

Magic Quadrant for AI Code Assistants

15 September 2025 - ID G00824517 - 55 min read

By Philip Walsh, Haritha Khandabattu, [and 3 more](#)

The AI code assistant market continues to burn hot as rapid innovation and intense competition is fueled by disruptive advances and new entrants. Use this research to compare vendors, navigate key trends, and select the right AI code assistant to enhance developer productivity, foster creativity, and enable robust, contextual software development workflows.

Strategic Planning Assumptions

- By 2027, 80% of software engineering leaders will consider prompt (or context) engineering a highly important skill to deliver software that meets the business needs of their organizations, up from 20% in 2024.
- By 2028, 90% of enterprise software engineers will use AI code assistants, up from less than 14% in early 2024.
- Through 2028, a 30% productivity gain in software development will be achieved across enterprises as a result of applying multiple AI-powered tools in the software development life cycle.

Market Definition/Description

Gartner defines AI code assistants as tools that generate and analyze software code and configuration. They use foundation models like LLMs, program-understanding technology, or both. Developers engage with these assistants to generate, analyze, debug, test, fix, refactor code, search dependencies, update libraries, create documentation, understand code,

upgrade versions, translate languages and review commits. They help developers learn and explore codebases and access related information, such as frameworks and tools. AI code assistants integrate with developer environments, code editors, command-line terminals, chat interfaces, project management tools, monitoring, logging and deployment tools. Some are customized to an organization's specific codebase and documentation.

AI code assistants enhance software developers' experience by boosting their efficiency, accelerating application development, minimizing cognitive overload, amplifying their problem-solving skills, enabling faster learning, fostering creativity and maintaining their state of flow.

Mandatory Features

- Context-aware multiline code completion, including natural language-based suggestions
- Unit test and documentation generation
- General purpose coding assistance across multiple programming languages and frameworks, rather than specialized support for a proprietary vendor ecosystem
- Contextual awareness of enterprise code repositories and other enterprise data sources (such as documentation capturing coding styles, design patterns and architectural decisions)
- Conversational chat interface integrated into the development environment with ability to reference or "mention" specific files and folders, and transfer code from chat into the editor
- Guarantee that base models will not be trained on customer code or documentation (excluding approved fine tuning)
- Enterprise administration support including code exclusions, codebase indexing, programming design styles, usage metrics and user management

Common Features

- Multiple deployment options including SaaS, on-premises or private cloud instances.
- Multifile edits and generation based on chat requests
- Code quality scanning and vulnerability detection

- Command completion in a terminal (command-line interface)
- Ability to fine-tune models on proprietary enterprise data for a custom model
- Adaptive personalization based on developer preferences, feedback and style
- Developer intent prediction (next action)
- Specialized agents that autonomously execute complex development tasks such as unit tests, feature development, code modernization, refactoring and remediation.
- Integration with popular project management and DevOps platforms (e.g., Jira, GitHub, CI/CD pipelines) to support workflow continuity
- Support for model context protocol (MCP) to enable standardized context sharing between tools and models
- Configurable AI guardrails for hallucination reduction, bias detection in code, content moderation, and enforcement of security and compliance policies
- Open-source code snippet attribution
- Support for multiple natural languages (such as English, Spanish, Hindi and Mandarin)

Magic Quadrant

Figure 1: Magic Quadrant for AI Code Assistants





Gartner.

Vendor Strengths and Cautions

Alibaba Cloud

Alibaba Cloud remains a Challenger in this Magic Quadrant. Its Lingma AI coding assistant increases developer productivity with features like code completion, chat, multifile edits, and agent capabilities across various integrated development environments (IDEs). Alibaba Cloud operates globally, with a strong presence in China and growing reach in North America; Europe, the Middle East and Africa; and Asia/Pacific. Clients include individual and enterprise developers across industries. With over 1 million new users in the past year, Alibaba Cloud is expanding its global cloud and AI infrastructure to drive further expansion.

Strengths

- **Operations:** Alibaba continues to offer strong service-level agreements (SLA), reliable historical uptime, and zero-downtime upgrade processes, maintaining its reputation for operational excellence at scale. Strategically placed regional service teams across all continents provide around-the-clock support, and its security posture includes multitenant SaaS offerings on public cloud infrastructure, dedicated instances with virtual private cloud (VPC) network isolation, and extensive compliance certifications.
- **Overall Viability:** As part of Alibaba Group, a publicly listed and financially strong company, Alibaba Cloud has substantial resources and market stability. The vendor demonstrates strong market presence in China, high customer retention, and continuous, heavy investment in its AI code assistant product by allocating hundreds of dedicated full-time employees and committing to increased future investments.
- **Geographic Strategy:** Alibaba Cloud enhanced its geographic strategy by officially commercializing its product internationally in 2024 and launching broad market expansion efforts in 2025. Although its growth remains primarily concentrated in China, Alibaba Cloud has successfully penetrated global markets, expanding its user base across the Asia/Pacific region as well as North America.

Cautions

- **Market Understanding:** Alibaba Cloud's market vision lacks distinct long-term differentiation. While Alibaba articulates important themes — such as progression toward multiagent orchestration, customer data security and integration with DevOps toolchains — most of its vision centers on capabilities already present in today's market or expected in the next 12 to 24 months (e.g., IDE plug-ins and AI-native IDE, retrieval-augmented generation [RAG], and supervised fine-tuning).
- **Marketing Strategy:** Alibaba Cloud's marketing approach provides limited emphasis on vertical-specific segmentation and features moderately targeted buyer personas, which may constrain its ability to differentiate competitively. Its marketing resource allocation, while adequate, remains lower than leading competitors.
- **Product or Service:** Lingma's roadmap and product architecture are heavily centered on Qwen, Alibaba's proprietary model, with limited emphasis on model optionality or broader ecosystem flexibility. While Qwen performs strongly on coding tasks, this tight

coupling may constrain developer choice and adaptability, particularly as leading vendors increasingly prioritize multimodel support to accommodate diverse use cases.

Amazon

Amazon remains a Leader in the Magic Quadrant. Its Amazon Q Developer enhances productivity across the software life cycle with AI-driven code generation, transformation, vulnerability scanning, and agentic capabilities, especially for Amazon Web Services (AWS) environments. AWS operates globally, serving a wide range of clients, from individual developers to large enterprises, across all industries, including those that build and operate applications on AWS. Its investments in AI include expanding generative AI, new agentic experiences for troubleshooting and resource analysis in the AWS console (including cost optimization), and enhancing its AI infrastructure and global data centers to support growing AI-driven development.

Strengths

- **Innovation:** AWS demonstrates notable innovation by surfacing its AI capabilities across multiple IDEs (including its new AI-native IDE, Kiro, which is still in preview and not included in this evaluation), terminals, and DevSecOps platforms — including integrations with GitHub and GitLab. It pioneered specialized agent solutions, notably its modernization-focused code transformation agent, and has invested in automated reasoning, further showing its ability to deliver solutions by combining diverse AI techniques.
- **Customer Experience:** Utilizing a comprehensive approach, including advisory boards, surveys, beta programs, and direct engagement, AWS integrates feedback into iterative releases. These customer responsiveness mechanisms directly fuel the innovation described above, continuously aligning AWS’s capabilities with customer demands.
- **Operations:** AWS leverages its extensive global infrastructure to offer operational support, including robust security controls, zero-downtime upgrades, and globally distributed Technical Account Managers (TAMs). Its “follow-the-sun” support model guarantees rapid response times, ensuring continuous and reliable service, crucial for enterprise-scale customers requiring consistent high performance and security.

Cautions

- **Business Model:** Amazon Q Developer's business model introduces complexity that may make cost planning and competitor comparisons more challenging. Additionally, the lack of fully defined pricing for some components (e.g., AI Teammates) and ongoing changes to the pricing model may create uncertainty for enterprise buyers seeking predictable long-term cost structures.
- **Sales Execution/Pricing:** AWS supports Amazon Q Developer through its broad sales organization and specialist teams. However, compared with some hyperscaler competitors that apply more concentrated sales efforts to their AI code assistant products, AWS' approach may present challenges in achieving equally strong traction within targeted market segments.
- **Product or Service:** Amazon Q Developer remains limited in enterprise customization and deployment options — such as private/self-hosted deployment, enterprise prompt repositories, and model fine-tuning — where some vendors currently offer stronger native support. Its deep optimization for AWS-native services benefits AWS customers but offers fewer features for non-AWS environments, which may raise switching costs. Many of these limitations have been addressed in Kiro, but were made after the cut-off for our assessment.

Anysphere (Cursor)

Anysphere is a Challenger in the Magic Quadrant. Its VSCode-based editor, Cursor, provides sophisticated autocompletion for code generation, advanced editing and context awareness, and agentic task execution. Operating globally, Cursor serves over 50,000 teams and over 10 million individual users. The company continues to apply its investor funding to grow its feature set, leveraging R&D to enhance its platform capabilities, developer experience, security controls, and proprietary models. Anysphere continues to invest in expanding the platform's features, improving developer experience, strengthening security controls, and developing its proprietary models as it scales its offerings to larger teams and organizations.

Strengths

- **Product or Service:** Cursor provides an AI-native IDE, incorporating powerful code indexing and advanced contextual autocomplete that leverages the current IDE state — including active files and open tabs — to predict and suggest edits. It maintains comprehensive agent capabilities, offering users the flexibility to select from multiple large language models (LLMs), including proprietary models.

- **Market Responsiveness Record:** Cursor has a strong track record as an early mover, introducing several capabilities now central to AI-native IDEs, while also actively engaging with and contributing to open-source communities to ensure alignment with developer expectations and industry standards. It employs diverse mechanisms to capture and promptly respond to customer feedback, continuously adapting its offerings to meet evolving market needs.
- **Market Understanding:** Cursor has articulated a compelling and clear vision for AI's role in software development over the next five years. Recognizing the shift toward agentic software engineering practices, Cursor positions itself at the forefront of this evolution by aligning its roadmap and feature development with a future where developer workflows are driven through AI task-delegation and agent orchestration.

Cautions

- **Business Model:** Cursor lacks transparency regarding potential changes to its pricing strategy. Recent pricing changes to its individual pro plan led to unexpected cost overruns for customers. This has raised concerns among some enterprises considering how to forecast and budget for Cursor's services.
- **Sales Strategy:** Cursor's go-to-market approach has historically relied predominantly on organic, word-of-mouth growth without a mature enterprise sales strategy. It currently lacks a robust network of indirect sales channel partners and, while currently expanding, maintains a smaller dedicated sales team compared with Leaders in this Magic Quadrant.
- **Vertical/Industry Strategy:** Cursor maintains a primarily horizontal approach with minimal strategic vertical partnerships. Additionally, as of the time of publication, it has yet to achieve industry-specific certifications, such as Federal Risk and Authorization Management Program (FedRAMP) authorization, potentially limiting its appeal to sectors like government and healthcare, which have heightened security and compliance requirements.

Augment Code

Augment Code is a Visionary in this Magic Quadrant. Its product is explicitly focused on empowering developers within large, complex enterprise codebases by offering AI-driven code generation, completion, and intelligent task execution, as well as robust documentation generation and multiagent orchestration for complex workflows. Augment

Code is headquartered in the U.S. and operates globally, primarily serving enterprises and developer teams managing large codebases. It's currently investing in advanced AI agent capabilities — including remote agents for concurrent tasks and expanded context for multirepo projects — and prioritizing enterprise-grade security and data privacy.

Strengths

- **Product or Service:** Augment Code excels at customization and context awareness for large complex codebases in enterprise environments. Its analysis of code behavior during execution and documentation capabilities are standout features.
- **Market Understanding:** Augment Code emphasizes the long-term evolution of AI assistants from simple productivity tools to integral development partners that automate significant portions of the software development life cycle (SDLC). Its focus on human-steered agents, deep enterprise integration and cross-functional team collaboration demonstrates foresight into how AI will transform complex engineering workflows over the next five years, from capabilities to how teams build, secure, and operate software.
- **Sales Strategy:** Augment Code has established an effective enterprise sales strategy that combines product-led growth with high-touch enterprise engagement. Augment Code employs migration incentives — including price matching and enhanced onboarding support — to successfully displace competitors, which has accelerated enterprise customer acquisition.

Cautions

- **Overall Viability:** Augment Code, a relatively small venture capital-funded firm, secured substantial Series A funding (\$227 million) at a near-billion-dollar valuation in April 2024. However, its lack of profitability, given its early growth-focused stage, raises potential viability concerns for enterprise buyers, particularly regarding long-term independence or potential acquisition scenarios.
- **Vertical/Industry Strategy:** Augment Code employs a purely horizontal strategy without industry-specific solutions or vertical specialization. While common among new entrants, this may limit appeal to enterprises seeking tailored vertical integrations.
- **Operations:** Augment Code offers a formal SLA in its enterprise agreements, but its support organization is primarily centralized in North America. Enterprises operating globally may find this limited geographic distribution of support resources a constraint

when compared with peers that have more mature, regionally distributed support models.

Cognition (Windsurf)

Windsurf (formerly Codeium) is a Leader in the Magic Quadrant. As of 14 July 2025, Windsurf announced its acquisition by Cognition. Its AI-powered platform, including Windsurf Editor and Plugins, accelerates developer productivity with AI-driven real-time code completion, generation and multistep agentic workflows across the software life cycle. Operating globally, Windsurf serves 4,000-plus clients, from individual developers to large enterprises, including major financial and defense contractors. It's currently investing in agentic AI, such as Cascade, for deeper codebase understanding, expanding its Model Context Protocol for more integrations and prioritizing flexible deployment and enterprise security.

Strengths

- **Product or Service:** Windsurf offers extensive programming language support and robust IDE integration through its proprietary AI-native IDE and popular industry-standard IDEs, such as Visual Studio Code, Visual Studio, JetBrains IDEs, Eclipse, Jupyter, Emacs, Xcode, and Vim. It excels at enterprise code customization, testing, debugging, document generation, and requirements management while demonstrating exceptional flexibility through agent integrations across SDLC events.
- **Innovation:** Windsurf's Cortex reasoning engine significantly enhances complex, multistep tasks, bolstered by powerful semantic retrieval mechanisms and extensive context window capabilities (up to 1 million tokens). This continuous and aggressive reinvention differentiates Windsurf from competitors.
- **Customer Experience:** Customers praise Windsurf's exceptional flexibility, ease of deployment, and responsiveness to evolving enterprise needs. Its hybrid deployment models and adaptable enterprise integrations consistently earn high satisfaction scores, contributing to strong retention rates and loyalty.

Cautions

- **Overall Viability:** Windsurf's viability faces uncertainty following Google's licensing of its core technology and hiring of its CEO and key founding research team members. This significant loss of leadership and foundational talent raises concerns about Windsurf's future capacity for sustained innovation and competitive positioning.

- **Business Model:** Windsurf's recent acquisition by Cognition creates uncertainty around its future pricing and licensing strategy, particularly as it integrates advanced agent capabilities and navigates the shift from seat-based to variable consumption-based pricing. This ambiguity complicates budget forecasting and cost planning for enterprises evaluating long-term investments in Windsurf's solutions.
- **Sales Strategy:** Despite demonstrating strong sales execution within its existing developer-centric markets, Windsurf's expansion into new market segments and strategic verticals faces limitations due to comparatively fewer dedicated go-to-market resources. Its sales approach is less mature than other leaders in this Magic Quadrant in scaling channel partnerships or deeply penetrating vertically specialized market segments.

GitHub

GitHub remains a Leader in this Magic Quadrant, with its flagship product (GitHub Copilot) focused on enhancing developer productivity across the software development life cycle. Copilot offers AI-driven code completion, generation and translation, runtime awareness (i.e., analysis of code behavior during execution), context enrichment, and automated documentation using models from Anthropic, Google, and OpenAI. GitHub serves a global client base of 77,000 enterprises, with 20 million GitHub Copilot users. It continues to invest in Copilot's agentic capabilities like Agent Mode for autonomous tasks and intends to embed AI throughout developer workflows.

Strengths

- **Product or Service:** GitHub Copilot excels in essential developer productivity tasks, including advanced capabilities in code translation, runtime understanding, enriched contextual suggestions, and automated documentation. Its integration of leading AI models further reinforces its strong technical offering.
- **Market Understanding:** GitHub keeps pace with market trends, clearly articulating its vision for evolving from AI-assisted to increasingly autonomous agentic software engineering. It emphasizes enterprise readiness, governance controls, and deep integration into existing codebases and developer workflows, positioning human developers as central, yet increasingly empowered by AI.
- **Offering (Product) Strategy:** GitHub's strategic evolution is marked by deep ecosystem integration and expansion across SDLC tasks, with GitHub Copilot natively embedded in GitHub and Visual Studio Code to provide agentic capabilities such as code review,

autofixes, and error remediation. This AI-native approach reflects a comprehensive product roadmap designed for enterprise-grade environments.

Cautions

- **Vertical/Industry Strategy:** GitHub remains predominantly horizontal in its approach, maintaining minimal industry-specific customization or targeted vertical strategies. This horizontal stance may limit the appeal to enterprises requiring highly tailored, industry-specific solutions.
- **Sales Execution/Pricing:** Despite leveraging Microsoft's substantial sales infrastructure, customers frequently encounter complexity and confusion in pricing structures related to Microsoft contractual arrangements. The multiplicity of SKUs (Copilot Free, Pro, Pro+, Business and Enterprise) and the interaction with GitHub Enterprise subscriptions contribute to perceived opacity and potential cost surprises for customers.
- **Sales Strategy:** GitHub's sales strategy includes Microsoft-driven sales channels. While beneficial in scale, the complexity of cross-platform pricing strategies create friction for customers seeking clear and direct negotiation paths, especially when differentiating between GitHub and Microsoft contractual terms.

GitLab

GitLab remains a Leader in this Magic Quadrant. Its offering, GitLab Duo, enhances developer productivity and automating tasks across the DevSecOps platform via AI-driven code suggestions, chat, security scanning, vulnerability resolution, and continuous integration/continuous delivery (CI/CD) troubleshooting, using a platform-level knowledge graph. Its globally diversified operations serve large organizations and teams in regulated sectors like financial services, the public sector, and e-commerce. GitLab is expanding Duo's agentic capabilities, including the new Duo Agent Platform for orchestrating AI agents across the SDLC, and is making enhancements to its context engine for more intelligent and relevant assistance. It also offers flexible self-hosted and cloud deployment options.

Strengths

- **Business Model:** GitLab maintains extensive dedicated resources spanning product strategy, development, marketing, sales, operations, support, and professional services. Partner-driven customer acquisition remains significant, enhancing customer growth through strategic partner references.

- **Offering (Product) Strategy:** GitLab's Duo AI offerings effectively leverage the established strengths of GitLab's widely adopted DevOps platform. Its integration across the existing DevOps platform infrastructure enhances efficiency, reliability, and depth of insights, particularly beneficial in highly regulated environments.
- **Operations:** GitLab excels in deployment flexibility — supporting on-premises, multitenant, and single-tenant SaaS models — with Duo Self-Hosted tailored for regulated industries through private cloud, air-gapped or fully on-premises solutions that ensure strong data sovereignty, security, and privacy controls. Its globally scalable infrastructure offers 24/7 support in all regions and complies with a wide range of industry certifications.

Cautions

- **Product or Service:** GitLab Duo scores lower than its competitors in Gartner's evaluation regarding enterprise code customization capabilities, code translation, and context enrichment, which are crucial to enterprise customers seeking highly customizable and contextually rich AI solutions. Recent releases by Gitlab are beginning to bridge that gap.
- **Marketing Strategy:** GitLab's marketing emphasizes platform breadth and AI for DevOps rather than highlighting core AI code assistant capabilities that appeal most to developers and drive bottom-up adoption. As a result, GitLab risks limited developer mind share in a market where leading tools succeed through strong grassroots advocacy from the developer community.
- **Market Responsiveness/Record:** GitLab has not been an early mover on key AI code assistant capabilities, with features such as chat functionality, AI impact dashboard, and self-hosted options arriving well after competitors' offerings.

Google Cloud

Google Cloud remains a Leader in this Magic Quadrant. Its Gemini Code Assist product focuses on deeply integrating AI into developer workflows across the software development life cycle, emphasizing productivity and Google ecosystem compatibility, and leveraging Google DeepMind's AI research. Its global, reliable operations primarily serve large enterprises and developer communities, but also support individual developers and small businesses. Google Cloud continues to expand Gemini Code Assist's agentic capabilities,

improve code quality and context management (e.g., Context Drawer), and broaden integrations across Google Cloud services and popular IDEs.

Strengths

- **Market Understanding:** Google effectively recognizes market demands for advanced developer productivity tools by offering Gemini 2.5 Pro and 2.5 Flash — which excel in coding, reasoning, and data analysis — with its extensive 2 million token context window. Google's Gemini Code Assist is strategically designed to integrate seamlessly into IDEs and Google Cloud Platform, delivering context-aware suggestions and agentic AI workflows alongside granular analytics that enable customers to clearly measure ROI and optimize their AI investments.
- **Business Model:** Google introduced a tiered pricing model, including a no-cost individual tier, to attract a broad user base to fuel long-term adoption and conversion. The launch of Gemini Code Assist Enterprise tier specifically targets enterprise needs such as security, compliance, centralized management, and custom code integration.
- **Operations:** Google's operational strengths are underscored by its global support infrastructure and strong SLAs. Through deep integration and collaboration with Google DeepMind, Gemini Code Assist optimally aligns AI model development with real-world coding use cases, enhancing speed, model performance, and comprehensive support capabilities throughout the AI life cycle.

Cautions

- **Product or Service:** Gemini Code Assist scores lower relative to other Leaders in critical capabilities such as code translation, runtime awareness, application delivery, documentation generation, and model choice flexibility. These gaps pose challenges for enterprises seeking comprehensive feature coverage across diverse development workflows. Recent releases, including Gemini CLI and new IDE and deployment integrations, begin to address some of these gaps, but these fell after the cutoff for this year's technical assessment.
- **Marketing Strategy:** Google's marketing strategy lacks clear differentiation and explicit targeting of vertical segments or specific customer sizes, contrasting with more strategically segmented approaches of other Leaders.
- **Sales Execution/Pricing:** While similar to other major competitors, Google's sales approach lacks transparency in discount structures and pricing details. This opacity in

pricing negotiation terms and discount visibility could complicate customer understanding of value and decision-making processes.

Harness

Harness is a Visionary in this Magic Quadrant. Its Harness AI product is broadly focused on infusing AI across the entire software delivery life cycle, from coding to building, deploying, and managing cloud costs. Its operations are mostly focused in North America, and its clients tend to be large enterprises across various sectors, including financial services, healthcare, and the public sector. Harness is investing in the advancement of its AI capabilities for automated testing, continuous verification, and intelligent incident response with the intent to further streamline and accelerate enterprise software delivery.

Strengths

- **Business Model:** Harness demonstrates leadership in the DevOps platform space through its proven business model and early AI adoption across the software development life cycle, including innovations like test intelligence and continuous verification that position it well for integration and scale. Its success in converting partner-sourced leads highlights strong go-to-market execution and growth potential.
- **Sales Strategy:** Harness executes a well-structured sales process, involving targeted outreach to development teams and IT leadership, and focusing on the holistic value of its solution across the SDLC. This strategy has proven to be successful in articulating the value proposition of Harness AI and securing new customer acquisitions.
- **Offering (Product) Strategy:** Harness targets the enterprise segment with a product strategy focused on platform coverage, governance, security, SDLC intelligence, and deep toolchain integration. It supports both Harness-managed and customer LLMs to prevent vendor lock-in and professional services (e.g., onboarding, migration, and education) to ensure enterprise success.

Cautions

- **Product or Service:** As a company specializing in DevOps tooling across the SDLC, Harness has access to a broad set of client data. However, its product falls short in leveraging this data to create rich dashboards from which clients can gain insight and take action.

- **Geographic Strategy:** Harness' geographic strategy currently lacks the strategic depth and aggressive expansion necessary to complement its otherwise strong business fundamentals. Its growth has been centered primarily in North America, necessitating the development of more tailored and assertive market entry plans into new regions, alongside a more nuanced approach to navigating diverse market needs and complex regulatory environments globally.
- **Marketing Execution:** Harness' execution of its marketing initiatives underperforms compared with those of its peers, especially in social media reach, resulting in limited brand awareness, poor lead generation, and inconsistent messaging in its go-to-market campaigns. This hampers market adoption and prevents full realization of its product potential.

IBM

IBM is a Niche Player in this Magic Quadrant. Its watsonx Code Assistant product is mainly focused on enhancing developer productivity through the modernization of mainframe applications (particularly COBOL modernization), IT automation (especially Ansible) and Java application life cycle improvements. Its operations are geographically diversified, and its clients tend to be large enterprises, especially those with significant investments in legacy mainframe applications, across sectors like banking, insurance, healthcare, and government. IBM is investing in expanding its generative AI capabilities to support broader programming languages and to further automate code generation, explanation, and testing across the development life cycle.

Strengths

- **Offering (Product) Strategy:** IBM continues to support enterprise-grade application modernization, with strengths in code understanding, documentation generation, and code transformation. Its product strategy is tailored to modernizing applications written in COBOL and RPG. Comprehensive consulting services further support legacy migration, augmenting IBM's tooling for holistic project delivery.
- **Business Model:** IBM's pay-per-use pricing model distinguishes it by providing flexible, transparent consumption-based billing. This approach supports enterprises with fluctuating workloads, avoiding rigid monthly charges.
- **Operations:** IBM maintains a strong global operational infrastructure, offering comprehensive uptime guarantees (up to 99.995%), rapid support response, and flexible

deployment options, including private cloud, public cloud, hybrid, and on-premises.

Cautions

- **Product or Service:** Despite significant strengths in legacy code modernization, watsonx Code Assistant still lags in critical areas like testing and debugging, context enrichment, and model choice compared with other vendors in the Magic Quadrant. Its setup process remains comparatively complex, limiting quick adoption compared with competitors.
- **Innovation:** IBM lags behind competitors in innovation due to limited investment in emerging capabilities, such as advanced AI agents and collaborative coding features.
- **Sales Execution or Pricing:** IBM's sales focus leans toward organizations managing legacy code and Java-based applications. Its sales orientation indicates that it is not pursuing the general, "greenfield" coding market.

JetBrains

JetBrains is a Niche Player in this Magic Quadrant. Its AI code assistant, JetBrains AI (including AI Assistant and Junie), enhances developer productivity with integrated AI-driven code generation, completion, refactoring, documentation, unit test creation, and commit message generation in its IDE ecosystem. JetBrains serves a global userbase of individual developers, teams and enterprises across industries. The company is investing in advanced AI agent capabilities with Junie for autonomous multistep task execution, expanding support for local and proprietary LLMs like Mellum, and improving context management in its IDEs for more proactive, intelligent assistance throughout the software development workflow.

Strengths

- **Overall Viability:** JetBrains' long-standing presence in the developer tools market, strong financial health and loyal customer base give it strong overall viability as a vendor. This stability provides a strong foundation for continued investment in AI code assistant development and ensures the long-term support and evolution of its offerings, which instills confidence in its enterprise customers.
- **Sales Execution/Pricing:** JetBrains has strong brand recognition and established channels to bring its code assistant to market. Its AI offering is priced lower than competitors, and the integration with existing IDE offerings will accelerate market penetration.

- **Innovation:** JetBrains has a strong foundation in developing innovative developer tooling. The strength of JetBrains' existing IDE and its widespread adoption in the developer community puts the company in a good position to optimize how AI capabilities are surfaced in the IDE.

Cautions

- **Product or Service:** Despite its foundational strengths, JetBrains AI's current product offering is perceived as playing catch-up in advanced AI assistant and agent capabilities compared with some leading competitors. While it provides core functionalities, its feature set in areas like proactive assistance and sophisticated AI agent behavior doesn't yet match the breadth or depth of more mature AI-native solutions in the market.
- **Customer Experience:** Our evaluation surfaced opportunities for JetBrains to improve its ease of onboarding, clarity of AI features, responsiveness of technical support and incorporation of user feedback. Enhancing these elements will be important for driving higher adoption rates and sustaining user satisfaction in a competitive market.
- **Operations:** JetBrains' internal processes related to the delivery and support of its AI code assistant must evolve to maintain parity within this market. It has no universal SLA and has had several major-to-critical outages in the past year, and its flagship AI agent solution, Junie, remains SOC 2 noncompliant.

Qodo

Qodo (formerly CodiumAI) is a Visionary in the Magic Quadrant. Its Qodo Gen platform focuses on a quality-first, enterprise-grade AI coding agent for the entire software development life cycle, offering intelligent code and test generation, automated code reviews (with Qodo Merge), and robust documentation. With teams in Israel, the U.S., and Europe, Qodo serves large enterprises and quality-focused development teams in industries like tech, finance, healthcare, and e-commerce. Qodo is investing in advanced agentic AI capabilities through new tools like Qodo Command, deeper codebase understanding, flexible deployment (SaaS, private cloud, and on-premises), and enterprise security.

Strengths

- **Product or Service:** Catered to security-conscious organizations, Qodo Gen offers precise-governance-friendly AI capabilities, and in Gartner's evaluations, demonstrated excellent runtime awareness, testing, debugging and context enrichment. Its analytics

dashboard and code review capabilities have also notably improved, positioning it competitively among enterprise-grade solutions.

- **Innovation:** Qodo successfully integrated AI agent deployments into its offerings, maintaining a clear quality-first strategy while emphasizing code integrity, robustness, correctness and adherence to standards. Its agentic capabilities leverage a powerful context engine and automated best-practice checks, streamlining code review.
- **Sales Strategy:** Qodo is aggressively scaling its direct sales organization, particularly in North America and EMEA, targeting both midmarket opportunities and strategic enterprise accounts. Investments in structured sales enablement frameworks and tailored vertical-specific playbooks position it effectively in regulated industries, and it strategically collaborates with cloud providers and DevTool marketplaces to accelerate market penetration.

Cautions

- **Geographic Strategy:** Qodo's internationalization and localization capabilities lag behind competitors, with the product's user interface currently available only in English. Although UI localization efforts are underway, this limited linguistic offering may impact broader global adoption compared with competitors with multiple language interfaces.
- **Operations:** Relative to market leaders, Qodo maintains smaller service and support operations that cover substantial business hours, but lacks regional specificity and comprehensive 24-hour coverage that is common among larger competitors. Although it holds SOC 2/SOC 3 and General Data Protection Regulation (GDPR) certifications, Qodo also trails competitors in achieving additional industry-standard certifications.
- **Overall Viability:** While Qodo has experienced commendable growth and strong customer retention rates, its reliance on venture capital funding and focus on expansion over profitability could raise viability concerns. Qodo's latest funding round was relatively modest compared with other startups in the space, marking potential constraints in sustainability and long-term operational stability from an enterprise client perspective.

Tabnine

Tabnine is a Visionary in this Magic Quadrant. Its AI code assistant delivers advanced, secure, and personalized code completion, generation, explanation, refactoring, and code review, with deep IDE integration and strong enterprise model customization. Tabnine serves

a global customer base, from individual developers to large enterprises across sectors like finance, software, and e-commerce. It continues to innovate with RAG (for global code awareness) and new AI agents (for documentation, validation, testing, and Jira-to-code), and it offers flexible deployment options (SaaS, VPC, on-premises, and on-device) with robust privacy standards.

Strengths

- **Product or Service:** Tabnine's product/service strength is driven by deployment flexibility, model choice, and enterprise code customization. Customers benefit from granular model controls, on-premises options for sensitive code and emerging multimodal support, such as image input and AI agents for test generation.
- **Market Responsiveness/Record:** Tabnine shows strong market responsiveness by integrating user feedback into new capabilities (e.g., analytics dashboards and context enrichment) and maintains a high annual contract value retention rate, indicating strong customer loyalty across 1.5 million developers.
- **Vertical/Industry Strategy:** Tabnine is expanding beyond its technology-centric roots by introducing vertical solution templates and compliance toolkits for regulated industries, supported by a global partner ecosystem (such as Accenture and Cognizant) and a strong geographic strategy.

Cautions

- **Overall Viability:** While Tabnine benefits from a longer operating history and strong customer retention relative to many newer entrants, its modest funding — most recently an \$8 million round in April 2025, bringing total funding to \$65 million — limits its ability to scale enterprise support or accelerate R&D at the pace of better-capitalized competitors.
- **Sales Execution/Pricing:** Despite transparent Basic/Pro/Enterprise tiers and automatic trials, Tabnine lacks usage-based or outcome-based pricing models, which have become an important aspect of pricing for agent-based offerings.
- **Marketing Execution:** Tabnine could improve investment in a more diverse array of marketing channels — including community forums, live chat, and social presence — to increase brand visibility and accelerate feedback cycles.

Tencent Cloud

Tencent Cloud is a Niche Player in this Magic Quadrant. Its product, Tencent Cloud CodeBuddy (formerly Tencent Cloud AI Code Assistant), enhances developer productivity across the software life cycle via intelligent code completion, generation, unit testing, code review, repair, and context-aware suggestions, integrating with major IDEs, Tencent Cloud Studio, and Xcode. Its global operation spans China, Europe, the Americas, and the Middle East, and it serves small business and large enterprise clients across sectors like gaming, financial services and web services. Investing in R&D, especially LLMs (Tencent Hunyuan and DeepSeek) and retrieval-augmented generation, deepen code understanding via agent-based features and comprehensive project analysis.

Strengths

- **Overall Viability:** Tencent's financial strength, technological resources, and diversified portfolio across social media, gaming, cloud, and fintech ensure long-term viability and profitability, directly supporting the future of its AI code assistant, CodeBuddy. Strategic investments in advanced technologies, ongoing innovation, and strong monetization underpin sustained growth and market leadership, providing a stable foundation for CodeBuddy's success and evolution.
- **Operations:** Tencent offers enterprise-grade operational reliability, backed by globally distributed support centers across North America, Latin America, Europe, and the Middle East. Its 24/7 support model ensures consistent service levels across regions, and its strong compliance posture is reinforced by a broad portfolio of international certifications, underscoring its readiness for global enterprise deployment.
- **Sales Strategy:** Tencent's strategic partnerships help it respond effectively to market demands. The transparent and flexible discount policy offers a clear, tiered structure for multiyear contracts, incentivizing long-term commitments. The company underscores a commitment to tailored solutions and dedicated support for enterprise clients.

Cautions

- **Market Understanding:** Tencent's vision for the AI code assistant market appears oriented toward incremental improvement of current capabilities, with an emphasis on post-training enhancements and back-end service integration. However, compared with leading peers, Tencent has shown less initiative in articulating or advancing a broader vision for background agents and increasingly autonomous forms of software development.

- **Market Responsiveness/Record:** Tencent's ecosystem engagement metrics remain comparatively lower than peers, with fewer customers and partners contributing directly to product evolution. In addition, Tencent has been slower than competitors in bringing certain early-mover capabilities, such as AI-native IDEs, to market.
- **Marketing Execution:** Tencent employs a broad and increasingly systematic marketing strategy — including partnerships, influencer campaigns, and high-profile product launches — but the impact of these efforts remains less consistent compared with leading peers. As a result, Tencent may face challenges in clearly differentiating its AI code assistant innovations and sustaining developer mind share in a competitive market.

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

- Augment Code
- Anysphere (Cursor)
- Harness
- JetBrains

Dropped

- **Sourcegraph:** Sourcegraph is significantly changing its product strategy, including discontinuing support for its code assistant (Cody) for most plans (Cody Free, Pro, Enterprise Starter), resulting in the code assistant no longer being broadly available to all customer segments through standard sales channels.
- **Refact.ai:** Refact.ai is an open-source, self-hosted AI coding assistant designed to integrate with IDEs like VS Code and JetBrains, offering on-premises deployment,

fine-tuning capabilities, and retrieval-augmented generation. It was included in the previous iteration but did not meet Gartner’s updated inclusion criteria for the 2025 evaluation.

Inclusion and Exclusion Criteria

To qualify for inclusion in this Magic Quadrant, providers need to:

- Offer a product that includes all of the mandatory features listed in the Market Definition and is generally available by 1 May 2025
- Have at least 100 paying customer organizations

Gartner defines general availability (GA) as the release of a product to all customers. When a product reaches GA, it becomes available through the company’s general sales channel — as opposed to a limited release, or beta version, used primarily for testing and user feedback purposes.

Evaluation Criteria

Ability to Execute

Product/service: This criterion evaluates the core capabilities, features, and quality of the goods and services offered in the defined market. It considers integration breadth across IDEs, editors, and code hosts, deployment flexibility across cloud and on-prem environments, and the ability to support multimodal interaction. It also assesses contextual intelligence, including prompt libraries, RAG configuration, model selection or fine-tuning, role-based access control, runtime awareness, and the assistant’s ability to perform language, framework, or architecture transformations.

Automation capabilities such as multiframe editing and agentic workflows are key areas of focus. In addition, the criterion examines security, governance, and compliance, including IP protection, license citation, legal indemnity, data ownership, and granular control over models, features, and developer access. It reviews the ecosystem of partnerships, integrations with SDLC and testing tools, available marketplaces, and advanced feature monetization. Lastly, it evaluates the vendor’s delivery of usage and value metrics, planned

analytics improvements, and the roadmap for enhancing AI-powered workflows through evolving agentic capabilities.

Overall viability: This criterion assesses the overall health and sustainability of the organization, including financial strength, the success of the business unit responsible for the product, and the likelihood of continued investment and prioritization within the broader portfolio. It includes a review of funding history and total capital raised, 2024 cash flow and debt-to-equity ratios, profitability status or path, and preparedness for economic downturns. Vendors should also provide fiscal-year 2024 revenue figures, 2025 projections, and the key assumptions underlying those forecasts.

The evaluation also considers market traction and customer economics, including 2024 customer retention by contract value and count, current satisfaction scores (e.g., NPS), acquisition costs, customer lifetime value, and any acquisitions in 2024 through 2025 that materially affect the roadmap or workforce. Lastly, it reviews the vendor's organizational capacity and structure, including the number of full-time employees assigned to the product as of May 2025, expected changes in headcount or investments across engineering, go-to-market, and customer success functions through 3Q25, and recent changes in senior leadership along with their cultural or strategic significance.

Sales execution/pricing: This criterion assesses the organization's capabilities in all presales activities and the supporting infrastructure, including deal management, pricing strategy, negotiation processes, presales support, and the overall effectiveness of the sales channel. It examines the customer base in detail, including the primary decision-maker roles in large enterprises, the current size and growth of the customer base, segmentation by company size, and the tenure of major enterprise clients on flagship product editions. The evaluation also reviews the vendor's monetization approach, including pricing models and rate structures, how they vary across deployment options (e.g., public cloud, private cloud, virtual private cloud, on-premises), and the minimum seat or usage thresholds required for each edition. Finally, it considers the availability and structure of free or time-limited trials, including eligibility, evaluation scope, typical conversion rates, and how these programs contribute to broader customer acquisition and expansion strategies.

Market responsiveness and track record: This criterion assesses the organization's ability to adapt to changing market conditions, customer needs, and competitive pressures, including its demonstrated agility and history of responding effectively to evolving demands. It evaluates the vendor's engagement with open-source communities, including the number

and type of contributions (e.g., code, sponsorships, leadership roles), measurable outcomes, and the percentage of customers and partners classified as active contributors, along with how this is defined and tracked.

The evaluation also considers marketplace vitality, such as year-over-year growth in active listings, or the rationale behind alternative ecosystem models when no marketplace exists. In terms of customer responsiveness, it examines mechanisms such as surveys, beta programs, advisory boards, in-product feedback tools, and social media monitoring — along with their frequency, reach, and documented outcomes. Finally, it looks at the vendor's ability to tailor offerings to specific industries, supported by representative examples, and assesses innovation leadership by identifying up to three platform capabilities launched ahead of the market, with supporting data on timing, competitive context, and business impact.

Marketing execution: This criterion assesses the clarity, quality, creativity, and effectiveness of the organization's efforts to shape market perception, promote the brand, increase product visibility, and foster a strong reputation among developer audiences. It examines how the product is positioned to development teams, the core reasons teams adopt it — with supporting metrics or examples — and the key differentiators that drive preference over competitors.

The evaluation also reviews the estimated marketing budget for 2025, the distribution of spend across channels, and highlights major campaigns or programs executed in the past year, including their reach and audience engagement strategies. Additionally, it considers the vendor's market presence across SEO, paid search, live TV, PR, and content marketing, with attention to spending levels and demonstrable outcomes. Finally, it reviews physical and virtual event participation, particularly conferences sponsored or presented at over the past 12 months, and revisits competitive positioning, focusing on what clearly sets the product apart in a crowded market.

Customer experience: This criterion assesses the products, services, and programs that help customers realize their intended outcomes, with emphasis on the quality of interactions, technical support, account management, and supporting resources. It examines developer training formats, target audiences, onboarding timelines and milestones, resource requirements at each stage, and the metrics used to confirm onboarding success.

The evaluation also includes the structure and capacity of the customer support function, including FTE headcount, escalation procedures, SLAs, hours and channels of availability,

geographic coverage of certified partners, and access to self-service tools or portals. Additionally, it considers the strength of the vendor’s customer success approach, including community engagement, dedicated success roles, ROI and value tracking metrics, and retention initiatives such as adoption campaigns, webinars, or customer events.

Operations: This criterion assesses the organization’s ability to meet goals and commitments through effective structures, systems, and practices. It reviews service-level agreements for system uptime, response times, and technical support, including any regional or customer-specific variations.

The evaluation considers release and upgrade practices, customer control over update timing, the handling and impact of recent outages, and changes in the number of FTEs dedicated to enterprise technical support. It also examines workforce management, including staff training in risk management, recent functional workforce reductions, onboarding timelines for new customers, and the structure and global coverage of support and operations centers. Additionally, it evaluates how the vendor communicates with customers across different stakeholder levels. Lastly, it considers hosting strategies that offer competitive differentiation, as well as the regulatory or industry certifications included with the product offering.

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	High
Marketing Execution	Low
Customer Experience	High

<i>Evaluation Criteria</i>	<i>Weighting</i>
Operations	Medium

Source: Gartner (September 2025)

Completeness of Vision

Market understanding: This criterion assesses the vendor’s ability to understand and anticipate customer needs and translate that insight into innovative, high-impact offerings. evaluates the clarity and credibility of the vendor’s forward-looking vision for the AI code assistant market, including expected evolution over the next five years, anticipated disruptions, and alignment of the innovation roadmap with future demand.

The criterion also examines product-market fit, highlighting the top developer needs addressed, differentiated capabilities, and the measurable business value delivered — whether through productivity gains, cost savings, or talent optimization. It considers how vendors respond to emerging challenges across technical, operational, and industry contexts, and adapt to shifting customer priorities, expectations, or adoption patterns. Lastly, it evaluates the vendor’s broader sense of responsibility to the market, including programs to engage users in strategic planning and efforts that align with long-term corporate sustainability goals.

Marketing strategy: This criterion assesses the organization’s ability to deliver clear, differentiated messaging that is consistently articulated internally and projected externally through social media, advertising, customer engagement programs, and positioning statements. It evaluates how well the organization defines its messaging and go-to-market approach, including how value propositions are tailored to distinct buyer personas — such as developers, CTOs, CIOs, and engineering leaders — and how messaging is aligned across technical and executive audiences.

The evaluation also considers visibility and adoption strategies, including developer engagement efforts, product marketing resource allocation, and campaign planning. It reviews the company’s segmentation by vertical and customer size, as well as outreach efforts and performance metrics tied to community engagement and user participation. Lastly, it assesses the vendor’s partnership strategy, with attention to newly targeted

partners for 2025, rationale for selection, and the expected outcomes — whether through expanded distribution, co-innovation, or entry into new markets.

Sales strategy: This criterion assesses the organization's ability to implement an effective and scalable sales strategy through appropriate channels, including direct and indirect sales, marketing, services, and partnerships. It reviews the primary forces shaping the sales pipeline, historical and projected lead-to-close conversion rates, current direct sales headcount, and growth plans, as well as enablement programs designed to address sales challenges and support team expansion.

The evaluation also considers forecast revenue and customer growth, along with any differentiated sales motions aimed at reaching new buyer personas or evolving usage patterns. It examines geographic and industry expansion plans for the next 12 months, identifying priority countries, regions, and verticals, and evaluates the contribution of channel partners to overall sales, including their strategic role in scaling adoption. Lastly, it reviews the commercial model, including current and planned licensing structures (e.g., SaaS, subscription, perpetual, enterprisewide, consumption-based), pricing by edition for SaaS and on-premises offerings, discounting practices (including multiyear tiers and pricing transparency), and how shifting market dynamics are influencing the vendor's overall pricing and sales approach.

Offering (product) strategy: This criterion assesses the vendor's approach to product development and delivery, emphasizing alignment with evolving market needs, differentiation, core functionality, and strategic methodology. It begins with a comprehensive understanding of the offering — its official product name, underlying technical architecture (including security design and integration points), the range of professional services available, and the vendor's network of qualified consulting or system integrator partners.

The evaluation also examines the user base by size and composition across free, paid-individual, and enterprise segments, as well as the vendor's involvement in open-source communities. It considers how adoption, satisfaction, and retention are measured through both quantitative metrics and qualitative insights. Finally, it reviews the vendor's forward-looking strategy, including top investment areas for the next 12 months, product differentiation efforts to avoid commoditization, strategic partnerships, and the formal processes used to capture, evaluate, and act on customer feedback in shaping the product roadmap.

Business model: This criterion assesses the design, logic, and execution of the organization's strategy to ensure sustained success in the AI code assistant market. It reviews any planned changes to the business or pricing model over the next 12 months, including the introduction of new SKUs, shifts from seat-based to consumption or tiered pricing, licensing adjustments, or bundling strategies, and how these align with the company's broader growth objectives.

The evaluation also considers the AI code assistant's current contribution to total company revenue and expectations for its future share. Additionally, it examines the vendor's partnership ecosystem, including the roles of resellers, hyperscalers, and services partners, along with the percentage of net-new customers influenced or referred by partners over the past year. Finally, it evaluates staffing commitments across strategy, product development, go-to-market, and support functions, and assesses whether usage or performance limitations — such as rate caps or conversation quotas — are applied at the edition level, including which editions are affected and the thresholds in place.

Vertical/industry strategy: This criterion assesses the organization's ability to strategically allocate resources, capabilities, and product development to meet the distinct needs of specific industries and market segments. It examines the depth and scope of vertical-focused partnerships, including key industry alliances, their geographic or sectoral orientation, the nature of joint initiatives (such as co-innovation or go-to-market programs), and measurable outcomes that demonstrate business value. The evaluation also reviews customer distribution by industry, with approximate counts, notable referenceable clients, and context where disclosure is restricted, to gauge vertical penetration and brand credibility. Finally, it considers the vendor's vertical growth strategy for the next 12 months, including prioritized industries, the presence of prebuilt models or customization frameworks, planned initiatives with defined KPIs and milestones, and — where applicable — how a horizontally positioned platform is tailored to meet vertical-specific requirements.

Innovation: This criterion assesses the strategic allocation of resources, expertise, and capital to create and sustain competitive advantage — whether through investment, consolidation, or preemptive positioning. It examines the vendor's product innovation strategy, including how new ideas are sourced, evaluated, and validated, and the structured methods and milestones planned to meet emerging developer and enterprise needs over the next 12 months.

The evaluation also considers the influence of key technology trends on the roadmap, the offering’s most differentiated capabilities supported by customer outcomes, and the percentage of AI assistant revenue reinvested in R&D to drive future innovation. Strategic partnerships that accelerate innovation are also reviewed. In addition, it assesses how innovation efforts translate into market differentiation — identifying the typical buyer, competing alternatives displaced, and the tangible outcomes that demonstrate a lasting advantage in performance, usability, or ecosystem strength. Lastly, it evaluates the vendor’s engagement with the broader technology community, including contributions to open-source projects or standards, and the governance structures that ensure those efforts remain relevant, active, and aligned with long-term industry progress.

Geographic strategy: This criterion assesses the vendor’s ability to strategically allocate resources, talent, and offerings to meet the distinct needs of markets outside its home region, whether through direct operations or via partners, channels, and subsidiaries. It reviews differentiated go-to-market strategies, including delivery models, sales and marketing approaches, and the structure of the partner ecosystem, along with the top three geographic initiatives or launches planned over the next 12 months and the rationale behind regional prioritization.

The evaluation also examines compliance with regional data sovereignty and privacy requirements, including technical and organizational safeguards such as data residency options, encryption protocols, administrative controls, and supporting certifications or audits. Localization capabilities are assessed across interface languages, documentation, support services, and the underlying language model, with attention to adaptations that enhance usability in non-native regions. Finally, it considers the vendor’s current and planned presence across global markets, including product availability by country, physical office footprint, headcount, customer base, channel partners, and the number of new customers acquired per region in the past year, along with stated geographic expansion goals for 2025.

Completeness of Vision Evaluation Criteria

<i>Evaluation Criteria</i>	<i>Weighting</i>
Market Understanding	High

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

Source: Gartner (September 2025)

Quadrant Descriptions

Leaders

Leaders in this Magic Quadrant for AI Code Assistants continue to set the benchmark for end-to-end developer enablement through advanced AI capabilities. They consistently integrate cutting-edge models into robust, agentic workflows that enhance productivity, code quality, and security across the SDLC. These vendors exhibit a sharp understanding of enterprise needs, translating customer feedback into iterative improvements and differentiated innovation — such as long-context reasoning, hybrid deployment flexibility, and seamless integration across IDEs and DevSecOps platforms. Leaders also stand out through their operational maturity, supporting global enterprises with scalable infrastructure, regulatory compliance, and strong service-level reliability. Their growing ecosystems, deepening product integration, and strategic clarity ensure sustained influence across diverse developer and enterprise environments.

Challengers

Challengers in this Magic Quadrant for AI Code Assistants demonstrate strong execution, underpinned by operational scale, commercial viability, and regional expansion. These

vendors exhibit high service reliability, strong SLAs, and effective go-to-market capabilities, resulting in increased user adoption and commercial traction. In 2025, Challengers continue to mature by expanding their global reach, growing their customer base across diverse markets while maintaining deep operational excellence. However, they remain limited when it comes to developing a comprehensive and long-term strategy for how the market is evolving and the business model needed to sustain innovation. Their growth is enabled more by execution than by pioneering product direction or differentiated market foresight.

Visionaries

Visionaries in this Magic Quadrant for AI Code Assistants are distinguished by their progressive innovation, emphasizing advanced capabilities such as agentic assistance, contextual awareness, and integrated quality assurance. These vendors show strong alignment with emerging developer needs — particularly around testing, security, and life-cycle-wide code intelligence — delivering differentiated capabilities that reflect a forward-looking product strategy. However, their execution continues to trail their ambition; issues such as immature go-to-market operations or scaling constraints often limit their broader market influence. Visionaries remain key players for organizations seeking cutting-edge AI assistants with a strong focus on correctness, security, and development rigor, even as they work to strengthen operational consistency and enterprise reach.

Niche Players

Niche Players in this Magic Quadrant for AI Code Assistants deliver focused solutions that align with specific enterprise needs, development environments, or modernization priorities. While these vendors may demonstrate considerable operational scale, ecosystem integration, or regional reach, they generally underperform in broad innovation or strategic vision. Their AI code assistant offerings often reflect strong foundations in developer tools, legacy system modernization, or targeted IDE experiences, but fall short in areas such as advanced AI orchestration, proactive agent functionality, or global developer momentum. Compared with Leaders and Challengers, Niche Players tend to advance more conservatively, with narrower product scope and slower feature evolution, limiting their influence on the wider trajectory of AI-driven software development. However, their specialized strengths can offer distinct advantages in certain enterprise contexts.

Context

Selecting an AI code assistant for your organization should involve a structured and collaborative effort. Software engineering leaders should establish a cross-functional task force comprising stakeholders from engineering, architecture, security, and legal teams. This diverse team is essential for addressing the nuanced implications and maximizing the strategic benefits of AI code assistant investments.

This task force should use the analysis presented in this Magic Quadrant to create an initial shortlist of AI code assistant vendors. After shortlisting, the task force should initiate pilot programs within their organizations to validate tool effectiveness and measure actual productivity impacts.

Effective measurement of productivity and return on investment (ROI) is critical. Organizations should first establish baseline metrics aligned with recognized frameworks such as DevOps Research and Assessment (DORA) metrics and the SPACE framework, encompassing satisfaction, performance, activity, communication, collaboration, efficiency, and flow (see [Market Guide for Developer Productivity Insight Platforms](#)). Evaluate the impact holistically, avoiding narrowly scoped task-level metrics alone, as they often fail to capture broader, system-level productivity gains. Instead, connect individual task-level efficiency improvements directly to strategic team and organizational outcomes, such as increased capacity for innovation and value delivery (see [How to Capture AI-Driven Productivity Gains Across the SDLC](#)).

A common challenge organizations face is low adoption rates despite investments in AI code assistant licenses. Often, less than half — and sometimes fewer than a third — of purchased licenses see active use after several months. To mitigate this, organizations should customize training to meet developers' specific roles, expertise, and team contexts, facilitating peer learning and collaboration to share successes and foster best practices. Moreover, partnering directly with developers to select and evaluate AI tools increases alignment with practical development needs and ensures higher adoption and satisfaction (see [Boost Developer Productivity by Improving Adoption of AI Tools](#)).

Finally, the proliferation of AI code assistants creates a skills-experience paradox, particularly affecting junior developers. Organizations must proactively address this gap by providing structured pathways for developing foundational skills needed alongside effective AI utilization. AI tools should be positioned as complementary learning aids, enhancing productivity through skill development rather than mere output generation. Clearly defined frameworks for evaluating traditional software engineering competencies alongside AI-

augmented skills — such as problem solving, critical thinking, and effective communication — will become essential components of successful upskilling strategies (see [How to Upskill Software Engineering Teams in the Age of AI](#)).

Market Overview

The AI code assistant market continues to burn hot as rapid innovation and intense competition is fueled by disruptive new entrants and rapid advancements. Venture capital firms and established technology companies continue to make significant investments, pouring resources into R&D and marketing as they strive to gain a foothold in the SDLC as a beachhead for future expansion. Gartner estimates the global AI code assistant market size for 2025 at \$3 billion to \$3.5 billion. ¹

This rapid growth and investment has come with notable volatility, shaped significantly by competitive pressures and the strategic influence of foundational model providers. After OpenAI's attempt to acquire Windsurf fell through, Google hired Windsurf's leadership and several key researchers, licensing their technology. Windsurf was subsequently acquired by Cognition. Cursor's rapid growth was tempered by pricing controversies driven by increased costs from its dependency on Anthropic's Claude models. The emergence of Anthropic as both a model provider and a direct competitor — through its Claude Code SWE agent — illustrates the strategic vulnerability facing startups reliant on external model providers. In response, vendors increasingly seek vertical integration and proprietary model development to mitigate dependency risks.

AI code assistants continue to be one of the most hyped technologies in software development history, as evidenced by a significant year-over-year increase in Gartner inquiry volume on this topic (a 71% increase since June 2023). Gartner's Technology Adoption Roadmap for Large Enterprises for 2025 Survey indicated that 78% of large enterprises would either be piloting or deploying AI code assistants by the end of 2024, which would suggest that well over 80% are in some adoption phase at the publication of this report. ²

However, while AI code assistants are rapidly becoming basic requirements for software developers, organizations continue to encounter challenges in upskilling teams and managing internal change effectively to achieve consistent adoption and maximum benefit.

Developer productivity remains a central promise and benefit of AI code assistants, yet the past year has seen a recalibration of expectations. Despite initial claims of productivity gains

ranging from 30% to 50%, Gartner's Developer Experience Assessment Survey shows that most enterprises report more modest gains. Specifically, 42% engineering staff report productivity gains of only 1% to 10% due to AI, while an additional 12% report no gains at all. ³. Enterprises now recognize that, while these tools significantly enhance productivity, their benefits align more realistically with incremental improvements rather than transformative ones.

Transformative gains in the AI tooling market will emerge as these tools integrate more deeply across the entire software development life cycle. These gains will be driven by AI-enablement in areas such as requirements analysis and decomposition, architecture and design, UI creation, testing, deployment, and operations. As major platform providers expand their offerings throughout the SDLC, vertical specialists may face increasing pressure to differentiate or consolidate.

Two key trends have notably shaped the AI code assistant market:

- **Enhanced Contextual Awareness:** While the size and capability of foundation models remains critical, the context layer has emerged as the primary site of innovation. The ability to dynamically surface relevant context from code repositories, developer environments, knowledge management systems, MCP servers and broader software development life cycle data is now a critical differentiator among vendors.
- **Rise of Agentic Technologies:** Agent-driven workflows have gained prominence, with leading vendors embedding synchronous and asynchronous agent capabilities within AI-native IDEs and plug-ins, enabling developers to automate and delegate complex development tasks effectively.

This current Magic Quadrant reflects an expanded market definition and an evolving set of critical capabilities and use cases, notably shifting from mandatory multi-IDE plug-ins toward AI-native or stand-alone IDEs. Over the past year, AI code assistants have evolved dramatically beyond basic code completion and debugging, achieving deeper contextual awareness across large and complex enterprise codebases. Enterprises now require solutions that address a comprehensive range of tasks across the software development life cycle, enabled by robust contextual capabilities and advanced functionalities, such as agentic workflows. Additionally, organizations increasingly demand enterprise readiness in deployment flexibility, administrative controls, security and compliance, and features for tracking adoption, productivity, and ROI.

⊕ 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.