Magic Quadrant for Conversational AI Platforms

7 July 2026 - ID G00841552 - 59 min read

By Gabriele Rigon, Justin Tung, and 3 more

The conversational AI platform market is rapidly evolving with new tech trends around multimodality and agentic AI, as well as changing governance needs and M&A activity. This Magic Quadrant helps application leaders select platforms for enabling complex and scalable AI assistants and conversational AI agents.

Market Definition/Description

Gartner defines conversational AI platforms (CAIPs) as platforms primarily used for developing applications simulating human conversation across multiple channels and on a mix of modalities such as text, voice and visual content. CAIPs leverage a composition of AI techniques, including classic natural language processing (NLP), and generative AI (GenAI) and agentic AI architectures. To support the building of conversational applications, CAIPs principally provide low-code and no-code coding options. Application areas include AI assistants and conversational AI agents.

Conversational AI platforms are designed to address the increasing demand for organizations to efficiently build, deploy and manage AI-driven conversational systems at scale, addressing the requirements of both employee experience and customer experience use cases. While they may offer some predefined AI assistants or AI agents that can be modified, the primary focus of a CAIP is to equip organizations with tools for building customized conversational AI applications.

CAIPs typically embed specialized, dedicated and feature-rich tooling for language-specific NLP and multimodal interactions, as well as conversational flow building and analytics. By offering a unified environment that supports low-code and no-code

development — and, in some cases, extends to pro-code and GenAI-assisted options — CAIPs empower technically savvy business users, including citizen developers, to create and orchestrate both customer-facing and internal AI assistants and conversational agents. Unlike AI engineering environments, which primarily serve AI engineers, software developers and data scientists, CAIPs are purpose-built for broad enterprise adoption, enabling strategic, scalable and organizationwide conversational AI initiatives.

The centralized approach of CAIPs eliminates the fragmentation seen when conversational AI is embedded as targeted extensions within other enterprise applications or delivered as isolated GenAI-native solutions. CAIPs democratize conversational AI development, allowing multiple teams to collaborate and manage diverse initiatives from a single platform, and the organization to govern such activities and deliverables. This leads to more cohesive and efficient operations, as well as the agility to respond to evolving business requirements. Tangible outputs include streamlined deployment of AI-powered, customer experience-focused applications — such as text-based conversational AI agents and voice AI agents — and enhanced employee productivity through intelligent assistants. Additionally, some CAIPs may be able to facilitate the orchestration and integration of multiple conversational applications, whether built natively or with third-party technologies, ensuring seamless user experiences across channels and business functions.

Mandatory Features

- **Multichannel connectivity:** Prebuilt and customizable integrations to a variety of different digital channels that use specific experience-rich features of different channels and operate multiple channels centrally. Channels may include messaging platforms, website chats, webhooks, telephony and smart speakers.

- **Conversational flow building:** Visual tools that enable the design and orchestration of conversations, combining deterministic logic with agentic AI planning and execution, which typically depends on large language model (LLM)-based components guided by instructional prompts provided during the coding phase. Flows embed logic related to specific user journeys or workflow process tasks, including the selection of tools needed to achieve goals or execute actions autonomously. Conversational flow-building tools support drag-and-drop configuration of nodes and their connections. Flows can be drafted using GenAI-assisted coding options.

- **Tool integration:** Tooling meant to set up or personalize integration and communication with any kind of application or system required to fulfill specific actions or goals. It includes Model Context Protocol (MCP) and API-based approaches. Tool integrations can connect the conversational AI application to 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 business intelligence (BI) systems in natural language.

- **RAG enablement:** Tooling to enable retrieval-augmented generation (RAG). CAIPs must, at minimum, provide tools suitable for implementing basic (or "naive") RAG by retrieving relevant information from a single data source using simple keyword- or vector-based search. Retrieved content is directly supplied to the language model for response generation, without advanced reranking, multi-step retrieval or iterative refinement.

- **AI guardrails:** Controls and mechanisms that enforce governance, security, and mitigation of AI-specific risks natively within conversational AI applications. AI guardrails address critical risks that pertain to GenAI models and/or AI-driven components such as prompt injection, data leakage, bias, hallucinations and unauthorized access. They include features like prompt injection mitigation, content moderation, identity and access management, PII/PHI protection, ethics and bias controls, hallucination reduction, and validation of RAG-generated content.

- **Analytics module:** Dedicated tools, including dashboards, to collect, monitor and analyze conversational interactions to get meaningful and actionable insight for reporting, oversight and improvement purposes, which is key for overall app life cycle management.

- **QA, observability and optimization tools:** Functionalities meant to facilitate conversational AI apps' quality assurance (QA) while building and when releasing new versions of the same, and tools ease maintenance and continuous optimization and enable an implementation to easily improve over time as more data is collected. QA and optimization tools include, among others, "preview" and A/B testing features, as well as advanced testing features that may be provided by specialized third parties, or other solutions in the vendor's product suite.

- **Baseline data security and privacy controls:** Functionalities and tools to handle privacy, enterprise compliance and security aspects in the platform itself, as well as when building and deploying conversational AI applications and at runtime. They include features such as encryption of communications in transit and at rest, PII redaction and/or anonymization, and functionalities that provide adherence to industry standards and regulations.

Optional Features

- **Pro-code coding options:** 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.
- **GenAI-assisted coding options:** GenAI-assisted coding options may include automatic generation of training examples, GenAI-assisted dialogue-flow drafting, automatic creation of FAQs and code generation.
- **Composite NLP:** Composite orchestration of AI techniques to enable the processing of natural language inputs through the integrated use of heterogeneous tools and methods, including rule-based systems, traditional machine learning models and LLMs. Leveraging dynamic routing strategies, the feature allows users to intelligently select, sequence and combine the most suitable AI tools based on input characteristics, task requirements and contextual factors such as language variety or interaction modality (voice or text).
- **Voice interaction support:** Native or deeply embedded functionalities that can be leveraged to implement sophisticated voice experiences (e.g., listening and speaking) on either dedicated devices or using telephony. It may include voice biometrics.
- **Video interaction support:** Native or deeply embedded functionalities aimed to process video content in batches or in real time. It may include support for gesture controls in input and the deployment of digital-human interfaces.
- **Image processing support:** Native or deeply embedded functionalities that allow CAI app users to upload images, enabling the CAI app to analyze, interpret and respond to visual content effectively.

- **Digital human enablement:** Tools for creating and deploying lifelike digital avatars that visually represent AI assistants and conversational AI agents in conversational applications. This feature includes 3D avatar design, animation tools for synchronizing gestures and facial expressions with speech, and support for real-time emotional and behavioral responses. Digital humans enable rich, interactive user experiences by rendering AI-driven agents as visually expressive, humanlike personas within digital environments.
- **Knowledge base management**: A centralized toolset for organizing, indexing, and maintaining enterprise knowledge sources to support conversational AI applications. Knowledge base management tools facilitate real-time retrieval of grounded responses by leveraging advanced search and RAG.
- **Advanced, modular and/or specialized RAG enablement:** Tooling to sophisticated RAG architectures with modular components for query processing, retrieval, reranking and context optimization. Enables advanced features such as multihop retrieval, feedback loops, summarization, knowledge graphs or agentic orchestration, delivering higher accuracy and adaptability for complex conversational tasks.
- **Prompt management:** Tools for continuous designing, organizing, versioning, and optimizing prompts within conversational AI platforms. Prompt management tools enable centralized governance of prompt libraries, support collaborative prompt engineering workflows and provide analytics for continuous monitoring of prompt performance and effectiveness. These features ensure consistent, high-quality AI outputs, facilitate compliance and security, and help mitigate risks such as prompt injection and unintended model behaviors.
- **Extended AI-agent communication protocols:** Standardized specifications that enable AI agents to securely discover, authenticate, exchange context and coordinate actions with external systems, tools and other agents. These protocols support both agent-to-system and agent-to-agent interactions, covering capabilities such as discovery, identity verification, authorization, context sharing and secure invocation of resources. Optimized for LLM and agent workflows, they facilitate tool access, multiagent collaboration and orchestration while addressing interoperability, security and governance requirements. Common protocols include MCP and Agent2Agent (A2A).

- **Native DSLMs:** Native domain-specific language models (DSLMs) that are pretrained or fine-tuned on data from a particular industry, discipline or functional domain to optimize performance for specialized tasks and terminology. DSLMs deliver higher accuracy, relevance and contextual understanding in areas such as healthcare, finance, legal or customer service, supporting advanced use cases like expert Q&A, compliance and workflow automation by leveraging domain-tailored knowledge and linguistic patterns. Being native, they ensure independence from third-party providers of models, thereby enhancing control over privacy and security aspects of the model deployment.
- **Diverse 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.
- **FinOps and cost-optimization tools:** Tools designed to maximize the efficiency and the cost-effectiveness of conversational AI deployments. These include counters to track LLM-token consumption and FinOps practices for real-time monitoring, flexible billing models, granular analytics, forecasting of AI workloads, and managing budget controls and alerts, enabling organizations to optimize cloud and application usage.
- **AI assistant orchestration:** Tooling for coordinating and managing multiple AI assistants to deliver unified, context-aware user experiences. Orchestration tools dynamically assign, sequence and delegate tasks across specialized conversational AI applications, whether those were built and deployed using in-CAIP native tooling or third-party tools. Orchestration tools also enable context and data sharing, and resolve conflicts or overlaps in responsibilities.

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. Its Avaamo Conversational AI Platform offers capabilities to build, deploy and govern agentic AI applications in enterprise environments. The product suite includes AutoQA, which evaluates customer interactions across channels to boost agent performance; Agent Copilot, to guide human agents through the call journey; and Live Agent for MyChart, to equip contact centers with electronic health records (EHRs) for better patient experiences. Avaamo's customer base is mainly in North America and Europe; it is concentrated mostly in healthcare, but extending into financial services, IT,

education and digital commerce. Roadmap items include a multiagent orchestration validator for complex multiagent routing, cross-AI agent governance, and global context management.

Strengths

- **Product usability:** Avaamo offers a particularly clean, minimalistic and user-friendly UI compared with competing platforms. Its CAIP utilizes thoughtfully designed builder features providing an intuitive interface that streamlines the agent coding experience.
- **Marketing strategy:** Avaamo executes a significantly stronger and more clearly differentiated marketing strategy than most other vendors evaluated in this report. By utilizing highly relevant channels and strategic approaches, it is well positioned to generate greater awareness and engagement compared with competitors.
- **Geographic strategy:** Avaamo provides broader native support for text and speech across languages than most competitors. This is reinforced by a distributed worldwide office presence and a well-distributed network of sales partners across multiple regions. Buyers are likely to benefit from satisfactory multilingual capabilities and reliable local support in diverse regions.

Cautions

- **Industry strategy:** Avaamo's customer base is less diverse than other vendors evaluated, with half its clients heavily concentrated in healthcare, but few clients in other sectors. Buyers in industries that are not in the core area of expertise of Avaamo may experience limited sector-specific capabilities and support compared with vendors with broader market penetration.
- **Market responsiveness:** Avaamo has made fewer significant strategic business and product pivots over the past two years compared with other peers we evaluated in this research, and has therefore adapted less flexibly and less promptly to disruptive market forces and opportunities.
- **AI trust, risk and security management (TRiSM) and governance:** Avaamo implements less granular, multilayered and agent-specific guardrails compared with top vendors in this report. Additionally, the platform meets fewer regulatory standards and security certifications than leading market competitors currently achieve.

Boost.ai

Boost.ai is a Challenger in this Magic Quadrant. It is a private company located in Sandnes, Norway. Its Boost.ai conversational AI platform focuses on orchestrating conversations across chat, voice and agent assist, specifically targeting regulated industries. The product suite includes Boost.ai Human Chat for real-time support; Agent Assist, to support human agents; and AI-Powered CX Insights for interaction analysis. Boost.ai's customer base is mostly in Europe and North America; it includes industries such as financial services, government, telco, IT and travel. Its roadmap includes enabling capabilities for "self-testing," to automate QA and identify jail-breaking attempts; "self-analyzing," to spot problems in conversational data and determine how to solve them; and MCP tools to build and improve AI agents via a built-in companion or third-party LLM clients.

Strengths

- **Customer experience:** Boost.ai ranks among the top vendors in this report for professional services, distinguished by rapid activation times and service levels. It has the highest percentage of clients that use a CAIP to enable multiple use cases in their organizations among evaluated vendors.
- **Market acumen:** The vendor demonstrates exceptional alignment with current enterprise customer requirements. Compared with most competitors included in this research, Boost.ai better addresses demands for cost optimization, pro-code capabilities and analytics within its platform.
- **Product strategy:** Boost.ai's value proposition stands out for its focus and effectiveness among those evaluated in this report. It is highly pragmatic, prioritizing real-world delivery and outcomes specifically tailored to the rigorous demands of highly regulated industries.

Cautions

- **Vertical industry strategy:** Boost.ai's customer base is less diverse than other vendors in this research. Over half of its clients are concentrated within financial services and other highly regulated industries. End users in other sectors are likely to experience comparatively less tailored support or industry-specific expertise.
- **Financial flexibility:** Boost.ai's customer base growth was more limited in 2025 than most competitors included in this report. Its latest funding round was smaller compared with other investor-backed vendors, and it has not pursued any

acquisitions in the past two years. This may affect Boost.ai's ability to quickly respond to market changes or scale its operations to meet evolving client demands.

- **Innovation:** Boost.ai makes relatively small investments in fundamental research compared with peers of similar size. Moreover, its R&D processes are significantly less driven by direct customer feedback compared with other vendors in this report.

Druid AI

Druid AI is a Niche Player in this Magic Quadrant. It is a private company located in New York City, New York. Its CAIP, Druid, focuses on orchestrating multiagent workflows, scripted flows and GenAI within a unified platform. Druid AI's product and service suite includes Conductor, an orchestration engine for goal-driven multiagent orchestration; Agentic Framework 2.0, a multiagent collaboration layer supporting skills, workflows and goal-based orchestration; and specialized professional services. Its customer base is mostly in EMEA and North America; it includes industries such as financial services, healthcare, government, retail and IT. Druid AI's roadmap comprises industry-specific Druid Voice, including a proprietary contact center connector with native CCaaS integrations; industry-specific prebuilt and authoring AI agents, featuring vertical agent templates; and multiagent orchestration via Conductor 2.0.

Strengths

- **Marketing strategy:** Druid AI executes a significantly clearer marketing strategy than peers in this research, by delivering highly relevant messaging that directly addresses enterprise pain points. Prospective buyers are then likely to gain better visibility into Druid AI's capabilities and value proposition, supporting more informed purchasing decisions.

- **Business model:** Druid AI maintains a diversified and well-distributed revenue growth across licensing, consulting and usage streams. Its strategic adaptations, including North American relocation, provide a distinctly stronger foundation than those of some other competitors included in this report.

- **Analytics:** Druid AI features a comprehensive, in-depth and feature-rich analytics module in its platform. Its detailed dashboards provide a good client-outcome-oriented view of key performance indicators (KPIs), outpacing the reporting capabilities of other evaluated vendors.

Cautions

- **Research and development (R&D):** Druid AI invests less in R&D compared with evaluated peers. The vendor is launching a formal IP & Patent Program, but over the last two years, it has produced few academic papers and intentionally avoided patents, which is likely to delay its capability to establish thought leadership in the CAIP market.
- **Product usability:** Druid AI's platform is less usable for nontechnical developers than competing solutions. For example, the orchestration-related features remain significantly pro-code, requiring users to manually configure APIs and webhooks instead of primarily utilizing GenAI-assisted approaches like leading competitors.
- **Product differentiation:** Druid AI's core platform features, including RAG and orchestration capabilities, are less differentiated than those of other vendors evaluated. This may make it difficult for the vendor to clearly establish its unique market value compared with leading vendors in this Magic Quadrant.

Google

Google is a Leader in this Magic Quadrant. It is a public company located in Mountain View, California. Its CAIP, Customer Experience Agent Studio (CX Agent Studio), focuses on native multimodal capabilities and multiagent orchestration for complex conversational enterprise tasks in customer service and support. CX Agent Studio is part of Gemini Enterprise for Customer Experience, which includes Agent Assist, a real-time, step-by-step AI companion for human representatives; Conversational Insights, which analyzes agent usage patterns and business outcomes; and Agent Search. Google's overall customer base is mostly in North America, EMEA and the APAC region; it includes industries such as retail, telecommunications, financial services, government and the automotive sector. Roadmap items focus on tooling to enable building of AI agents via agentic IDE to minimize manual configuration, alongside visual console usability enhancements for multiteam collaboration via Git-style version control and configuration management.

Strengths

- **AI innovation:** Google's product innovation capabilities are differentiated by Google DeepMind, which drives Google's AI advancements in LLMs and multimodality, and impacts Google's thought leadership and vision. Prospective buyers seeking access to the latest AI technologies will find Google's offerings aligned with their requirements.

- **Complete tech stack:** Google provides a comprehensive, end-to-end technology stack unlike many other vendors in this Magic Quadrant. By operating its own cloud services and utilizing proprietary chips like Tensor Processing Units (TPUs), Google is in the position to directly control performance, reliability, cost and scalability, thus potentially reducing deployment complexity and improving efficiency for clients.
- **AI operations breadth:** Google possesses a significantly larger organizational capacity than most other vendors evaluated in this research. It also offers among the widest range of educational programs for its employees, thus increasing the likelihood that customers and partners can access skilled support and guidance.

Cautions

- **Deployment options:** Gemini Enterprise for CX runs exclusively on Google Cloud infrastructure and is unavailable for isolated, air-gapped deployment, or for being deployed on competitors' clouds. On-premises is offered selectively for specific out-of-the-box deployments, such as Automotive Agents. Buyers in highly regulated industries should validate if these deployment options are compatible with their requirements.
- **CAI portfolio complexity:** Conversational Agents, which was the vendor's flagship CAIP in mid-2025, has evolved into CX Agent Studio. In Google's current CAI suite, CX Agent Studio coexists with Dialogflow CX to support deterministic state-machine workflows. Product pivots — which are more frequent than most other vendors' in this Magic Quadrant — are likely to make it difficult for end users to select and evaluate desired capabilities.
- **Analytics interfaces:** CX Agent Studio's core runtime monitoring and performance analytics module is less advanced and less usable than those provided by other vendors featured in this research. It features a basic dashboard that forces users to access a separate interface (Conversational Insights) for detailed business outcome analytics and advanced business intelligence.

IBM

IBM is a Visionary in this Magic Quadrant. It is a public company located in Armonk, New York. Its CAIP, watsonx Orchestrate, focuses on providing a central platform to build, deploy, manage and govern conversational AI agents. Additional products and services

in IBM's suite include watsonx.data, a hybrid data lakehouse for clients' scaled-up multiformat data repositories and analytics; watsonx.governance, a platform to support responsible, transparent and explainable AI; and the Technology Expert Labs, which provide expertise in AI, automation and security. IBM's overall customer base is primarily in North America and EMEA, although a relevant portion resides in APAC and LATAM. The CAIP is mainly used by industries such as healthcare, manufacturing, financial services, government and IT. Roadmap items comprise agentic control plane capabilities, unified policy layers across AI agents, and new GenAI-automated testing capabilities.

Strengths

- **Operations:** IBM possesses a significantly larger developer employee base compared with most competitors, supported by comprehensive talent acquisition and retention programs that ensure sustained technical firepower and far greater operational stability across deployments.
- **Geographic strategy:** IBM delivers wider global presence and broader language support than most peers, boasting a highly distributed customer base and extensive penetration across North American, LATAM, EMEA and APAC markets. Buyers located in regions that are underserved by other vendors are likely to benefit from IBM's coverage.
- **Scalability:** IBM offers advanced methods for model fine-tuning and optimization, along with more robust AI workload scaling strategies than other peers evaluated in this Magic Quadrant. It also allows reduction of inference costs by including all native LLM models within the SaaS entitlement without imposing additional per-token usage charges.

Cautions

- **Activation times of high-complexity deployments:** For advanced enterprise use cases, IBM demonstrates an overall service activation time ranging from one to four months, depending on the complexity of the enterprise use case. These average activation times are longer than those projected by some other vendors included in this Magic Quadrant for complex enterprise deployments.
- **Business model:** IBM's product strategy is comparatively less integrated than that of some of its competitors. In some scenarios, customers may have to acquire additional portfolio components (e.g., IBM Bob, an AI coding tool, and watsonx.governance) to

access additional features that other vendors included in this report offer within a unified platform.

- **AI TRiSM and guardrails:** watsonx Orchestrate offers less sophisticated native guardrails compared with other platforms evaluated in this report. The launch of the Agentic Control Plane within watsonx Orchestrate was announced only in May 2026. Some more robust risk and compliance features require accessing IBM watsonx.governance.

Kore.ai

Kore.ai is a Leader in this Magic Quadrant. It is a private company headquartered in San Mateo, California. Its CAIP, the Kore.ai Agent Platform, focuses on customer service, employee productivity and process automation. The product suite includes the Agent Management Platform, for managing and governing AI agents built on Kore.ai and other platforms; AI for Work, an employee productivity suite with enterprise search; and AI for Service, a suite of features to serve a range of CX use cases such as support, loan management and appointments. The company's customer base is mostly in North America, EMEA and APAC; it includes industries such as financial services, healthcare, manufacturing, retail and travel. Kore.ai's roadmap includes the development of a centralized governance console with adaptive guardrails; real-time voice coaching overlays; and predictive analytics to forecast AI agent performance degradation before it impacts users.

Strengths

- **Product innovation:** Kore.ai maintains a higher level of R&D staffing compared with most peers evaluated in this report, and its OS for AI Agents concept is both visionary and ambitious. Distinctive builder tools such as Arch and Agent Blueprint Language (ABL) further set Kore.ai apart.

- **Service and delivery capabilities:** Kore.ai's SLA agreements are more closely aligned with enterprise requirements than those of other vendors included in this research. Its North American and European implementation and service partners are numerous and well distributed across regions, and its professional services are structured and comprehensive.

- **Operation scalability:** Kore.ai's substantial organizational size and operational capacity make it one of the largest vendors in this Magic Quadrant. The company supports

employee growth and education through well-developed programs, systems and development practices, which increases the likelihood that clients will receive consistent service delivery and have access to skilled personnel.

Cautions

- **Product complexity:** Despite extensive feature richness, Kore.ai's complex platform ranks lower in overall usability than competing solutions featured in this Magic Quadrant. The high volume of configuration options and sophisticated features are likely to create a longer and steeper learning curve than alternatives.
- **Marketing execution:** Kore.ai's marketing quality and messaging differentiation appear weaker than other evaluated vendors. This strategic shortfall is likely to limit the vendor's ability to clearly showcase its highly differentiated value proposition against competitors, whose marketing execution is more effective.
- **Sales strategy:** Kore.ai utilizes a more resource-intensive, region-specific sales expansion approach through its partners than other competitors. This approach is likely to slow down customers' adoption, particularly in regions where partner enablement and alignment are still maturing.

Netomi

Netomi is a Challenger in this Magic Quadrant. It is a private company headquartered in San Mateo, California. Its conversational AI platform, the Netomi CAIP, focuses on resolving customer issues by operating as an orchestration layer across enterprise systems. The product suite includes Conversation Viewer for interaction analysis; agent assist for human agent support; and a Knowledge Management Hub to manage the content that powers AI responses. Netomi's customer base is largely in North America, with some presence in EMEA; it includes industries such as travel and hospitality, media, financial services, utilities and retail. Roadmap items are autolearning AI agents, which improve through "observing" human agent behavior in live workflows; adaptive experiences via generated UI (e.g., cards, maps, flows); and speech-to-speech runtime for streaming speech-native turns with orchestration, fallback and handoff.

Strengths

- **Sales execution:** Netomi's customer retention rates are consistently higher than most competitors, supported by a transparent, value-driven pricing model and ongoing

expansion into more complex enterprise deployments, which drives long-term client engagement.

- **Market responsiveness:** Netomi demonstrates greater market responsiveness than many peers, having rapidly adapted its business model and product architecture to address evolving customer needs (e.g., Governed Agentic OS), expanding into new use-case domains (i.e., from CX-only to EX) and investing in building a wider partner network.
- **Marketing strategy:** Netomi's marketing strategy is more clearly differentiated and more directly addresses current AI market trends than the marketing strategies of most vendors in this report. This enables the company to more effectively communicate its unique value proposition and positioning to enterprise buyers.

Cautions

- **Delivery strategy:** Netomi's platform is less frequently adopted as a private label solution than those of leading competitors, so clients seeking to offer Netomi's capabilities under their own name are likely to find fewer options or less flexibility. Its approach to enabling service partners is less structured than that of other vendors included in this research.
- **Geographic spread of customer base:** Netomi's geographic presence is less extensive than most competitors, with the majority of its customers concentrated in North America and limited presence in EMEA, APAC and LATAM, potentially restricting its appeal for global enterprises.
- **Multimodality:** Netomi relies entirely on third-party technology for speech transcription and synthesis, and offers fewer prebuilt connectors to key channels. Buyers with advanced multimodal needs should verify whether the vendor is able to meet their specific requirements.

NiCE Cognigy

NiCE Cognigy is a Visionary in this Magic Quadrant. It is a public company, post acquisition of Cognigy by NiCE in 2025. NiCE's headquarters are located in Hoboken, New Jersey. Its CAIP focuses on deploying conversational AI agents for customer support and CX. The platform includes Cognigy Voice Gateway, to enable voice connectivity to any telephony system; Simulator, providing automated agent evaluation with prebuilt

scenario tests; and Knowledge AI, an end-to-end RAG pipeline. Its customer base is mostly in North America and Europe, spanning industries such as financial services, telco, travel and hospitality, manufacturing, and utilities. Planned roadmap items include an embedded AI assist coding assistant for end-to-end AI agent life cycle management, an open agent ecosystem, and agent observability to aid agent debugging.

Strengths

- **Service level and deployment diversity:** NiCE Cognigy demonstrates higher service levels and deployment effectiveness across a broader range of channels and use-case domains than most competitors in this Magic Quadrant. Customers report using the CAIP to enable multiple complex use cases seamlessly across diverse environments.
- **Vertical and industry expertise:** NiCE Cognigy features a significantly broader and more widely distributed customer base across industries than other vendors evaluated in this report. This large footprint makes the platform highly industry-agnostic and a strong fit for specific vertical enablement.
- **Product innovation:** NiCE Cognigy invests a higher percentage of its employee base in product development and R&D than peers of similar size in this market. Backed by recent patents granted in 2025, the vendor exhibits a highly structured, customer-led R&D roadmap that outpaces the innovation processes of many competitors.

Cautions

- **Post acquisition product strategy:** NiCE Cognigy presents a higher risk of roadmap deviations and R&D reductions compared with other competitors. Buyers must perform rigorous technical evaluations to identify any potential discontinuity in the integration of the broader NiCE suite and NiCE CXone.
- **Market responsiveness:** NiCE Cognigy has projected less agility and market responsiveness over the last 12 months than its competitors. Its market differentiation remains largely unchanged post acquisition, relying heavily on a broader positioning as a holistic, agentic CX platform.
- **Post acquisition operations:** NiCE Cognigy's operations are more vulnerable to acquisition-related impacts than those of other vendors in this Magic Quadrant. After acquisitions, buyers should always assess increased risk of service discontinuity, requiring strict monitoring of missed SLAs and unannounced leadership or workforce changes.

Omilia

Omilia is a Visionary in this Magic Quadrant. It is a private company located in Cyprus. Its CAIP, Omilia Cloud Platform, focuses on automating customer service across voice and digital channels using automated approaches to monitor interactions and continuously improve AI agents. The product suite includes CSR CoPilot, for real-time human agent assistance; TalkGuard, for contact center antifraud; and Workforce AI, for automated analysis of every call on all channels in contact centers. Omilia's customer base is mostly in North America and Europe; it includes industries such as financial services, telecommunications, healthcare, education and food services. Roadmap items include contact center digital twins for training simulations; desktop intelligence, to record human agent screens and conversations; and autonomous self-learning agent optimization.

Strengths

- **Risk management and governance:** Omilia offers some of the most comprehensive and structured guardrails for risk management in this report, in particular for privacy protection and cybersecurity, compared with other vendors evaluated. This capability is further enhanced by the complementary antifraud product, TalkGuard, and broad compliance coverage.
- **R&D:** Omilia's R&D spending is higher than most other companies in this report. The vendor's dedication to research and its numerous publications in academic journals keep it at the forefront of research compared with other competitors in this report, positioning its clients to benefit from ongoing innovation and state-of-the-art technology.
- **Market understanding:** Omilia addresses current needs with distinctive features like automated agent building from call clustering. Its roadmap stands out for its dedication to "self-learning agents," demonstrating a stronger visionary focus than most other vendors in this report.

Cautions

- **Geographic strategy:** Omilia's geographic footprint is more limited compared with other vendors in this report. The majority of its customer base is concentrated in North America and the remainder located almost entirely in Europe, alongside an emerging presence in APAC. Buyers in other geographies should confirm the availability and level of support for Omilia's services within their specific region.

- **Diversity of deployments:** Omilia's conversational AI platform is utilized to enable a much narrower diversity of deployment types compared with other competitors in this Magic Quadrant. Its deployments are almost exclusively concentrated in customer service and CX automation, predominantly operating on voice and telephony channels.
- **Financial flexibility:** Omilia's customer base is more limited compared with other vendors of similar or smaller size in this report. Its growth has been entirely organic, with no acquisitions in the past two years, and its VC funding is lower relative to competitors. Buyers should consider that these factors may constrain Omilia's ability to rapidly adapt to market shifts.

PolyAI

PolyAI is a Niche Player in this Magic Quadrant. It is a private company headquartered in London, U.K., and San Francisco, California. Its CAIP for Enterprise CX focuses on autonomous orchestration of end-to-end customer journeys across omnichannel and multimodal digital touchpoints. The product suite includes Supervisor Suite, an AI-powered solution that integrates analytics and human agents to evaluate conversations; Owl, a proprietary speech-transcription service for contact centers; and Agent Development Kit (ADK), a developer-facing building solution. PolyAI's customer base is mostly in North America and Europe; it includes travel and hospitality, retail, healthcare, financial services and government. Roadmap items are Agentic Builder, for natural-language-driven creation of AI agents in-browser; PolyAgents Suite for autonomous analytics; and Raven v4 Omni, a proprietary speech-to-speech LLM.

Strengths

- **Offering and product strategy:** Built natively on proprietary voice technology and language models for multimodal CX, PolyAI offers a clearer vision of its current product than most peers. Its evolution toward an open platform demonstrates one of the most cohesive strategic directions seen among the vendors participating in this report.
- **Marketing execution:** PolyAI invests a higher percentage of its budget and workforce in marketing than most vendors evaluated in this report. By utilizing diverse channels like digital content and alliance referrals, the vendor communicates its differentiated value robustly and extensively.

- **Usability for builders:** PolyAI outpaces many other vendors in this report in coding flexibility, from pro-code to GenAI-assisted coding options. Advanced capabilities such as “automated onboarding scrapers” make the platform more appealing to users of varying technical expertise than other solutions in this report.

Cautions

- **Geographic strategy:** PolyAI’s customer footprint is primarily concentrated in North America and Europe, and the vendor maintains a more limited network of implementation partners across these regions compared with most evaluated peers. This strong regional focus makes PolyAI a less established choice for organizations with operations in other parts of the world.
- **Niche voice AI expertise:** PolyAI’s platform supports fewer channels via prebuilt connectors and demonstrates lower deployment diversity across channels than most other vendor solutions in this Magic Quadrant. This is primarily due to its focus on deployments for messaging platforms and voice-based telephony. PolyAI’s main expertise remains heavily linked to its voice AI capabilities, requiring more work to differentiate its product, service and pricing compared with competitors.
- **Limited deployment options:** Unlike other vendors in this report, PolyAI’s core platform components are restricted to cloud deployments. Although virtual private cloud deployments are allowed for enterprise contracts, the lack of other deployment options limits its ability to support stringent enterprise use cases requiring fully air-gapped or on-premises environments.

Salesforce

Salesforce is a Leader in this Magic Quadrant. It is a public company located in San Francisco, California. Its CAIP, Agentforce 360, focuses on unifying enterprise context, work and engagement by integrating conversational AI agents with native CRM data and workflows. The product suite is available through Salesforce Foundations, which provides access to other capabilities, such as Agentforce Builder, Agentforce Studio and Data 360. Salesforce’s customer base is mostly in North America, EMEA and LATAM; it includes industries such as retail and consumer goods, financial services, travel and hospitality, IT, and healthcare. Roadmap items include extending its voice capabilities with support for outbound interactions and additional channels and languages; Agentforce Contact

Center global availability; and enhancements to Slackbot, the prebuilt AI agent embedded in Slack.

In June 2026, Salesforce signed a definitive agreement to acquire Fin, a "customer agent platform" vendor focused on automating omnichannel and multimodal customer interactions.

Strengths

- **Pricing:** Salesforce offers a differentiated and transparent tiered pricing model compared with most competitors. It provides a scalable, universal consumption currency with transparent estimation tools (e.g., a free provisioning SKU via Salesforce Foundations to enable frictionless POCs, and Agentforce Flex Credits).
- **Industry strategy:** Salesforce demonstrates a stronger cross-industry presence than many peers included in this report, underpinned by a portfolio of over 300 prebuilt AI agent templates. Purpose-built solutions such as Agentforce Financial Services, Agentforce Health, and Missionforce for the public sector allow for significantly faster enterprise deployments in specific verticals.
- **Marketing strategy:** Salesforce strongly positions Agentforce 360 as the platform for the "Agentic Enterprise," unifying data, work and agency better than most vendors. Events like Dreamforce and robust content marketing create wide market visibility. This provides buyers with strong visibility into product capabilities and future direction.

Cautions

- **Multimodality:** Salesforce relies exclusively on third-party models for speech recognition and synthesis, offering less native customizability than leading competitors. Moreover, speech-to-speech support remains on the roadmap, putting Salesforce behind other vendors that already utilize it broadly in the market.
- **Talent growth:** Salesforce's initiatives for AI talent acquisition, employee retention and continuous learning are less structured and granular than those of leading vendors in this research. This could hinder its capacity to scale its talent pool to meet evolving enterprise requirements compared with other peers in this Magic Quadrant.
- **Business model:** Salesforce positions its commercial model to eliminate isolated, point-solution purchasing for CAIP add-ons. Pathways to purchase CAI capabilities —

such as Agentforce 360, Agentforce Trust Layer, Data 360 and MuleSoft — include Salesforce Foundations and the Agentic Enterprise License Agreement (AELA). This model is not typical of most other CAIP vendors included in this report. Buyers should ensure they fully understand its implications and the capabilities delivered by the different tiers.

SoundHound AI

SoundHound AI is a Leader in this Magic Quadrant. It is a public company located in Santa Clara, California. In the CAIP market, SoundHound AI acquired Amelia and Interactions over the last two years and announced plans in April 2026 to acquire LivePerson. The CAIP featured in this Magic Quadrant, Amelia 7, is a voice-native solution that combines governed orchestration, proprietary speech technology and an LLM-agnostic architecture. The product suite includes Polaris, a proprietary speech-recognition technology leveraged throughout the vendor's portfolio; Smart Answering, for handling inbound calls; and Dynamic Drive-Thru, a voice-ordering system. SoundHound AI's customers are mostly in North America and Europe, primarily in financial services as well as healthcare, telco, travel and hospitality, and retail. Roadmap items include an AI testing and simulation suite, AI Builds AI, which is intended to enable a self-building and self-improving agent loop, and a unified analytics platform.

Strengths

- **Extended product suite:** SoundHound AI offers a broader and more flexible product suite than most vendors evaluated. Expanding through strategic acquisitions, it natively incorporates advanced conversational analytics, robust voice-related tools and distinctive human-in-the-loop orchestration capabilities.
- **Responsiveness and adaptability:** SoundHound AI has rapidly integrated capabilities from recent major vendor acquisitions. Additionally, it has adapted by evolving its business model to a partner-first revenue architecture and unifying its marketing under a single masterbrand strategy.
- **Multimodality and voice experiences:** SoundHound AI's voice capabilities are more robust compared with other vendors in this research. It offers a fully tunable, native automatic speech recognition (ASR) engine for speech-to-text (STT) and text-to-speech (TTS), real-time voice-to-voice models in production, and advanced support for digital humans, while retaining third-party flexibility.

Cautions

- **Geographic customer base:** SoundHound AI's customer base remains predominantly concentrated in North America compared with the more globally distributed footprints of other vendors in this report. Its expansion into other regions, despite maintaining a relevant presence, is proceeding at a slower pace than its peers.
- **Sales strategy:** SoundHound AI's sales partner network is less established among global enterprise clients compared with other vendors in this research. Additionally, its reliance on inorganic growth via acquisitions is likely to present adoption risks for enterprise customers.
- **Sales effectiveness:** SoundHound AI exhibits lower overall sales effectiveness than other vendors evaluated in this report. Its customer base size and customer acquisition growth rate lag behind those of competing vendors of similar or smaller size, which poses a risk to its long-term revenue growth.

Sprinklr

Sprinklr is a Niche Player in this Magic Quadrant. It is a public company located in New York City, New York. Its CAIP, Sprinklr AI Agent, focuses on enterprise workflows to deliver autonomous and AI-assisted experiences across omnichannel touchpoints. The product suite includes Quality Management and Conversational Analytics for AI agent conversations; AI+ Studio, for LLM management and governance; and a knowledge base module, for facilitating information management to enable RAG. Sprinklr's customer base is mostly in EMEA and North America; it includes financial services, retail, IT, telecom, and travel and hospitality. Roadmap items are features to support video processing natively; a product recommendation engine with support for commerce protocols (e.g., UCP, ACP); and a memory management system, for brands to configure and manage AI agent memory for advanced controls.

Strengths

- **Research:** Sprinklr outpaces most similarly sized peers in pure research spending and relevant academic publications. Its R&D process is well-structured and customer-led, ensuring a progressive validation of new features that keeps its innovation more closely aligned with market needs than other competitors.
- **Operations:** Sprinklr's conversational AI practice is backed by one of the largest bases of employees working in the CAI practice, among all vendors in this Magic Quadrant.

It distinguishes itself from competitors through comprehensive, highly effective educational programs and progression systems targeting employee growth, ensuring talent retention and scaling.

- **AI agent building:** Sprinklr offers a more logical and intuitive builder framework than most market competitors. Its AI agent orchestration logic, packaged as a router, provides universal instructions that customers can customize at will to effectively handle ingress, internal tasks, and external AI agents via APIs or MCP.

Cautions

- **Marketing strategy:** Sprinklr's overall marketing strategy, specifically in the areas of relevancy, content differentiation and channel utilization, is not as strong as other vendors evaluated in this research. This is likely to restrict its ability to effectively communicate its value to buyers.
- **Diversity of deployments:** Sprinklr exhibits less diversity in its deployments across channels compared with competitors, relying predominantly on website chat UIs and messaging interfaces. Additionally, a lower percentage of its customers utilize the CAIP across multiple distinct use cases, limiting its potential enterprisewide impact.
- **Deployment options:** Sprinklr offers less flexibility in deployment architectures than other vendors in this Magic Quadrant, restricting deployments almost exclusively to the public cloud. This strict limitation hinders its capability to deploy in highly sensitive and heavily regulated enterprise scenarios.

Yellow.ai

Yellow.ai is a Niche Player in this Magic Quadrant. It is a private company located in San Mateo, California. Yellow.ai's CAIP focuses on AI-driven automation and multiagent systems for omnichannel customer experiences across chat, email and voice channels. The product suite includes a live chat module, a real-time web/mobile messaging solution with AI-to-human handoff; agent assist, for omnichannel helpdesks; and Enterprise Search, a cognitive search engine that leverages agentic RAG. Yellow.ai's customer base is mostly in APAC, North America and EMEA; it includes industries such as financial services, healthcare, retail, utilities and manufacturing. Roadmap items include features for autobuilding bots via agentic discovery; a customizable analytics module, to

add business-specific pivots for enhanced visibility and insights; and improvements to the Yellow.ai AI Trust Center.

Strengths

- **Marketing strategy:** Yellow.ai demonstrates adaptability by targeting enterprise buying groups, strengthening its partner ecosystem and integrating AI into its operations. Its metrics for measuring marketing effectiveness are among the clearest and most structured of all vendors evaluated in this report.
- **Diversity of deployments:** Yellow.ai supports a diverse breakdown of deployments across channels like website chat, messaging apps and telephony. Compared with peers, a significantly higher percentage of its customers leverage the CAIP to enable multiple use cases, proving superior omnichannel versatility.
- **Research investment:** Yellow.ai's spending on product development and R&D is significantly higher than that of similarly sized organizations in this evaluation. This puts the vendor in a favorable position to outpace competitors of similar size in innovation within this market.

Cautions

- **Operations:** Recent organizational changes are likely to affect Yellow.ai's operations. The vendor's delivery, service capacity and responsiveness must be evaluated carefully by prospective buyers to mitigate potential risks to service quality and project outcomes.
- **Business model:** Yellow.ai's business model is less differentiated and specific compared with other vendors in this research. Additionally, it offers a much narrower suite of additional add-on products (e.g., live chat, agent assist, enterprise search, and AI Agentic Discovery) than most competitors evaluated in this report.
- **Process management features:** Yellow.ai's AI agent building hierarchy provides less visual clarity than other CAIPs evaluated in this report, requiring developers to spend more time structuring agents than on peer platforms. Its fragmented orchestration tools are scattered, offering a weaker UX compared with leading competitors.

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

- Netomi
- Salesforce

Dropped

- LivePerson

Inclusion and Exclusion Criteria

To qualify for inclusion, in addition to Gartner client relevance, and as determined by analyst expertise and opinion, providers need to meet the following criteria:

Product

The company must offer a single, stand-alone product that meets Gartner’s market definition for conversational AI platforms (CAIPs). All the following conditions must apply:

- The product must include all of the mandatory features that a CAIP is 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 is 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. “Stand-alone” in this context means that the CAIP can be purchased and is marketed as a stand-alone platform, and is a viable alternative when clients are looking for a best-of-breed conversational AI solution. The product should be stand-alone from the perspective of the technical

installation, and should not require the client to replace or modify other aspects of, for example, their chat operations or contact center software in order to build a fully functioning solution. Thus, the CAIP should be compatible with third-party live chat, interactive voice response (IVR) systems, ticketing systems, telephony systems, CRM systems, ERP systems, ITSM platforms, human agent dashboards and RPA systems. References of customers that purchased and used 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 2024.
- The company must have been actively operating in the market of conversational AI platforms (CAIPs) with a generally available CAIP prior to 1 January 2025.

Customer Base Size, Growth and Weight

All of 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:

- The provider must have a minimum of 200 customers at the date of submission.
- The provider must have a minimum of 35 new customers in calendar year 2025.
- The provider must have at least 35 active licensed customers that spend a minimum of USD $100K annually.

Geographic Spread

All of the following conditions must apply:

- No more than 80% of the provider's customer base can be in a single region. Regions are defined as North America (NA); Latin America (LATAM); Europe, the Middle East and Africa (EMEA); Asia/Pacific (APAC). Regional breakdowns as follows:
    - Sub-Saharan Africa includes countries such as Ghana, Kenya, Nigeria, Senegal, South Africa, and Uganda. For the purpose of assessing the geographic spread of customers in this research, Sub-Saharan Africa is part of EMEA.

- Northern Africa and the Middle East includes countries such as Türkiye, Egypt, Jordan, Lebanon, Iraq, Tunisia, Israel, United Arab Emirates, and Saudi Arabia. For the purpose of assessing the geographic spread of customers in this research, Northern Africa and the Middle East is part of EMEA.
- Europe includes countries from both Western and Eastern Europe, such as the U.K., France, Germany, Italy, Spain, Norway, Poland, and Czech Republic. Europe is part of EMEA.
- Central and Southern Asia includes countries such as Kazakhstan, Turkmenistan, Afghanistan, India, Pakistan, Iran (Islamic Republic of), and Bhutan. For the purpose of assessing the geographic spread of customers in this research, Central and Southern Asia is part of APAC.
- Eastern and South-Eastern Asia includes countries such as China, Japan, Republic of Korea, Malaysia, Philippines, and Singapore. For the purpose of assessing the geographic spread of customers in this research, Eastern and South-Eastern Asia is part of APAC.
- Oceania includes countries such as Australia, New Zealand, Fiji, Guam, and the Cook Islands. For the purpose of assessing the geographic spread of customers in this research, Oceania is part of APAC.
- North America includes Canada and the United States of America. For the purpose of assessing the geographic spread of customers in this research, North America is part of NA.
- Latin America and the Caribbean includes countries such as Mexico, Dominican Republic, Jamaica, Costa Rica, Brazil, Chile, and Argentina. For the purpose of assessing the geographic spread of customers in this research, Latin America and the Caribbean is part of LATAM.

Employee Base

All the following conditions must apply:

- The company must have a minimum 150 employees working on the CAI practice or business unit across all roles and functions.

- The company must have a minimum of 20% of the employees in product development/R&D.

Use Case Focus

- A minimum of 33% of the company's customers licensed to use the CAIP must be using the CAIP for external-facing use cases in customer service and support.

CAIP Features

- The CAIP must provide low-code, no-code, GenAI-assisted, and pro-code development options, plus the following features, which are mandatory in Gartner's definition of the market, as per the following thresholds:
    - AI guardrails: GA as of 31 March 2026, with over 10% of customers using it in a prod environment
    - Analytics module: GA as of 31 March 2026, with over 10% of customers using it in a prod environment
    - Baseline data security and privacy controls: GA as of 31 March 2026, with over 10% of customers using it in a prod environment
    - Conversational flow building: GA as of 31 March 2026, with over 10% of customers using it in a prod environment
    - Multichannel connectivity: GA as of 31 March 2026, with over 10% of customers using it in a prod environment
    - QA, observability and optimization tools: GA as of 31 March 2026, with over 10% of customers using it in a prod environment
    - RAG enablement: GA as of 31 March 2026, with over 10% of customers using it in a prod environment
    - Tool integration: GA as of 31 March 2026, with over 10% of customers using it in a prod environment
- The following features, which are optional in Gartner's definition of the market, must be GA as of 31 March 2026:
    - AI assistant orchestration

- Composite NLP
- Voice interaction support

Honorable Mentions

The providers that are most relevant to our clients were selected for evaluation in this Magic Quadrant. However, the decision not to evaluate a provider does not mean that the provider lacks viability. The following are noteworthy providers not included in the formal analysis. These providers could be appropriate for clients, contingent on their requirements:

Decagon: Founded in 2023 and headquartered in San Francisco, California, Decagon offers a platform for building, optimizing and scaling conversational AI agents, primarily in customer service and support. The platform includes modules such as “Experiments” for A/B testing and “Watchtower” for conversation analytics and monitoring. Decagon’s customer base is predominantly located in North America.

Microsoft: Headquartered in Redmond, Washington, Microsoft offers Copilot Studio, a conversational AI platform launched in November 2023. Copilot Studio enables organizations to build, manage and deploy custom conversational AI agents, and works in conjunction with other AI agent products in Microsoft’s broader suite, such as Foundry (AI engineering for agent design, deployment and governance) and Agent 365 (a control plane for management, governance and security of AI agents). Copilot Studio supports a wide range of internal- and external-facing use cases.

Parloa: Parloa is headquartered in Berlin, Germany and New York City, New York. Its CAIP, the Parloa AI Agent Management Platform, is designed to help organizations build custom conversational AI agents while efficiently scaling conversational volumes, predominantly in customer service and support use cases. The vendor’s R&D efforts are supported by Parloa Labs, which focuses on advancing research in voice and augmentation of agentic AI capabilities.

Rasa: Rasa is headquartered in San Francisco, California, and entered the CAIP market in 2016. Its enterprise solution, the Rasa Platform, enables organizations to build, deploy and manage conversational AI agents, mainly in customer service and support. Rasa’s open-source offering complements its enterprise CAIP and is recognized as a distinctive

and established toolkit for building custom conversational AI applications, providing users with full ownership and control.

Sierra: Based in San Francisco, California, Sierra was founded in 2023 and specializes in building, managing and optimizing conversational AI agents for enterprise customers. Sierra's platform allows for use cases across the entire customer experience spectrum, including customer service, but also marketing and sales via no-code options with its Agent Studio, alongside its Agent SDK for advanced implementations.

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 current and future 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 critical, reflected in a high weighting for the product criterion. Organizations seek platforms that offer sophisticated features for AI agents and GenAI enablement, advanced QA and analytics tooling, superior usability for nontechnical users and enterprise-grade reliability, both in terms of governance and compliance. These attributes are challenging to deliver and differentiate in a highly competitive landscape.

Sales execution — which includes effective pricing models, a deep understanding of buyer personas and evidence of proven sales performance — is also rated as a high-importance criterion in such a competitive landscape. The same holds true for marketing execution, as the strategic allocation and effectiveness of marketing spending, along with the quality of marketing initiatives, are key to maximizing reach and impact and standing out from the crowd.

Criteria such as overall viability and market responsiveness/track record are assigned medium weightings. 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.

Similarly, customer experience — including the efficiency of service activation, reliability of service levels, availability of service partners and user feedback — serves as an important differentiator, but remains secondary to product, sales and marketing execution.

Operational excellence — including leadership quality and organizational size and capability — also supports a vendor's ability to scale and sustain successful deployments, but is considered less critical than core product and go-to-market capabilities.

Ability to Execute Evaluation Criteria

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

Source: Gartner (July 2026)

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 and highly competitive 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 thought leadership, research and advances in natural language technologies and AI can quickly redefine competitive dynamics.

The remaining evaluation criteria — 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.

Completeness of Vision Evaluation Criteria

Evaluation Criteria	*Weighting*
Market Understanding	High
Marketing Strategy	Medium
Sales Strategy	Medium
Offering (Product) Strategy	High

Evaluation Criteria	*Weighting*
Business Model	Medium
Vertical/Industry Strategy	High
Innovation	High

Source: Gartner (July 2026)

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 to respond 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 are likely to focus more on incremental improvements and established practices, rather than on driving innovation or shaping new market directions. While they are effective at meeting current customer needs, Challengers are not likely to 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 are not likely to have achieved the same level of market presence, operational maturity, product comprehensiveness 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 are likely to excel in specialized areas or provide strong capabilities for targeted scenarios, but generally have a narrower scope compared with Leaders or Visionaries. Niche Players 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 CAIPs rapidly evolve to incorporate the latest technology trends and expand their feature sets, buyers face increasing complexity in the selection of such platforms, as well as heightened cost considerations. Key insights to understand this market in 2026 include the following:

- **Regulatory and governance needs:** The regulatory landscape is in flux, with compliance requirements adapting to new AI regulations and shifting geopolitical realities. On risk mitigation in general, the absence of standardization in tooling further complicates governance and compliance, especially when orchestrating multiple systems.
- **M&A trends:** The CAIP market is being reshaped by accelerated mergers and acquisitions (M&A), introducing new risks and uncertainties regarding product roadmaps, service continuity and long-term vendor viability.
- **Scalability:** In this environment, the scalability advantage of CAIPs — compared with, for example, prebuilt AI assistants and AI engineering environments — remains a

critical differentiator through 2026. Organizations that opt to build conversational AI solutions internally must recognize that design and scoping of AI assistants and CAI agents, advanced analytics, quality assurance, deep integration with workflow systems, reporting and continuous optimization are essential for demonstrating ROI, driving user adoption and ensuring optimal user experience. These functions, along with robust governance, AI guardrails and cost optimization, require sophisticated, productized features that are difficult to replicate without the support of specialized vendors.

- **Cost optimization:** Buyers should pay particular attention to cost optimization, as lack of clarity around cost escalation remains a significant concern. Additionally, M&A activity heightens the need for proactive vendor management.

To mitigate risks and maximize value, enterprise application leaders, in collaboration with enterprise architects, and sourcing, procurement and vendor management (SPVM) teams, should:

- Ensure compliance and governance capabilities are sufficient to meet evolving regulatory requirements, especially in complex multiagent environments (see **7 Key Guardrails for Conversational AI Agents and Assistants** and **Governance Challenges and Solutions for AI Agents**).

- Proactively manage vendor contracts and develop contingency plans to maintain operational continuity amid ongoing M&A activity to mitigate the risk of incurring unexpected costs, loss of AI assets and stalled innovation as M&A-driven disruptions can lead to product gaps, service outages and potential data security issues (see **Prepare for M&A Disruptions in the AI Vendor Landscape**).

- Prioritize CAIP vendors that offer clear scalability via flexible model and tool choices and smart composition of AI techniques, robust governance, and advanced analytics as out-of-the-box features. (see also **Toolkit: RFP Template for Conversational AI Platforms**).

- Rigorously evaluate cost structures and cost optimization features and best practices, and demand transparency to avoid unexpected cost escalation (see **Optimize Costs When Scaling AI Assistants and Conversational AI Agents**, **Improve AI Efficiency Beyond Cost Optimization** and **AI Tokenomics: Measuring LLM Usage in AI-Enabled Applications**).

Market Overview

The conversational AI platform (CAIP) market continues to demonstrate strong relevance and growth through 2026, fueled by enterprises' increasing demand for scalable, low-code and no-code solutions to develop tailored conversational applications. Gartner executed a preliminary analysis of the market as a prework to this Magic Quadrant research. Survey responses from 21 vendors indicated that, collectively, they added over 1,000 new customers in the past 12 months (from February 2025 through January 2026). Additionally, over the past year, the market has undergone significant shifts, particularly in response to evolving buyer priorities and the rapid advancement of adjacent technologies, such as agentic AI and multimodal interaction capabilities.

Mergers & Acquisitions

M&A activity has accelerated, significantly reshaping the vendor landscape (see **Prepare for M&A Disruptions in the AI Vendor Landscape**). While M&A can enhance product suites and integration capabilities, it also introduces uncertainty regarding product roadmaps, service continuity and long-term vendor viability. As a result, organizations must carefully evaluate not only the technical capabilities of CAIP solutions but also the stability and strategic alignment of vendors within their broader digital transformation initiatives.

Early Market Maturity

The presence of hyperscalers like Google and Salesforce among the Leaders, IBM as a Visionary, combined with an increasing rate of acquisitions in this space, indicate that the CAIP market is entering early maturity. Market differentiation is shifting from product innovation to other dimensions, such as operational sustenance, pricing and scale, with smaller vendors being acquired or becoming Niche Players as the market matures.

Build Versus Buy

Gartner continues to observe that CAIPs are being challenged in their value proposition by end users that attempt to build their own conversational applications internally, often leveraging AI engineering tools. Vendors — particularly hyperscalers — that have invested in this area have advanced their platforms by providing robust pro-code options and expanding adjacent code generation tools to enable the creation of AI agents.

Governance

Buyers are placing greater emphasis on robust governance, compliance and risk mitigation, including the implementation of comprehensive AI guardrails. However, challenges persist, particularly regarding the standardization of these approaches. Many governance and guardrail features are still delivered in an ad hoc manner — often relying on prompt engineering (i.e., instructions embedded within prompts that are sent to the LLMs) — rather than as multilayered policy controls or composite AI (integrating rules, ML models and prompts). This is especially problematic for orchestrating multiagent deployments, which require governance capabilities that extend beyond the needs of a single application.

Cost Optimization

Many organizations continue to face uncertainty around cost escalation, and the lack of standardized tooling complicates scalability and cost management. While there is some convergence around best practices — such as integration with cloud FinOps tooling or the use of caching strategies for GenAI responses — the maturity and readiness of CAIPs in supporting these practices varies significantly across vendors. See **Optimize Costs When Scaling AI Assistants and Conversational AI Agents**, **Improve AI Efficiency Beyond Cost Optimization**, and **AI Tokenomics: Measuring LLM Usage in AI-Enabled Applications**.

Technology Trends

Fast-paced technological innovation remains a defining characteristic of the CAIP market:

- There is a clear shift from traditional workflow automation to more **agentic AI approaches**, including the adoption of AI agent communication protocols such as Model Context Protocol (MCP) and agent to agent (A2A).
- **Multimodality** is becoming increasingly important, extending conversational capabilities beyond text and voice to include additional modalities such as image and video within omnichannel experiences. This enables the seamless rendering of contextual, synchronized multimodal content across different channels — even in real time during interactions with end customers. However, the reliability and sophistication of these features differ widely among vendors.

- The emergence of **self-building** and **self-learning** capabilities is also notable, with some CAIPs now able to autonomously build, maintain and test conversational applications through AI-powered functionality. GenAI-assisted coding and pervasive AI-driven analytics are further expanding the self-learning value proposition noted in roadmaps of CAIP solutions.

Differentiation

Ultimately, many CAIP vendors have yet to fully recognize that differentiation in this market extends beyond core product functionality. True differentiation encompasses adjacent products, service offerings, pricing models and the alignment of vendor brands with evolving end-user needs (see **AI Vendor Race: Crush the Competition: Differentiate With 4 Dimensions in Conversational AI**). As the definition of the CAIP market continues to evolve in response to shifting buyer priorities and rapid technological innovation, organizations are advised to assess both the functional breadth and strategic fit of solutions, ensuring alignment with their risk posture and long-term digital transformation objectives.

Acronym Key and Glossary Terms

ACP	Agentic Commerce Protocol
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
UCP	Universal Commerce Protocol

⊕ Evaluation Criteria Definitions