterprise search, and CUI-driven process automation across business units

## Mandatory Features

- **Coding options:** CAIPs provide multiple coding options, which may include low-code, no-code, GenAI-assisted and pro-code. Low-code, no-code and GenAI-assisted tools support rapid application development, deployment, execution, testing and management using declarative, high-level programming abstractions via dedicated graphical UIs (GUIs) or GenAI-driven assistants. Pro-code tools enable the use of scripting and programming languages, such as Java or Python, to develop highly configured functionalities and integrations, giving technical users access to deep configuration and customization of the CAI applications.
- **Composite multilingual NLP:** Predefined and customizable pipelines of techniques that allow the processing of the user's natural language leveraging the multilayer composition of, for example, rule-based and machine learning techniques, including large language models (LLMs).
- **Workflow building:** GUIs for visual conversational or nonconversational workflow building, including tools to easily manage any relevant fulfillment logic during the dialogues. The resulting scripted workflows should be able to work in combination with, and back when appropriate, agentic AI planning and execution features, if those are available (see "Common Features").
- **Integration with back-end systems and data sources:** Tooling meant to set up or personalize communication with critical back-end systems, such as services (e.g., cloud services, AI frameworks), data sources (e.g., CRM, customer data platform) or applications (e.g., contact center as a service, martech, and analytics and business intelligence). It includes natural language query (NLQ) functionalities to interface with underlying BI systems in natural language.

- **Support for LLM prompt engineering:** Support for LLM prompt engineering via dedicated low-code/no-code modules or embedded functionalities. Prompt engineering is the discipline of providing inputs (e.g., in the form of text or images), to GenAI models to specify and confine the set of responses the model can produce. Retrieval-augmented generation (RAG) is a common composite prompt engineering architecture. It includes GenAI content-anomaly guardrails.
- **Analytics module:** Dedicated tools, including dashboards, to collect, monitor and analyze performance data to get meaningful and actionable insight for reporting, oversight and improvement purposes, which is key for the overall app life cycle management.
- **Data security and privacy controls:** Dedicated functionalities and tools to handle privacy, enterprise compliance and security aspects in the platform itself, as well as when building and deploying the CAI app and at runtime.

## Common Features

- **Voice interaction support:** Native or deeply embedded services leveraged to implement sophisticated voice experiences (e.g., listening and speaking) on either dedicated devices or using telephony. It may include voice biometrics.
- **Knowledge graph support:** Native tooling to build and maintain knowledge graphs (KGs). KGs are machine-readable representations of the physical and digital worlds. They include entities (people, companies, digital assets) and their relationships, which adhere to a graph data model — a network of nodes (vertices) and links (edges/arcs).
- **CAI app orchestration:** Dedicated tools that enable the orchestration of multiple CAI applications, either within a single use case or targeting multiple use cases. This functionality includes different collaboration patterns, ranging from scripted CAI app pass through to CAI apps escalating and routing to other CAI apps, or AI agents. It also includes the orchestration of language assets and dialogue assets across several CAI implementations.
- **Multichannel connectivity:** Prebuilt and customizable integrations to a variety of different digital channels, use specific rich features of different channels and operate multiple channels centrally. Channels may include messaging platforms, website chats, webhooks, telephony, voice speakers and others.

- **Custom integration with non-native LLMs services:** Tooling to enable custom integrations to platform-external LLMs, whether hosted by the deployer or provided by third parties.
- **Video interaction support:** Native or deeply embedded functionalities aimed to process video content in batch or in real time. It may include support for gesture controls in input and the deployment of digital-humans interfaces.
- **Support for numerical data processing:** Numerical-data-processing features that enable the CAI apps to interpret, analyze and respond to numerical inputs, facilitating data-driven interactions and insights.
- **Support for image upload:** Native functionalities that allow CAI app users to upload images, enabling the CAI app to analyze, interpret and respond to visual content effectively.
- **GenAI engineering:** GenAI engineering tools enable enterprises to operationalize models faster, balancing both governance and time to market. AI engineering tools can be subdivided into model-centric and data-centric tools. We consider DataOps, LLMOps, LangOps or FMOps, or more broader terms, such as ModelOps or MLOps, as a subset of AI engineering.
- **Vendor-specific LLMs:** The ability vendors may have to provide a self-hosted domain-specific LLM service, such as LLMs being built by the vendors in-house, or a fine-tuned version of open-source or third-party proprietary LLMs.
- **Agentic AI planning and execution:** Includes tooling for enabling autonomous process planning and execution in CAI applications. These features are emerging, but are not standardized in the broad market of CAIPs yet. Despite this, Gartner observes a clear trend toward them becoming more mainstream for CAIPs, and to be distinctively associated with the emergence of conversational AI agent applications. The processes automated by agentic AI planning and execution should be mostly planned and designed by the relevant engine, typically LLM-based, and autonomously optimized leveraging heuristics based on past experiences or user behavior. However, they may partially rely on scripted workflows (see “Workflow building” under “Mandatory Features”). CAI applications infused with such features are capable of doing the following:
  - Understand the user's goal, without an explicit prompt, and autonomously plan a process (a series of steps) that can lead to its fulfillment

- Execute the process step by step, taking the actions (e.g., querying underlying systems or data sources) required to fulfil the goal.
- **Continuous optimization tools:** Functionalities meant to facilitate CAI apps' maintenance and continuous optimization, to enable an implementation to easily improve over time as more data is collected.
- **Deployment options:** Functionalities that allow to install or grant access to the platform, as well as to deploy applications, on-premises, in the cloud or on the edge.

## Magic Quadrant

Figure 1: Magic Quadrant for Conversational AI Platforms





**Gartner.**

## Vendor Strengths and Cautions

### Avaamo

Avaamo is a Niche Player in this Magic Quadrant. It is a private company headquartered in Los Altos, California, U.S. Its conversational AI platform (CAIP) enables conversational AI (CAI) applications for customer experience (CX) and employee experience (EX) use cases. It is complemented by services such as Avaamo Ambient, which uses generative AI (GenAI) to document patient-clinician conversations, and AutoQA, which evaluates customer interactions using LLaMB, a proprietary framework for building GenAI apps. Avaamo's customer base spans North America, Europe and Asia, primarily in healthcare, but also in

banking, insurance and retail. Its roadmap includes advancing domain-specific AI agents and providing real-time insights into CAI usage and performance.

### *Strengths*

- **Market understanding:** Avaamo demonstrates a strong understanding of the buyer personas within targeted industries, enabling it to consistently focus on meeting the needs of the enterprises it serves. Additionally, it offers a solid framework for collaboration to engage customers in planning its roadmap and developing new features.
- **Cost of ownership clarity:** Avaamo's pricing model is less complex and provides more cost predictability compared to pricing models adopted by many vendors included in this Magic Quadrant.
- **Ease of use:** Avaamo's CAIP features a particularly clean, minimalistic and usable low-code/no-code user interface, making it easy for nontechnical users to get up to speed quickly while offering options for developers who require more advanced customization.

### *Cautions*

- **Vertical focus:** Avaamo's strong emphasis on the healthcare sector may limit its R&D bandwidth to develop innovative products beyond that domain and its ability to provide the same degree of comprehensive services and support for customers in other industries.
- **Product roadmap:** Since 2023, Avaamo has evolved its product strategy, putting more emphasis on specific industries and skills but, in general, introducing less differentiating features, compared to other vendors in this Magic Quadrant. This relatively targeted differentiation may pose a risk to its competitive positioning.
- **Geographic presence:** Although Avaamo has established a presence in the U.K. and expanded its partnership networks in EMEA, its sales and operations remain predominantly concentrated in North America, particularly in the U.S., which may limit its service capabilities outside the aforementioned markets.

## **Boost.ai**

Boost.ai is a Leader in this Magic Quadrant. Headquartered in Sandnes, Norway, it is a private company specializing in CAI solutions for regulated industries. Its CAIP enables secure multi-AI-agent interactions across chat and voice, and it works in combination with



products such as Agent Assist and AI-powered CX Insights. Boost.ai's customer base is primarily in Europe, North America and Oceania, and it comprises banking, insurance, government, telecommunications and technology clients. Its roadmap includes developing capabilities to orchestrate AI agents across voice and chat, as well as a marketplace to share reusable assets across instances and platforms.

### *Strengths*

- **Responsiveness to market changes:** Boost.ai demonstrates strong adaptability to evolving market demands through continuous investment in differentiating features. Examples include its enterprise-grade security, advanced collaboration tools, Test Studio, which features Persona-Based Testing, and recent focus on regulated industries and government use cases.
- **Customer experience:** Boost.ai's customer service is distinguished by rapid activation times and high service levels. The platform's usability, combined with its ability to support scalable and diverse deployment models, positions Boost.ai among the top vendors in this research for customer experience.
- **Sales strategy:** Boost.ai demonstrates a highly focused and transparent sales strategy, further strengthened by a robust presence on third-party marketplaces. It emphasizes providing prompt and customized responses to prospects during CAIP evaluations.

### *Cautions*

- **Research investment:** Boost.ai invests less in fundamental research compared to peers of similar size in this research, and it does not hold patents. This approach may limit its ability to establish thought leadership and maintain a competitive edge in innovation.
- **Talent strategy:** Boost.ai's employee base has experienced modest growth over the past five years, and its approach to talent acquisition and retention is less formalized and structured compared to leading vendors in this report. This may hinder its ability to scale.
- **Customer base diversity:** Boost.ai's customer base is less diverse compared to other vendors in this research, primarily concentrated in financial services and government. Organizations in other industries should ensure alignment with their industry-specific needs.

## **Cognigy**

Cognigy is a Leader in this Magic Quadrant. It is a private company headquartered in Düsseldorf, Germany. Its CAIP, Cognigy.AI, enables multimodal and multichannel use cases primarily related to contact center automation and CX. Complementary products include Cognigy Voice Gateway, to enable voice connectivity to any telephony systems, and Cognigy xApps, which are custom microweb applications that enable rich multimodal experiences. Cognigy's customer base is mostly in Europe and North America, across industries such as finance, insurance, retail, e-commerce and manufacturing. Roadmap items include 360-degree AI assistance for developers and a large language model (LLM)-driven scenario-based testing.

*NiCE announced its intention to acquire Cognigy on 28 July 2025. At the date of publication, Cognigy met the inclusion criteria for this Magic Quadrant and continued to operate as a going concern. Gartner will provide additional insight and research to clients as more detail becomes available.*

#### *Strengths*

- **Product capabilities:** Cognigy's CAIP has robust and flexible coding options, which significantly increase its adaptability to a wide range of developer profiles. Its multimodal capabilities and process management features also stand out for feature richness and usability compared to most other vendors in this research.
- **Market understanding:** Cognigy demonstrates a strong grasp of its target market, engaging not only contact center leaders but also C-suite AI decision makers. Its structured approach to customer involvement ensures that product development closely aligns with real user needs.
- **Customer base diversity:** Cognigy serves one of the most diverse customer bases among vendors evaluated in this report, generating significant revenue across multiple industries. This demonstrates the platform's industry-agnostic design and its ability to support a wide range of use cases.

#### *Cautions*

- **Business model flexibility:** Cognigy's growth objectives primarily focus on financial targets, placing less emphasis on visionary detail than some competitors. This may impact the vendor's ability to quickly adapt its offerings or business strategy in response to shifting market demands.

- **Research investment:** Cognigy invests less in pure research than peers of similar size included in this research, and it focuses its R&D on applying existing advancements rather than innovative research.
- **Strategy differentiation:** Cognigy's most recent major strategic shifts — moving to LLM orchestration in 2023 and agentic AI in 2024 — are less distinctive and more generic than the strategic moves made by most of the vendors covered in this study, which may limit its ability to stand out.

## DRUID AI

DRUID AI is a Challenger in this Magic Quadrant. It is a private company headquartered in New York, New York, U.S. Its CAIP combines GenAI capabilities with classic automation to orchestrate processes and enable composable CAI design via the DRUID Conductor. It integrates with services such as RPA (UiPath), analytics (for example, Power BI and Tableau) and voice/chat (for example, AudioCodes and Genesys), as well as Microsoft Copilot 365. DRUID AI's customer base spans Europe and North America, and it comprises industries such as banking, insurance, healthcare, higher education and retail. Roadmap items include enhancing platform multiagent orchestration with GenAI and improving voice AI capabilities, including real-time analytics.

### *Strengths*

- **Business agility:** DRUID AI's business model balances diverse revenue streams and has evolved to effectively incorporate several technological advancements. The company also adapted it in the last three years — for example, through geographic expansion, relocation to the U.S. and the launch of a Customer Excellence department.
- **Vertical partnership strategy:** DRUID AI has strategic partnerships with third-party products and consulting firms that allow it to penetrate various business units. Its marketplace offers industry-, use-case- or task-specific prebuilt conversational apps or modules, catering primarily to IT, healthcare and financial services while also supporting a broad range of other industries and use cases.
- **Customer experience:** DRUID AI scored high relative to the majority of the vendors in this Magic Quadrant for reported customer satisfaction. The vendor is distinguished by its service activation time.

### *Cautions*

- **Research investment:** Due to its primary focus on applied AI, DRUID AI scored lower compared to other vendors in this Magic Quadrant for CAI research spending, and it has fewer R&D staff than similarly sized peers. The company also lacks patents and academic contributions.
- **Regional focus:** DRUID AI's implementation and sales partners cover a more limited regional scope on average compared to the partners of most vendors in this Magic Quadrant, and DRUID AI's partners delivered fewer projects and won fewer deals. Expanding the geographic focus and the effectiveness of the partnership network is necessary for the vendor to grow beyond Europe, North America and the Middle East.
- **Voice automation tools:** DRUID AI utilizes third-party integrations for speech transcription and synthesis, rather than offering built-in, native capabilities. This approach may result in less control over the configuration of speech-to-text (STT) and text-to-speech (TTS) services in voicebots compared to other vendors in this research.

## Google

Google is a Leader in this Magic Quadrant. It is a public company headquartered in Mountain View, California, U.S. Its CAIP, Conversational Agents, enables multimodal CAI apps powered by the latest DeepMind innovation. The platform is part of Customer Engagement Suite with Google AI and works in combination with additional Google products, such as Vertex AI, Google's AI development platform. Google's customer base for CAI is global and includes a wide range of industries such as telecommunications, retail and e-commerce, financial services, government, and the public sector. Google's roadmap for CAI includes expanded evaluation capabilities and further enhancing multimodality capabilities for voice and video interactions.

### *Strengths*

- **AI research investment:** Google is at the forefront of innovation with its DeepMind initiative, a robust R&D organization and a strong portfolio of patents and academic contributions in natural language processing and AI.
- **Geographic presence:** Google has a global presence with offices in all regions. Its multilingual capabilities, and the way they are tailored to specific geographical locales, are among the most extensive in the sample of vendors included in this research. These capabilities are further supported by a well-distributed global partner network.

- **Market responsiveness:** Google has demonstrated exceptional market responsiveness. With the launches of LaMDA in 2021, Vertex AI in 2023 and Gemini LLMs in 2024, it has positioned itself at the forefront of the AI wave and as a provider of a unified AI stack.

### *Cautions*

- **Composability of some CAIP components:** Google's CAIP is less composable than other products we evaluated in the context of this research. For example, the Generative Playbooks depend more heavily on LLMs than other Q&A modules we assessed in this research, and the CAIP does not natively support knowledge graphs. This may reduce configuration flexibility in some Q&A use cases.
- **Overall sales strategy:** Google's sales strategy is globally defined, with less precision than the strategies outlined by other vendors, especially for target roles and tangible objectives, which may limit growth in verticals. The vendor has less sales strategy differentiation for industry verticals or geographical markets than other vendors in this report.
- **Product portfolio:** Google's AI and CAI offerings cover a broad range of interoperable products and services throughout the AI tech stack, such as Customer Engagement Suite, Agentspace, Vertex AI Agent Builder and Vertex AI Search. End users may find it difficult to navigate and assess different Google offerings.

## **IBM**

IBM is a Challenger in this Magic Quadrant. It is a public company headquartered in Armonk, New York, U.S. Its CAIP, watsonx Orchestrate, focuses on enhancing productivity and mostly targets employee-facing use cases with customizable prebuilt agents. Complementary products include watsonx.data, to manage large data repositories and indexes, and the Technology Expert Labs, which provide expertise in AI, automation and security. IBM's CAIP customer base is global and spans domains such as HR, sales, procurement, customer care and finance. Roadmap items include advanced speech-to-speech models and a continued focus on prebuilt AI agents for common domains.

### *Strengths*

- **Geographic presence:** IBM has a broad and deep international footprint with widespread office locations, a well-distributed customer base across multiple geographies and a robust global network of sales and service partners. It also features IBM Consulting. This

network facilitates consistent service delivery and easier access to IBM's expertise, regardless of region.

- **Research investment:** IBM received some of the most favorable evaluations in this research for its commitment to innovation and research. It consistently invests in producing a high volume of patents and academic contributions in the fields of natural language technology and AI.
- **Marketing execution:** IBM's marketing organization is one of the largest among vendors evaluated in this research. The company leverages its extensive developer ecosystem to maximize its market reach. It also effectively communicates its value proposition through a high volume of blogs and case studies, as well as through webinars and the TechXchange community.

### *Cautions*

- **Co-existing CAI offerings:** watsonx Orchestrate and watsonx Assistant currently operate in parallel, with watsonx Orchestrate including all the capabilities of watsonx Assistant. Overlapping features may cause confusion and uncertainty for buyers and developers.
- **Product maturity:** watsonx Orchestrate demonstrates a lower level of maturity as a CAI solution compared to other vendors' solutions evaluated in this report. For example, clients may find that the platform demonstrates early-stage channel integration and language coverage.
- **Product differentiation:** watsonx Orchestrate offers less differentiation as a CAI product than that of other vendors assessed in this report. End users may find fewer unique CAI features in watsonx Orchestrate, especially in relation to the enablement of specific use cases in contact center automation.

### **Kore.ai**

Kore.ai is a Leader in this Magic Quadrant. It is a private company headquartered in Orlando, Florida, U.S. Its CAIP, the Kore.ai Agent Platform, enables enterprises to build CAI and multiagent systems for service, work and process automation across different functions and business units. Additional offerings include AI for Service, to streamline omnichannel multimodal support in contact centers, and AI for Work, which covers enterprise search and employee-facing AI Assistant. Kore.ai's customer base is primarily in North America and

Europe and spans industries such as financial services, healthcare, retail and media. Roadmap items include features to enhance agent orchestration and governance.

### *Strengths*

- **Product strength:** Kore.ai delivers a feature-rich platform that stands out for its comprehensive and well-balanced capabilities. The platform excels particularly in GenAI enablement and process management, including tools and functionalities that support the development, deployment and management of conversational AI agents.
- **Business model adaptability:** Kore.ai operates a scalable business model focused on tiered licensing and usage fees with limited reliance on professional services. It has blended technology upgrades, partner enablement, leadership updates and operational restructuring to align with evolving business needs.
- **Employee development programs:** Kore.ai offers structured and extensive educational programs and learning resources for employees. Its commitment to continuous learning and innovation is reflected in its rapidly growing workforce, which has tripled since 2021.

### *Cautions*

- **Geographic reach:** Kore.ai's customer base is heavily concentrated in North America and Europe, with minimal presence elsewhere. Additionally, the regional coverage of Kore.ai's implementation partners is less extensive compared to that of other vendors evaluated in this research.
- **Product complexity:** The platform's extensive feature set and complexity may lead to a steeper learning curve — especially for citizen developers — and could require additional upskilling or reliance on partners to effectively navigate and utilize.
- **Marketing strategy:** Kore.ai's overall approach to communicating its differentiating features to the market is less clear, and the changes applied to its marketing strategy in the last 12 months are more generic or tactical, compared to those of other vendors included in this research. Clients may find it challenging to quickly grasp Kore.ai's unique strengths.

## **LivePerson**

LivePerson is a Niche Player in this Magic Quadrant. It is a public company headquartered in New York, New York, U.S. Its CAIP, Conversational Cloud, enables brands to manage AI-



driven engagements across digital and voice channels, mostly in customer service and support, and digital commerce. It offers capabilities such as Conversational Intelligence, used to extract insights from customer interactions, and Conversation Assist, which provides human agents with real-time guidance. LivePerson's customer base is primarily in North America and Europe, and in industries such as banking, retail, healthcare and telecommunications. Roadmap items include major upgrades to its Copilot product and enhanced support for third-party GenAI models.

### *Strengths*

- **Industry strategy:** LivePerson has a strong cross-industry presence with one of the most evenly distributed customer bases among vendors in this report. This balanced breakdown of clients and revenue highlights its ability to serve diverse industry needs effectively.
- **Operations improvement:** LivePerson has implemented significant operational enhancements since 2023, including the formation of specialized GenAI teams and the establishment of a GenAI center of excellence. These initiatives will likely improve efficiency, foster innovation and strengthen the delivery of LivePerson's solutions.
- **Strategy adaptability:** LivePerson has made several significant strategic pivots in the past two years, including expanding from a pure digital solution to a comprehensive end-to-end omnichannel offering and evolving its pricing model for greater simplicity and flexibility.

### *Cautions*

- **Product differentiation:** LivePerson's heavy strategic focus on agentic AI and GenAI has resulted in a less differentiated offering compared to that of other vendors in this report. Planned roadmap items for 2025, including Copilot 2.0 and enhancements to existing channels, do not introduce major innovations.
- **Marketing execution:** LivePerson has a comparatively small number of employees dedicated to supporting its CAIP and ran among the fewest marketing campaigns related to its CAIP in 2024 compared to those of other vendors of similar size in this report. As a result, it may face challenges in driving market awareness and customer acquisition.
- **Composable architecture:** LivePerson's CAIP offers limited deployment flexibility for specific modules — restricted to public and private cloud environments. It also lacks an



ontology-like natural language understanding component and does not provide built-in native machine translation capabilities.

## **Omilia**

Omilia is a Visionary in this Magic Quadrant. It is a private company headquartered in Athens, Greece. Omilia Cloud Platform (OCP) is a voice-first CAIP primarily targeting customer service use cases. Additional products in Omilia's suite include quality management (Workforce AI) and Contact Center Security through passive voice biometrics and behavioral signals. Omilia's customer base is mostly in North America and Europe, across industries such as financial services, insurance, utilities and food services (e.g., to automate order taking in drive-thrus). Roadmap items include the addition of new speech-to-speech models and a workflow assist solution to let human supervisors guide bots in real time.

### *Strengths*

- **Enterprise customer understanding:** Omilia's "unified AI" approach enables CAI applications to learn and improve from interactions in all channels. In addition, its strong focus on enterprise security and antifraud measures closely aligns with the priorities of enterprise customers in many specific industries.
- **Pricing model transparency:** Omilia's pricing model is more transparent and straightforward than that of most others evaluated in this research. Key pricing factors align with the unique consumption patterns of voice automation services; containment, latency and performance guarantees influence costs based on actual outcomes.
- **Research investment:** Omilia's strong R&D organization has achieved a high number of patents and academic contributions in AI and natural language technology (NLT) relative to its size, making it one of the top performers in research and innovation among vendors in this evaluation.

### *Cautions*

- **Geographic coverage:** Most Omilia offices are branch locations, and the size of its implementation partners' network is smaller compared to that of other vendors in this research. The scope of its language support, although granular in terms of relating to specific locales, is also more limited compared to that of other leading vendors included in this research.

- **Marketing Investment:** The number of targeted campaigns Omilia has run in 2024 was lower compared to most of the vendors in this research, as was its scope, which related primarily to financial services and utilities. Omilia's event participation and publications are among the lowest reported, which may limit its visibility to broader audiences.
- **Number of prebuilt connectors:** Omilia provides fewer prebuilt connectors to channels and back-end systems compared to the average number offered by other vendors in this Magic Quadrant. Although custom connectors can be built as needed, this runs the risk of incurring delays.

## PolyAI

PolyAI is a Niche Player in this Magic Quadrant. It is a private company headquartered in London, U.K. Its Agent Studio is a voice-first omnichannel CAIP that builds, manages and analyzes CAI applications — in particular, multimodal conversational AI agents for customer-facing use cases. The platform includes security features for regulated industries and comes with advanced conversation analytics features. PolyAI's customer base is primarily in North America and Europe, spanning industries such as travel, healthcare, financial services, retail and consumer services. Roadmap items include faster and more flexible AI agent creation, including a Supervisor Suite for enabling agentic workflows for analytics and continuous self-improvement, and improved, scalable voice AI.

### *Strengths*

- **Voice experiences:** PolyAI's capabilities for voice automation and voicebot use cases are more advanced than those of other vendors in this research. They provide natural-sounding voice experiences with low latency, as well as distinctive voice-specific testing tools.
- **Research investment:** PolyAI distinguishes itself through significant investments in pure research, delivering a disproportionately high volume of academic contributions compared to that of other vendors of similar size.
- **Customer retention:** Strong contract growth based on renewals and low customer churn demonstrate PolyAI's effective sales execution and its ability to foster customers' loyalty and trust in its product.

### *Cautions*

- **Consumption-based pricing:** PolyAI's pricing model is primarily based on conversation volume and professional services, although it is distinctive in including an outcome-based logic. This approach offers fewer options to optimize costs and can make long-term total cost of ownership less predictable compared to that of other vendors in this research.
- **Geographic coverage:** PolyAI has fewer offices compared to the average number of offices among the vendors included in this research. It also lacks implementation partners in Africa, the Middle East and Asia, and has no sales partners outside North America and Europe.
- **Channel partner network:** PolyAI's implementation delivered fewer projects, and its sales partners closed fewer deals in fewer regions than the average among vendors in this research. This limits its ability to expand into new markets and regions.

## SoundHound AI

SoundHound AI is a Visionary in this Magic Quadrant. It is a public company headquartered in Santa Clara, California, U.S. Its CAIP, Amelia 7, was integrated into its broader offering following the acquisition of Amelia in 2024. The current version comprises dedicated features to enable conversational AI agents and advanced voice capabilities. Amelia is used in industries such as banking, financial services and insurance (BFSI), healthcare, hospitality, utilities, and retail. Its customer base is primarily in North America and Europe but also in the Middle East and Latin America. Roadmap items include agentic AI enhancements for testing and analytics and a redesigned Amelia Voice Gateway to dynamically orchestrate multiple STT and TTS engines.

### *Strengths*

- **Offering and delivery strategy:** SoundHound AI has a broad and well-integrated CAI portfolio. Amelia is widely adopted as a white-labeled solution, serving as the foundation for third-party offerings such as CXone Mpower Autopilot, DXC Assure Platform and Fujitsu Next Generation Service Desk.
- **Market responsiveness:** SoundHound AI demonstrates a well-rounded and proactive approach to market trends, investing in technology upgrades and shifting from a one-size-fits-all model to vertical-specific solutions.
- **Voice automation capabilities:** SoundHound AI offers advanced voice technology, including proprietary solutions. By owning and operating key components of its voice

tech stack, SoundHound AI controls the quality and performance, supporting its focus on voice as a core strength in CAI.

### *Cautions*

- **Postacquisition status:** The 2024 acquisition of Amelia introduces potential uncertainty regarding the future direction of the product roadmap and technology integration. Customers and prospects should closely monitor developments to ensure their requirements for stability, service levels and long-term investment protection continue to be met.
- **Sales partnerships:** SoundHound AI's partner network currently focuses on implementation and lacks dedicated partners focused exclusively on sales activities. Clients may experience fewer localized sales touchpoints or a lack of tailored sales support.
- **Marketing activity and visibility:** SoundHound AI conducted fewer marketing campaigns and produced less thought leadership content in 2024 compared to similarly sized vendors, and it has a relatively small marketing team focused on CAIP. Clients may encounter fewer educational resources.

### **Sprinklr**

Sprinklr is a Niche Player in this Magic Quadrant. It is a public company headquartered in New York, New York, U.S. Its CAIP is part of the Unified-CXM (Customer Experience Management) platform and works in combination with Sprinklr's AI Agent Studio. Other products in Sprinklr's suite include Live Chat for web and mobile, and a Unified Agent Desktop to manage conversations from CAI and human agents. Sprinklr's customer base includes North America, Europe, the Middle East and Southeast Asia, with industries spanning finance, retail, technology, telecommunications and travel. Roadmap items include platform Copilots for multiple use cases and the enhancement of conversational AI agents' multimodal skills for voice and image interactions.

### *Strengths*

- **Marketing execution:** Sprinklr has a relatively large marketing team compared to other vendors in this report, with its marketing spend in line with tech industry standards. This likely contributes to increased product visibility and credibility.

- **Partner network:** Sprinklr has established a broad global network of sales and implementation partners with a good track record for successful project delivery. This partner ecosystem enables it to serve regions, such as Southeast Asia and Latin America, which are less commonly covered by competitors
- **Breadth of language support:** Sprinklr supports a wide range of languages with granular locale options. It offers native capabilities for both text and voice interactions, including STT and TTS.

### *Cautions*

- **Product architecture:** Sprinklr's CAI capabilities are provided via two UIs: Conversational AI and AI Agent Studio. Developers must select the interface that aligns best with their needs in relation to some GenAI and agentic AI features, whose sophistication varies between the two UIs. This distribution of features is unusual and may be detrimental to the coding experience.
- **Features pricing:** AI Agent Studio is only available in the highest tier of Sprinklr's Conversational AI offering, whereas most vendors in this report embed advanced GenAI features by default. Additionally, Sprinklr's pricing model is less transparent compared to other vendors included in this Magic Quadrant, making cost predictability more difficult.
- **Employee growth:** Sprinklr shows less employee growth in CAI compared to the other vendors included in this report, potentially impacting its ability to innovate, support customers and keep pace with evolving market needs in CAI.

### **Yellow.ai**

Yellow.ai is a Challenger in this Magic Quadrant. It is a private company headquartered in San Mateo, California, U.S. Its CAIP focuses on AI-powered automation to support omnichannel CX through chat, voice, SMS, email and social media. Yellow.ai's suite includes an AI agent builder, to deploy multichannel conversational AI agents, and Agentic Discovery, which infers trends and insights from historical tickets. Its customer base is mainly in Asia, with limited presence in North America and the Middle East, and it comprises industries such as BFSI, retail, consumer goods/services, hospitality and telecommunications. Roadmap items include enabling autobuilding omnichannel CAI applications in Agentic Discovery and a voice agent assist solution.

### *Strengths*

- **Marketing strategy:** Yellow.ai utilizes a well-balanced array of channels to distribute its marketing content and a mix of key performance indicators (KPIs) to assess the effectiveness of its marketing efforts. This approach has led to significant gains in brand awareness and large-scale customer acquisition.
- **Cross-industry feedback:** Yellow.ai's CAIP received high customer ratings for both overall usability and customer experience. This feedback is notable, given Yellow.ai's well-balanced customer distribution across industries compared to the more concentrated industry presence of that of most other vendors in this report.
- **Pricing model:** Yellow.ai's pricing model stands out for its clarity, predictability and flexibility compared to the approaches of other vendors in this research. It provides good visibility regarding annual costs and offers reasonable opportunities for negotiation.

### *Cautions*

- **Research investment:** Yellow.ai's R&D organization is comparatively limited in both budget and personnel, and its academic publication rate lags behind other vendors featured in this report. Additionally, the company has not secured any new patents in recent years.
- **Geographic strategy:** Yellow.ai has a limited presence in Europe and North America, with most of its customer base located in Asia. Prospects should carefully evaluate whether the vendor's geographic focus aligns with their own regional requirements.
- **Regulatory certifications:** Yellow.ai holds fewer certifications related to global regulations and industry standards compared to most vendors included in this research. Prospects should assess these factors when considering Yellow.ai for regulated environments.

## **Vendors Added and Dropped**

We review and adjust our inclusion criteria for Magic Quadrants as markets change. As a result of these adjustments, the mix of vendors in any Magic Quadrant may change over time. A vendor's appearance in a Magic Quadrant one year and not the next does not necessarily indicate that we have changed our opinion of that vendor. It may be a reflection of a change in the market and, therefore, changed evaluation criteria, or of a change of focus by that vendor.

## Added

As this is a new Magic Quadrant, no vendors were added.

## Dropped

As this is a new Magic Quadrant, no vendors were dropped.

# Inclusion and Exclusion Criteria

To qualify for inclusion, providers need to meet the following inclusion criteria:

- Single, stand-alone offering that meets the market definition. All the following conditions must apply:
  - The product must meet Gartner's market definition for conversational AI platforms (CAIPs). This includes offering all the mandatory features CAIPs are expected to provide as per the definition itself.
  - The product must be a single, unified offering. If the vendor has multiple product offerings that conform to the definition, only one CAIP will be considered. If the vendor cannot avoid more than one single offering to be evaluated for the aims of this Magic Quadrant, then this whole criterion is not met.
  - The product must constitute a stand-alone offering; that is, the product may be purchased and is marketed as a stand-alone platform. "Stand-alone" offering means that the CAIP may be purchased and is marketed as a stand-alone platform — and be a viable alternative when clients are looking for a best-of-breed conversational automation. It should not require the client to replace other aspects of, for example, their chat operations or contact center software. Thus, it should be able to work with third-party live chat, IVRs, ticketing systems, telephony systems, CRM platforms, ERP platforms, ITSM platforms, agent dashboards and RPAs. Up to five references of customers that purchased the CAIP as a stand-alone product are required.
- Tenure in the market. All the following conditions must apply:
  - The company must have been established prior to 1 January 2021.

- The company must have been actively operating in the market of conversational AI platforms (CAIPs) with a generally available CAIP prior to 1 January 2024.
- Customer base size, growth and weight. All the following conditions must apply — by “customers,” we imply “customers that purchased a license to use the CAIP, not other products in the vendor’s suite:”
  - A minimum of 70 customers at the date of submission
  - A minimum of 20 new customers in 2024
  - At least 35 active licensed customers that are using the CAIP in a production environment and spend a minimum of \$100,000 annually.
- Geographic spread. All the following conditions must apply:
  - No more than 75% of the customer base is in a single region, with regions defined as North America (NA), Latin America (LATAM), Europe, the Middle East and Africa (EMEA), and Asia/Pacific (APAC).
  - The company must have a significant geographic presence in at least two regions. “Significant” applies to both regions, meaning the company must have either headquarters (HQs) or, for example, multiple full-function (FF) offices, multiple branch offices, multiple resellers, or an R&A center, not a single full-function, branch office or reseller.
    - HQs: The primary location where a company’s executive management and key managerial and support staff are located. This is usually the central hub for strategic decision making and corporate governance.
    - Full-function: A location where the company operates with a complete set of functions, such as sales, marketing, customer service and possibly production or R&D. This is akin to a regional office or subsidiary that operates semi-independently.
    - Branch office: A smaller, local office that handles sales, support or other functions but is not as comprehensive as a full-function office.
    - Reseller: A third-party company (a sales partner) or individual that sells the company’s products or services, typically without being owned by the company. Resellers often have a formal agreement with the company to distribute its products.



- R&D center: A specialized facility where research and development activities are conducted, focusing on innovation and product development.
- Employee base. All the following conditions must apply:
  - The company must have a minimum 150 employees working on the CAIP practice or business unit across all roles and functions.
  - The company must have a minimum of 25% of the employees in product development/R&D.
- Use case focus. A minimum of 33% of the customers licensed to use the CAIP must be using the CAIP for use cases in customer service and support.
- CAIP features. All the following conditions must apply:
  - The CAIP must provide the features, which are mandatory in Gartner's definition of the market, as per the following thresholds:
    - Coding options: Must include at least low code and no code.
    - Analytics module: General availability (GA) in 1Q25 — over 50% of the customers using it in a production (prod) environment.
    - Composite multilingual natural language processing (NLP): GA in 1Q25 — over 50% of the customers using it in a prod environment.
    - Workflow building: GA in 1Q25 — over 50% of the customers using it in a prod environment.
    - Data security and privacy controls: GA in 1Q25 — over 50% of the customers using it in a prod environment.
    - Integration with back-end systems and data sources: GA in 1Q25 — over 50% of the customers using it in a prod environment.
    - Support for LLM prompt engineering: GA in 1Q25 — independently of the number of customers using it in a prod environment.
  - All the following features, which are common in Gartner's definition of the market, must be generally available in 1Q25:
    - Agentic AI planning and execution

- CAI app orchestration
- Continuous optimization tools
- Multichannel connectivity
- Voice interaction support

## Honorable Mentions

**Inbenta:** Headquartered in Allen, Texas, U.S., Inbenta offers a CAIP built on a proprietary composite AI framework, combining symbolic reasoning, NLP and GenAI. Inbenta's platform supports use cases for customer and employee experiences, with a strong focus on accuracy and extensibility, and its product suite includes solutions for search and knowledge management. Inbenta is not included in this Magic Quadrant because its agentic AI planning and execution features were not generally available in 1Q25.

**Microsoft:** Based in Redmond, Washington, U.S., Microsoft has a large product portfolio. Its CAIP, Copilot Studio, was launched in November 2023, and it incorporates many of the features offered by its previous CAIP named Power Virtual Agents. The solution enables organizations to build conversational AI assistants and agents integrated with Microsoft 365 apps and the Power Platform. Copilot Studio is not included in this research because, at the time of this writing, Gartner was unable to determine whether 33% of the licensed customers are using the CAIP for use cases in customer service and support. Gartner also could not confirm whether 50% of licensed customers are leveraging the analytics module, composite multilingual NLP and workflow building features.

**Parloa:** Headquartered in Berlin, Germany, but with a growing presence in North America, Parloa provides the AI Agent Management Platform. The solution offers features to build and deploy CAI applications, plus an agent assist module that features real-time translation for customer service representatives. Parloa was excluded due to its limited geographic presence, with the vast majority of its customer base located in Europe.

**Uniphore:** Uniphore is based in Palo Alto, California, U.S. Its Business AI Cloud was recently unveiled and now delivers a set of AI-powered capabilities for various business units — including contact centers. The AI-powered capabilities include Self-Service Agent, to build and deploy CAI applications; Real-time Guidance Agent; Conversation Insight Agent; and Communication Recording Agent. The overall vendor's suite spans modules for marketing, sales, HR, and general AI-agent-building capabilities. Uniphore was excluded because, in

1Q25, its agentic AI planning and execution and CAI app orchestration features were not generally available, and in 1Q25, fewer than 50% of its customers were using its composite multilingual NLP features in a production environment.

## Evaluation Criteria

### **Ability to Execute**

In evaluating providers' Ability to Execute within the CAIP market, Gartner emphasizes the factors that most directly impact a vendor's capacity to deliver robust, scalable and high-quality conversational solutions that meet evolving enterprise needs. The evaluation criteria are designed to capture not only the technical strengths of a platform, but also the organizational, operational and customer-facing aspects that collectively determine success for both vendors and their clients.

Within the CAIP market, product capabilities are paramount. Organizations seek platforms that offer advanced features, flexibility and enterprise-grade reliability — attributes that are challenging to deliver and differentiate in a highly competitive landscape. As such, the robustness and scalability of the core product are given the greatest weight in the assessment.

Criteria such as overall viability, sales execution/pricing and marketing execution are assigned moderate weightings relative to the platform itself. While these factors are important for understanding a vendor's long-term stability and market presence, they tend to play a secondary role for buyers. For instance, financial health and pricing flexibility are relevant considerations, but they are evaluated in conjunction with, rather than above, the ability to deliver a differentiated product.

Other key evaluation criteria include market responsiveness and track record, which are essential, given the rapid pace of innovation and shifting user expectations in conversational AI. Customer experience also serves as an important differentiator, encompassing not only the usability and performance of the platform, but also the quality of vendor support, onboarding and ongoing partnership. Operations excellence — including leadership quality and organizational maturity — further underpins a vendor's ability to scale and sustain successful deployments.

**Table 1: Ability to Execute Evaluation Criteria**

<i>Evaluation Criteria</i>	<i>Weighting</i>
Product or Service	High
Overall Viability	Medium
Sales Execution/Pricing	Medium
Market Responsiveness/Record	Medium
Marketing Execution	Medium
Customer Experience	Medium
Operations	Medium

Source: Gartner (August 2025)

## Completeness of Vision

In assessing the Completeness of Vision of providers in the CAIP market, Gartner focuses on the extent to which vendors demonstrate a deep understanding of customer needs, anticipate future trends and articulate a compelling strategy for long-term differentiation. Vision is not simply about having an ambitious roadmap; it is about aligning product development, go-to-market strategies and organizational resources to deliver sustained value in a rapidly evolving landscape.

For this market, the most heavily weighted criteria are market understanding, offering (product) strategy, vertical/industry strategy and innovation. These elements are critical because they reflect a vendor’s ability to anticipate and respond to the complex, shifting requirements of enterprise customers. Providers that excel in these areas are those that not only listen to their customers, but also shape the direction of the market through

differentiated capabilities, forward-looking product strategies and deep industry expertise. Innovation, in particular, is essential in conversational AI, where advances in natural language technologies and AI can quickly redefine competitive dynamics.

Criteria such as marketing strategy, sales strategy, business model and geographic strategy are assigned medium weightings. While these factors are important for scaling and sustaining growth, they are less directly tied to a vendor’s ability to set the pace of change or deliver breakthrough solutions. For example, a strong marketing or sales strategy can accelerate adoption but does not necessarily equate to visionary leadership in product or technology. Some criteria, such as geographic strategy, are less emphasized because the market is not yet globally mature, and regional differentiation is not a primary driver of customer decision making. Similarly, business model flexibility is important, but in a market where technical differentiation is paramount, it is not the leading indicator of visionary leadership.

**Table 2: Completeness of Vision Evaluation Criteria**

<i>Evaluation Criteria</i>	<i>Weighting</i>
Market Understanding	High
Marketing Strategy	Medium
Sales Strategy	Medium
Offering (Product) Strategy	High
Business Model	Medium
Vertical/Industry Strategy	High
Innovation	High
Geographic Strategy	Medium

## Quadrant Descriptions

### Leaders

Leaders demonstrate a strong balance between effective execution and a well-defined vision for the future of conversational AI (CAI). They provide robust, feature-rich platforms that are trusted by a wide range of customers and have established a significant presence in the market. Leaders consistently exhibit high levels of capability across product development, marketing and sales, which enables them to drive widespread market acceptance and adoption. Their ongoing investment in innovation ensures that their offerings remain relevant as technology and customer needs evolve. Leaders are also recognized for their agility in responding to changing requirements and their influence on the overall direction and advancement of the CAIP market.

### Challengers

Challengers perform well in the market and have established a solid customer base with dependable CAIPs. They demonstrate operational strength and financial stability, and their platforms are proven in real-world deployments. However, Challengers may focus more on incremental improvements and established practices rather than driving innovation or shaping new market directions. While they are effective at meeting current customer needs, they may not be as proactive in introducing new features or setting future trends as Leaders or Visionaries.

### Visionaries

Visionaries stand out for their innovative ideas and differentiated features that anticipate future needs in CAI. These providers often introduce new approaches or advanced capabilities that set them apart from competitors. Visionaries have a good understanding of emerging customer requirements and are recognized for their forward-thinking product strategies. However, they may not yet have achieved the same level of market presence, operational maturity or execution as Leaders or Challengers.

### Niche Players

Niche Players focus on specific segments, industries, or use cases within CAI, offering valuable solutions tailored to particular needs. They may excel in specialized areas or provide strong capabilities for targeted scenarios, but generally have a narrower scope compared to Leaders or Visionaries. Niche Players may lack the breadth of features, market reach or long-term vision needed to compete broadly across the market. While they deliver value for select customers, their overall impact and influence on the direction of the market are more limited.

## Context

As the CAIP market rapidly evolves with GenAI and agentic AI, buyers must take a strategic approach to maximize value and minimize risk.

To lay a foundation for CAIP success, organizations should start by investing in CAI design and CAI-ready data to future-proof deployments. Furthermore, they should clearly define use cases, customization and scalability needs to match the right solution — whether CAIPs, GenAI-native apps or targeted extensions.

Platforms with robust governance, compliance and security features should be prioritized, especially for regulated industries or when leveraging GenAI, as they allow to establish strong guardrails for privacy and response validation. Organizations should assess each platform's ability to scale, support multiagent ecosystems and automate complex workflows across the enterprise. Clear vendor differentiation and transparent roadmaps for agentic AI features are additional key factors to consider when evaluating vendors. It is then crucial to seek for additional competencies such as voice biometrics, content generation or process transparency to ensure the platform meets both current and future business needs in specific use cases that may be pursued.

This comprehensive approach will help organizations deploy scalable, compliant and innovative conversational AI solutions.

## Market Overview

The conversational AI platform (CAIP) market has undergone significant transformation over the past two years, driven by the rapid advancement and adoption of generative AI (GenAI) and, more recently, the emergence of AI agents. As a result, contemporary CAIPs are

increasingly hybrid in nature, integrating GenAI, agentic AI, natural language technologies and multimodal capabilities to address a broader spectrum of enterprise needs.

However, the market continues to face notable challenges. Buyer maturity remains relatively low, solution differentiation is unclear and more nuanced, with many providers struggling to differentiate their solutions, and the competitive landscape is crowded with both GenAI-native applications and CAI-focused extensions in a number of enterprise applications (see [Generative AI Brings Opportunity and Risks to the Conversational AI Market](#)). These challenges complicate vendor positioning and customer decision making. In response, many vendors are repositioning as AI agent platforms providers and rearchitecting their product portfolios to incorporate generative and agentic features, marking a transition toward agentic AI.

CAIPs are one possible way to build AI agents, as illustrated in [Innovation Insight for the AI Agent Platform Landscape](#). GenAI has expanded the boundaries of conversational technology, enhancing traditional virtual assistants and enabling conversational AI agents. In [Answering the Top Trending AI Agent Questions, Part 1: Foundational Knowledge](#), Gartner outlines the difference between AI assistants and agents when observed from the perspective of the agentic AI continuum. In essence, AI assistants can integrate with more autonomous systems, and they do not typically engage in self-directed actions, whereas AI agents do. Furthermore, AI assistants are primarily designed to be reliant on human feedback and interaction. AI agents, on the other hand, do not necessarily need to have a user interface or depend on human interaction, although those that CAIPs allow to build tend to be primarily “conversational” and automate human-led processes (for the different process types that AI agents are intended to support, see [Innovation Insight: AI Agents](#)).

Despite ongoing disruption, CAIPs maintain distinct value, particularly in scenarios requiring high levels of accuracy, scalability to handle large volumes of conversations, regulatory compliance and governance — areas where vertical GenAI applications may not suffice. CAIPs were among the first market categories to embed GenAI and early large language models (LLMs) into their offerings, laying the groundwork for today’s agentic platforms. These platforms are differentiated by their ability to support multiagent ecosystems, orchestrate a wide range of use cases, and balance business unit outcomes with enterprisewide orchestration and scalability.

Strategically, CAIPs enable organizations to deploy conversational AI solutions at scale, from straightforward chatbots to complex, process-driven virtual assistants and conversational AI



agents. They support both customer-facing and internal workflows, empowering enterprises — particularly in regulated sectors such as financial services and healthcare — to automate large-scale interactions and mission-critical processes while maintaining high standards for customer experience, compliance and service excellence.

Vertical GenAI applications can deliver autonomy for specific, task-oriented AI agents. CAIPs blend the best of deterministic (for example, encoding business rules in the form of workflows) and probabilistic AI (non-generative machine learning [ML] models and LLM) capabilities to address different use cases and customer requirements. CAIPs are then uniquely positioned to support multiple business processes and larger user bases, including customers. This broader exposure necessitates a greater emphasis on governance, compliance and orchestration, reinforcing the strategic importance of CAIPs in the evolving AI landscape.

## Acronym Key and Glossary Terms

CAI	conversational AI
CAIP	conversational AI Platform
CX	customer experience
EX	employee experience
LLM	large language model
ML	machine learning
NLP	natural language processing
NLT	natural language technology
STT	speech-to-text

TTS	text-to-speech
-----	----------------

⊕ Evidence

⊕ Evaluation Criteria Definitions

© 2025 Gartner, Inc. and/or its affiliates. All rights reserved. Gartner is a registered trademark of Gartner, Inc. and its affiliates. This publication may not be reproduced or distributed in any form without Gartner's prior written permission. It consists of the opinions of Gartner's research organization, which should not be construed as statements of fact. While the information contained in this publication has been obtained from sources believed to be reliable, Gartner disclaims all warranties as to the accuracy, completeness or adequacy of such information. Although Gartner research may address legal and financial issues, Gartner does not provide legal or investment advice and its research should not be construed or used as such. Your access and use of this publication are governed by [Gartner's Usage Policy](#). Gartner prides itself on its reputation for independence and objectivity. Its research is produced independently by its research organization without input or influence from any third party. For further information, see "[Guiding Principles on Independence and Objectivity](#)." Gartner research may not be used as input into or for the training or development of generative artificial intelligence, machine learning, algorithms, software, or related technologies.



Get The App

