Gartner

Licensed for Distribution

Magic Quadrant for AI-Augmented Low-Code Application Platforms

28 September 2026 - ID G00841123 - 43 min read

By Adrian Leow, Kyle Davis, and 2 more

AI-augmented low-code application platforms help developers and business users build enterprise apps faster through AI generation, agentic workflows, and governance. This Magic Quadrant assesses vendors based on their platforms' AI capabilities, execution, and vision to help IT leaders choose the right fit.

Strategic Planning Assumption

By 2028, 40% of enterprises will use autonomous AI agents in production application delivery on AI-augmented low-code platforms, up from less than 5% in 2025.

Market Definition/Description

Gartner defines AI-augmented low-code application platforms (AILCAPs) as software platforms for the accelerated development and maintenance of applications, using model-driven development tools, prebuilt component catalogs, AI code assistants, and AI agents. These platforms leverage prompt-to-app generation including the workflows, database schemas, and UI as well as vibe coding to automate application life cycles across the technology stack. AILCAPs provide collaborative environments where autonomous AI agents assist in development, runtime scaling, and governance. This approach allows users to translate high-level intent into scalable applications, featuring integrated DevOps and

automated monitoring. AILCAPs feature governance controls and insights, self-service capabilities, APIs for integration with external DevOps tooling, success management with exhaustive technical documentation, training programs, and a comprehensive global partner network.

AILCAPs provide the foundation for developing a wide range of applications and application components with distributed data architectures, including complex multimodal front ends, business workflows, agentic AI, and integration capabilities.

The most popular use cases include:

- **Building and managing autonomous AI agents:** This is the primary use case for AILCAPs in 2026. Enterprises focus on building and managing AI agents at scale to create new generative AI solutions or augment existing applications. Requires robust agent implementation and strict governance across the entire AI agent life cycle
- **AI-augmented application modernization —** Professional developers use AILCAPs to modernize legacy systems, creating minimum viable products (MVPs) and new features via AI agents and vibe coding to augment or replace old business apps. Developers can use vibe coding and AI assistants to handle complex API integrations and back-end logic based on high-level intent, rather than manual line-by-line coding.
- **Intent-driven custom application development:** Building sophisticated web, mobile, and chat applications for both internal and external users. The focus has shifted from manual UI/UX building to "vibe coding." Instead of mapping every data field, developers use natural-language prompts to describe high-level intent, which AI code assistants then translate into functional, full-stack line-of-business applications.
- **Agentic workflow orchestration:** Automates workflows involving multiple application systems and human actors to accomplish business goals. It includes process modeling, orchestration, and decision modeling for both simple workflows and complex, document-oriented processes. AI agents now orchestrate background tasks and asynchronous data processing between internal domains and external APIs.
- **Autonomous governance guardrails:** With the rise of decentralized "citizen development," AILCAPs are now the primary tool for AI-powered governance. The platform continuously monitors all AI-generated code and agent actions for bugs, "AI slop," and compliance with regulations and policies.

Mandatory Features

The mandatory features for the enterprise AILCAP market include:

- AI agent development and life cycle management, including the ability to build, deploy, evaluate, and orchestrate autonomous AI agents that execute complex, multistep tasks and handle nonlinear reasoning.
- Visual development environments and model-driven IDEs that support minimal coding to create high-performance web, mobile, and chat interfaces for internal and external users.
- Agentic workflow orchestration to automate background tasks and complex end-to-end processes involving multiple application systems, external APIs, and human-in-the-loop interactions.
- Data virtualization and governance to support the creation and maintenance of complex distributed data architectures, integrating seamlessly with both internal and external enterprise data sources.
- Extensibility and legacy refactoring tools, providing component catalogs and SDKs that allow professional developers to modernize legacy code and extend platforms with custom code.
- Enterprise-grade security and automated AI governance, ensuring identity and access management (IAM), threat protection, and strict compliance enforcement for all AI-generated assets.
- Integrated runtime and self-service operations, offering high-performance environments with full documentation and tools for autonomous deployment, monitoring, and platform maintenance.

Optional Features

The optional features for this market include:

- Support for every stage of the software development life cycle, including automated deployment to multiple environments, environment-specific policies, and AI-driven observability for agentic behaviors.
- Intent-driven “vibe coding” and prompt-to-code capabilities, leveraging generative AI to translate natural language into functional application components, including UIs, logic, and

data structures.

- Collaborative development environments featuring visual versioning and real-time synchronization for fusion teams working in parallel on complex application logic and agent assets.
- Tools for the creation and maintenance of packaged business capabilities (PBCs), allowing shared, agent-ready APIs and components to be stored in a central catalog for enterprisewide reuse.
- Support for embedding prebuilt AI functional components, such as conversational chat interfaces and interactive AI agents, to provide intelligent business capabilities in end-user applications.
- An expansive library of prebuilt connectors for popular application platforms, niche legacy systems, and decentralized database management systems to ensure strong application and data connectivity.
- Internal implementation of design systems or deep integration with external systems to ensure high-fidelity “vibe-to-UI” consistency across web, mobile, and chat channels.
- Support for modern architectures, including event-driven streaming, microservices, and micro-frontends, to handle enhanced scalability and high-performance enterprise workloads.
- Inbuilt AI-powered test automation or seamless integration with external test management tools to validate generative code, complex workflows, and autonomous agent logic.
- Support for the development of native or “cross-platform-native” mobile apps that deploy to public or company internal app stores.
- Enablement of platform usage for B2B and independent software vendor (ISV) use cases, with support for client multitenancy and revenue-sharing agreements.
- Advanced process and decision modeling tools that allow for the orchestration of background tasks and asynchronous data processing for both simple and complex workflows.
- Dedicated prototyping tools that allow professional developers to rapidly create active prototypes and MVPs to validate new features before replacing legacy business applications.

- Automated reporting and data visualization tools, enabling the creation of management information system (MIS) dashboards through natural language prompts and autonomous data aggregation.

Magic Quadrant

Figure 1: Magic Quadrant for AI-Augmented Low-Code Application Platforms



Gartner

Vendor Strengths and Cautions

Appian

Appian is a Leader in this Magic Quadrant. Its product is the Appian Platform, an AI-augmented low-code application development and process orchestration platform solution delivered via Appian Cloud or self-managed deployment on Kubernetes. Appian operates globally through geographically diversified operations and serves enterprise and government clients, primarily in the public sector, financial services, insurance, and life sciences.

Appian embeds AI agents directly into its process and workflow engine, allowing agents to operate as governed steps inside business processes. Its AI design tool, Appian Composer, converts natural language requirements and documents into structured application plans that Dev Agents build into working Appian objects. Appian's future investment centers on this spec-driven approach, extending Dev Agents and adding MCP-based integrations that enable third-party tools to convert requirements into governed application changes.

Strengths

- **Product or service:** Appian unifies application development, process automation, and data fabric within a single platform. Its agent architecture applies existing user permissions and granular controls to agent actions and workflows. Clients can build applications and workflows that incorporate AI agents while retaining consistent security, data access, and governance controls.
- **Operations:** Appian offers managed cloud operations with service-level agreement options of up to 99.99% and supports high-availability configurations. The platform also supports self-managed, hybrid-cloud, and on-premises environments. Therefore, clients can choose between vendor-managed operations and greater infrastructure control based on availability, regulatory, and data-sovereignty requirements.
- **Overall viability:** Appian grew its cloud subscription revenue by 19% to $437.4 million in 2025 and improved its GAAP operating performance. This positive transition to profitability reinforces its business sustainability and recurring revenue model. Clients can rely on secure, long-term engineering investments in platform and cloud operations.

Cautions

- **Marketing execution:** Appian focuses its external positioning heavily on process automation rather than general application development. This narrow messaging underemphasizes capabilities in custom development and agent life cycle management.

Clients must request custom demonstrations that cover data modeling and testing to verify full platform utility.

- **Market responsiveness/record:** Self-managed deployments are only updated quarterly and some advanced AI capabilities remain restricted or unavailable in those deployments. Organizations must carefully assess feature parity and dependencies before choosing a self-managed deployment.

- **Sales execution/pricing:** Appian reserves capabilities such as AI agents and Composer for higher licensing tiers and includes defined AI Action allowances within those tiers. Clients should estimate production AI usage across applications and agents and clarify overage pricing, included capacity, and scaling assumptions before committing to a multiyear contract.

Creatio

Creatio is a Niche Player in this Magic Quadrant. Its product is the Creatio platform, comprising Creatio Business Studio and Creatio AI Studio, delivered as a managed cloud service or through customer-managed on-premises deployments. Headquartered in Boston, Creatio serves midsize and large organizations globally, with notable adoption in financial services, manufacturing, and public-sector organizations.

The platform combines composable no-code development, BPMN workflow automation, and CRM. Its AI App Development Toolkit supplies MCP tools that enable coding agents to generate and modify applications from natural language requirements. Creatio's future investment strategy centers on combining conversational AI-assisted development with its established no-code application development model.

Strengths

- **Product or service:** Creatio brings visual application development, process automation, and CRM on a single architecture. Business users build workflows visually, while technical teams extend solutions through APIs and AI coding agents. Clients can deploy cohesive operational systems without integrating separate products.

- **Offering strategy:** Creatio aligns application development, workflow automation, and agent life cycle management within a subscription framework. It bundles its Business Studio and AI Studio modules with an unlimited licensing option. Clients can scale their low-code adoption without facing unpredictable seat-based costs.

- **Customer experience:** Creatio delivers dedicated customer-success management, role-based academy learning, and standardized certifications. These resources support developers and business users, while its partner ecosystem assists with deployment. Clients can combine expert guidance and education to secure platform governance.

Cautions

- **Operations:** Creatio offers a vendor-managed cloud service on Amazon Web Services (AWS) and Microsoft Azure, alongside customer-managed cloud and on-site deployment options. In customer-managed deployments, clients assume responsibility for operating and maintaining the underlying infrastructure, including application, database, and caching components. Clients that select these deployment models should clarify responsibilities for availability, upgrades, backups, disaster recovery, and scaling, and assess whether they have the platform-engineering resources required to operate the environment.
- **Offering (product) strategy:** Creatio does not offer a free edition or free developer tier, so it lacks the no-cost onboarding entry point that many competitors use to let developers and citizen builders trial and adopt the platform before committing. Clients wanting to pilot broadly or seed grassroots adoption before a purchase decision may face a higher barrier to entry and slower internal buy-in.
- **Sales execution/pricing:** Creatio includes standard AI capabilities within its platform licensing, but it meters AI consumption through AI Actions. As usage scales, clients may find costs harder to predict if they do not establish consumption assumptions up front. Clients should model expected production usage and clarify included AI Actions, additional consumption pricing, and scaling terms before broader deployment.

Globant

Globant is a Niche Player in this Magic Quadrant. Its product is GeneXus, a model-driven low-code platform with AI-assisted capabilities provided via Glob.AI. Organizations install GeneXus as a development environment, and the platform packages applications for deployment on client-managed infrastructure, containers, and cloud platforms. Globant operates globally and serves large enterprises across multiple industries.

Globant's platform offers Nexa, an AI assistant that interprets natural language to create or modify GeneXus objects within an application's knowledge base before deterministic

generation. Globant's investment strategy centers on combining this model-driven development approach with AI agents and its subscription-based AI Pods model.

Strengths

- **Market understanding:** Globant helps enterprises accelerate with generative AI without sacrificing maintainability. Its Nexa assistant operates on structured objects, while the underlying knowledge base retains core business logic. Clients can modernize systems rapidly while retaining a fully governed, future-proof model.
- **Product or service:** Globant combines model-driven design with AI-assisted object creation and conversational controls. Its platform supports exposing .NET API objects as tools through an MCP server. Clients can build conventional business applications and embed agentic workflows within a single development model.
- **Marketing execution:** Globant aligns GeneXus with its broader enterprise AI portfolio, including AI Pods and industry-specific solutions. This unified approach provides access to enterprise buyers through both product and consulting engagements. Clients can utilize a structured AI delivery model to scale adoption across diverse sectors.

Cautions

- **Offering (product) strategy:** Globant distributes its capacities across multiple modular products and add-on suites. This fragmented structure can make licensing boundaries, functional entitlements, and life cycle responsibilities less apparent. Clients should map required capabilities to specific contracts before selecting the platform.
- **Sales execution/pricing:** Globant does not publish complete enterprise pricing for GeneXus, and instead relies on its distributor and partner network. Commercial proposals combine development licenses, professional services, and token-metered AI capacity. Clients should request a consolidated total-cost model before making multiyear commitments.
- **Operations:** Globant designs its cloud facility primarily for prototyping, while production deployments require client-managed runtimes. This architecture transfers SLA, scaling, and disaster recovery responsibilities directly to clients. Organizations should budget for infrastructure support when comparing vendors.

Huawei

Huawei did not respond to requests for supplemental information or to review the draft contents of this document. Gartner's analysis is therefore based on other credible sources.

Huawei is a Niche Player in this Magic Quadrant. Its product is Huawei Cloud Astro Zero, a vendor-managed low-code application platform. Huawei's operations focus mostly on the Asia/Pacific region, followed by the Middle East, Africa, and Latin America. Its clients are primarily in the government, finance, and manufacturing sectors.

The platform features native integration with industry-specific Pangu large models and can also connect to third-party and open-source models, such as DeepSeek and Alibaba Cloud Qwen, through the Astro connector based on client needs. The platform surfaces these models through an AI Assistant widget that supports knowledge-based grounding, natural-language-to-query capabilities, and AI-assisted generation of application components. Huawei's future investment strategy focuses on deepening the integration of Pangu models and agents across its cloud infrastructure stack.

Strengths

- Overall viability: Huawei Cloud, the cloud business unit of Huawei Technologies, operates through its wholly owned subsidiary, Huawei Cloud Computing Technology. It provides Astro Zero the backing of a large, well-resourced provider with sustained engineering and operational investment, assuring clients of continuity and long-term support behind their platform commitment.

- Operations: The underlying Huawei Cloud infrastructure provides the platform with broad compliance certifications, published service-level agreements, and multiregion resilience. These capabilities help clients secure high availability for critical applications and meet strict reliability and regulatory mandates.

- Innovation: Huawei differentiates Astro Zero by embedding industry-tuned Pangu models, AI that generates application components from natural language descriptions, and natural-language-to-query tools. Clients have access to AI capabilities aligned to specific verticals, such as government, finance, and manufacturing.

Cautions

- Offering (product) strategy: Huawei positions agentic capabilities in its separate Pangu and ModelArts platforms rather than building them natively into Astro Zero, which offers only a knowledge-base AI Assistant widget. As a result, the platform lacks native support

for capabilities such as visual agent workflow modeling, multistep agent orchestration, agent evaluation, and agent-specific governance and monitoring.

- **Product or service (testing):** Astro Zero lacks comprehensive native test automation and specialized front-end validation beyond basic sandbox testing. Organizations must source and integrate their own testing suites to verify accessibility and security when scaling Huawei deployments, which increases overall deployment costs.
- **Geographic strategy:** Global organizations operating outside China and the Asia/Pacific region may face limited regional availability, support, and reach. Astro Zero is available primarily in Huawei Cloud's China and Asia/Pacific regions rather than broadly across global regions, and its documentation is weighted toward China-region deployments.

Microsoft

Microsoft is a Leader in this Magic Quadrant. Its AI-augmented low-code application platform (AILCAP) offering is Microsoft Power Platform, which includes Microsoft Power Apps, Power Automate, Power Pages, Copilot Studio, and Dataverse, and is delivered as a Microsoft-managed cloud service offering. The platform offers multiple licensing options, including per-user, pay-as-you-go, and usage-based models.

Its services have AI capabilities for planning, building, and maintaining applications, automation, and AI solutions. Microsoft's most notable investments focus on Copilot Studio for building agents and AI automation, alongside the Power Platform admin center for AI governance. Key investments include a new GitHub Copilot harness in Copilot Studio, and FinOps and AI governance controls in the Power Platform admin center.

Strengths

- **Offering (product) strategy:** Microsoft's Power Platform admin center provides many governance controls within its managed-environments offering to address common governance risks, such as oversharing apps or agents, connecting to unapproved data sources, and the overconsumption of AI capacity. Clients can scale governance adherence through the use of platform controls.
- **Product or service (platform enablement):** Microsoft's AILCAP provides a strong set of development and governance capabilities for apps, automation, and AI agents. Although Microsoft licenses each product separately, the portfolio delivers strong technology

synergies that support the creation of more complex solutions. Clients benefit from products being aware of others, enabling invocation of agents from automation, and so on.

- **Innovation:** Microsoft continues to provide a large set of integration options for combining apps, automation, and AI agents with data across enterprise systems and APIs. These options have evolved to account for context data required to successfully build AI agents and automation. Clients benefit from having options for transaction, knowledge, and context-related integration requirements.

Cautions

- **Product or service (agent governance):** Fine-grained governance and admin controls, such as restricting access to specific features during agent development are inconsistent. For example, there is no way to limit what harness a user, group, or environment can use in Copilot Studio. Clients will need to invest in training and other tactics to address any missing fine-grained controls.
- **Market understanding:** Microsoft's Power Automate and Copilot Studio automation capabilities do not provide adequate features for more complex, long-running automations. These tools also struggle with automations that require a strong mixture of code-based execution within automation steps. Clients will need to rely on code-first, Microsoft options (such as Azure or Foundry) or use a multivendor strategy for more advanced automation requirements.
- **Business model:** Microsoft does not offer a single licensing option for Power Platform. In addition, inferencing must occur through Microsoft Azure, which prevents clients from reducing costs by directing products or solutions to models deployed on alternative platforms. Clients that scale AI consumption may encounter challenges with capacity planning and cost allocation that can contribute to overspending.

Oracle

Oracle is a Challenger in this Magic Quadrant. Its product is "Oracle APEX AI Application Generator", an enterprise-grade low-code application platform accessed via database entitlements or cloud-based subscriptions. Oracle APEX operations are geographically diversified, and its enterprise clients are from priority sectors, including financial services, government, manufacturing, and healthcare.

The platform recently introduced APEXlang, an application definition language that enables pages, workflows, and business rules to be managed as source-controllable specifications. Oracle's future investment strategy centers on governed generative development by extending APEXlang so that AI coding agents can generate, validate, and evolve source-controllable APEX applications under enterprise controls. Oracle also plans to add native DMN-compatible decision tables as a workflow capability.

Strengths

- **Product or service (DevOps):** Oracle APEX provides native environment separation and workspace isolation across complex estates. With APEXlang, released in APEX 26.1, AI-generated applications are produced as readable, source-controllable definitions that developers and AI agents can review and validate within a governed development workflow. This capability gives teams a governed path to AI-assisted generation without sacrificing version control or review discipline.

- **Customer experience:** Oracle backs APEX with mature onboarding, support services, and an extensive customer ecosystem. The platform draws on deep operational tooling to ensure reliable adoption pathways within established enterprise infrastructures. Clients benefit from a highly dependable and well-supported transition when adopting the low-code platform.

- **Operations:** APEX runs on Oracle AI Database and managed Autonomous infrastructure across OCI, multicloud, and customer-managed or on-premises environments. The platform provides high availability, automated backups, point-in-time restore, and mature workspace isolation and administration. Drawing on the widely available Oracle AI Database skills base, IT leaders gain enterprise-grade resilience and life cycle control without assembling separate systems.

Cautions

- **Sales execution/pricing:** Clients evaluating APEX as a low-cost option should model the full deployment spend because its total cost scales with the underlying Oracle AI Database and infrastructure consumption rather than being capped by APEX's own no-fee status. While APEX is a no-cost feature of Oracle AI Database and the separately priced APEX Service adds no per-user, developer, or workload fees, the database and infrastructure that APEX depends on still drive cost as usage grows.

- **Vertical/industry strategy:** APEX provides limited vertical orientation as a low-code platform, with fewer industry-specific accelerators, templates, and references than some competitors. Clients seeking prebuilt, domain-aligned capability within the low-code platform may need to build it themselves or look to Oracle's separate packaged applications instead of APEX.

- **Product or service (testing):** APEX does not provide a stand-alone native record-and-playback test automation suite and instead relies on external frameworks such as Selenium and Oracle's Application Testing Suite, although it supports governed, AI-assisted test generation and validation. Clients that want verification capabilities built natively into the platform should account for the required integration effort when planning AI-accelerated delivery.

OutSystems

OutSystems is a Leader in this Magic Quadrant. Its product is the OutSystems Developer Cloud (ODC), a low-code application platform delivered as an evergreen service across SaaS, hybrid, and on-premises environments. OutSystems' operations are geographically diversified, and its enterprise clients are primarily in financial services, professional services, and manufacturing.

The platform features an innovative plan mode for prompt-driven ("vibe coding") development that validates user intent before generating any artifact and allows manual edits to survive through iterative AI refinement. A key capability on its roadmap is workflow generation from AI/natural language, which it expects to launch in late 2026.

Strengths

- **Product or service (development capability):** OutSystems unifies visual, low-code, and pro-code development over a single governed application model that keeps AI-generated artifacts reviewable, version-controlled, and secure. Clients can scale AI-accelerated delivery without accruing maintainability or governance debt.

- **Market responsiveness/record:** OutSystems shows a strong record of tracking market shifts and shipping substantive product updates. This agility is illustrated by its native support for interoperability standards such as MCP. Clients benefit from a technology roadmap that keeps pace with the rapid advancement of AI.

- **Product or service (application security):** OutSystems applies permission-aware controls across AI retrieval, generation, and tool use, while incorporating prompt-injection mitigation and enterprise identity management for users and services. Clients gain assurance that AI-enabled features respect existing entitlements rather than bypassing them.

Cautions

- **Sales strategy:** OutSystems outlines a less differentiated growth and go-to-market plan compared with market competitors. Clients will find fewer signals about the vendor's trajectory and commercial commitment when building a multiyear platform case.
- **Customer experience:** OutSystems offers less comprehensive training, customer retention, and value-measurement programs than leading market competitors. These gaps in structured support require organizations to manage onboarding independently. Clients that scale adoption across large builder populations can invest heavily in internal enablement and success programs or rely on OutSystems partners for support.
- **Sales execution/pricing:** OutSystems offers weaker controls than competitor products to cap or enforce spending, such as budget thresholds that throttle or block usage before costs exceed defined limits. The platform also provides limited tooling to forecast how costs scale from pilot to production. Gartner also observes pricing and total-cost concerns in client inquiries about OutSystems as usage scales. Prospective buyers should model total cost of ownership carefully when comparing OutSystems with consumption-based alternatives.

Pega

Pega is a Visionary in this Magic Quadrant. Its offering is the Pega Infinity Platform, which serves as the low-code application development, workflow automation, and decisioning foundation for all its offerings. Pega Infinity 26 is delivered via Pega Cloud or client-managed cloud environments. Pega operates globally and serves large enterprises and government organizations, with a strong presence in financial services, insurance, healthcare, communications, and the public sector.

A key differentiator is Pega's workflow-governed agent architecture, which combines AI-driven actions with deterministic processes and orchestrates agents across applications through Pega Agentic Process Fabric. Pega Blueprint complements this approach by using AI

to analyze legacy assets and generate starting application designs. Pega's investment strategy focuses on governed agentic automation, AI-assisted modernization and cross-application orchestration.

Strengths

- **Sales strategy:** Pega incorporates its free Blueprint tool directly into sales and partner enablement processes. Its companion partner program measures contribution across co-sell engagements to incentivize collaborative designs. Clients can use these blueprints to validate requirements, align stakeholders before committing to development, and provide an initial buildout of a new application.
- **Product or service:** Pega combines low-code development, workflow automation, and real-time decisioning on a single platform. Its Predictable AI approach uses generative AI tools during design and guides runtime agents through structured business rules. Clients can safely apply AI to regulated processes while retaining full traceability.
- **Customer experience:** Pega delivers structured client support through dedicated success and service assurance programs. These resources cover value realization, technical-risk management, and role-based enablement through online and instructor-led training. Clients can combine vendor services and self-learning to establish the skills required for enterprise-scale implementations.

Cautions

- **Market responsiveness/record:** Pega supports open standards such as MCP and A2A and contributes open-source SDKs, deployment tooling, and data science utilities. However, these contributions are primarily tools and components rather than open AI models or an open application runtime. Clients prioritizing open-source application models and portability independent of Pega should assess dependencies on Pega-specific rules and runtime services before standardizing on the platform.
- **Sales execution/pricing:** Pega implementations can combine platform or Pega Cloud licensing with consulting, training, and other services, increasing the number of cost components that clients must consider when estimating total implementation costs. Clients should use scenario-based total-cost models that account for platform entitlements, implementation services, training, and production scaling.
- **Vertical/industry strategy:** Pega discontinued its manufacturing and transportation partner specializations for 2026 while retaining formal specializations in industries including

financial services, healthcare and life sciences, insurance, government, and communications and media. Clients in manufacturing and transportation should verify the availability of partners with recent Pega implementation experience, relevant industry solutions, and reference customers before committing.

Retool

Retool is a Visionary in this Magic Quadrant. Its product is the Retool platform, a low-code application platform delivered as Retool Cloud or Retool (self-hosted). Retool's operations are primarily in North America and Europe, and its clients range from startups to large global enterprises across technology, financial services, retail, manufacturing, and other industries.

The platform supports building and hosting applications, agents, and workflows, with full parity between its deployment options. Historically, Retool targeted professional development teams but has expanded support for citizen developers and fusion teams while actively evolving into an AI-centric platform.

Strengths

- **Product or service:** Retool provides multiple platform deployment options, offering its customers the flexibility to meet a range of privacy and compliance requirements, particularly those related to AI. Clients benefit from public, private, or hybrid deployment options to support data privacy and residency requirements.
- **Market understanding:** Retool differentiates itself through the ability to seamlessly mix code-based execution steps, model inferencing, and traditional workflow logic within AI and automation solutions. Clients can scale AI automation with more control and reliability by mixing deterministic and AI logic steps within solutions.
- **Business model:** Although Retool is evolving its pricing model for AI usage, it continues to offer persona-based licensing to help customers avoid licensing obstacles, particularly for end users. Clients benefit from a pricing model that combines established persona-based licensing with AI consumption-based pricing.

Cautions

- **Product or service (testing):** Retool's native testing and AI evaluation capabilities are limited and lack a unified framework. These AILCAP features are critical capabilities required to ensure applications, automation, and AI solutions are built for token efficiency,

reliability, and accuracy. Clients will need to implement and pay for third-party solutions to ensure the accuracy of solutions they build and maintain.

- **Marketing strategy:** Retool has developed a growing partner strategy centered on hyperscaler and data platform alliances, including AWS, Databricks and Snowflake, as well as system integrator partners; however, its ecosystem remains less extensive than those of some competitors' partner ecosystems. This limitation can reduce Retool's ability to compete with vendors that leverage alliances to expand both geographically and within vertical industries. Clients will need to make greater investments in internal enablement efforts.
- **Overall viability:** Retool's client distribution leans heavily toward startups and small and midsize organizations. Larger enterprise clients may find fewer features specifically designed to support AI governance at scale.

Salesforce

Salesforce is a Leader in this Magic Quadrant. Its AILCAP consists of Headless 360 and includes Salesforce Platform, Agentforce, Data 360, MuleSoft, Flow, and security services. With the launch of Headless Experience Layer, Salesforce offers a unique capability for rendering agent experiences across popular user channels such as Slack and Microsoft Teams. This capability distinguishes the platform from the "build anywhere" capabilities that most AILCAP leaders now provide.

Geographically diversified, Salesforce serves organizations of all sizes across all sectors. Historically, Salesforce has operated as a CRM vendor that provides turnkey solutions for sales, service, and other CRM subsegments. Salesforce continues to invest in more AI building capabilities to balance its portfolio between prebuilt applications and custom development tools.

Strengths

- **Marketing strategy:** Salesforce has a diverse and global footprint of strategic partners. The ecosystem includes vertical partners that offer implementation services, as well as AppExchange solutions for CRM products and Agentforce. Clients can leverage partner knowledge and solutions to accelerate the deployment of AI solutions.
- **Innovation:** Salesforce's Headless Experience Layer (HXL) provides a cohesive way to build AI agent experiences hosted in Agentforce that render across different UI surfaces, such as

Slack, Voice, WhatsApp, and Microsoft Teams. Clients will be able to focus on solution development because HXL simplifies the delivery of experiences across the channels where users already work.

- **Customer experience:** Salesforce's AppExchange marketplace and its training site called Trailhead distinguish the vendor from other AILCAP providers, both in quality and quantity of options available for customers to develop skills or find extensions. Clients can scale user and builder adoption by leveraging existing content and solutions.

Cautions

- **Offering (product) strategy:** Deploying Salesforce's AILCAP for large, multicloud enterprise scenarios can still require significant implementation effort and partner involvement. Although Salesforce has shortened the path to initial deployment through investments in agentic tooling, including the Mesh acquisition, prebuilt agent templates, and portable Git-based builder skills, organizations pursuing broader cross-functional rollouts should plan for additional coordination and implementation resources.
- **Sales strategy:** Clients may face challenges persuading technical teams to adopt Salesforce's AILCAP alongside or instead of existing development tools because Salesforce has traditionally excelled at selling to business buyers in domains such as sales, service, and marketing. As Salesforce expands into technical buyer segments, it must continue adapting its CRM-oriented sales approach to address the needs and preferences of those audiences.
- **Business model:** Salesforce has completed more acquisitions than many other AILCAP vendors. This acquisition strategy can increase risks related to product, platform, and licensing consistency. Clients may encounter pricing or product changes as Salesforce continues integrating acquired products into its AILCAP portfolio.

ServiceNow

ServiceNow is a Leader in this Magic Quadrant. Its product is the ServiceNow AI Platform, comprising App Engine, AI Agent Studio, AI Control Tower, Process Mining, ServiceNow Otto (formerly known as Now Assist) and Workflow Data Fabric, delivered via a multitenant cloud or customer/partner-operated private instances.

Geographically diversified, ServiceNow serves organizations of all sizes across all sectors. It sells its offerings under App Engine Prime for all capabilities with fixed assists, while Workflow

Data Fabric is sold as a separate SKU. While App Engine is its core low-code offering, ServiceNow provides embedded AI capabilities through ServiceNow Otto for Creator, which includes its build agent. AI Agent Studio enables custom agent development, AI Control Tower provides agent governance, and Workflow Data Fabric enables data connectivity and grounding of agents and workflows. Its future investment strategy focuses on launching Project Warpspeed, an upcoming intent-driven development environment sold outside standard licensing.

Strengths

- **Overall viability:** ServiceNow delivers strong financial stability, reporting 27.2% year-over-year growth and consistent profitability. Its high customer retention rate and the growth of ServiceNow Otto products to more than $600 million in contract value in 2025 indicate strong customer readiness and market-leading stickiness. Clients can benefit from its long-term commitments and viability of its services and products.

- **Customer experience:** ServiceNow offers 24/7 support under all subscription tiers, with response times varying by issue criticality, and schedules two family releases annually supported by monthly releases of AI developer tooling. It also fosters a large developer community through structured training, events such as CreatorCon, and comprehensive training programs via ServiceNow University. Clients benefit from dedicated support, simplified platform upgrades, and a large developer community.

- **Product or service:** The ServiceNow AI Platform delivers a comprehensive portfolio with strong capabilities in AI-augmented application development; agentic process workflow orchestration; and application and agent governance and security. Clients preferring AILCAPs can benefit from the platform's alignment with their use case.

Cautions

- **Vertical/industry strategy:** ServiceNow relies on global system integrators (GSIs) to deliver on certain vertical solutions. This model has resulted in limited evidence of deep industry customizability and depth. Clients looking for domain-specific solutions will be required to rely on their internal development teams or GSIs for customization.

- **Sales strategy:** ServiceNow announced limited growth in direct sales teams. Although it consolidated its portfolio under ServiceNow AI Platform, it continues to sell AI-augmented low-code features under App Engine-branded SKUs with fixed assists. This dual-brand structure can create confusion around commercial terms, licensing, and entitlements.

Clients may encounter challenges when negotiating contracts or evaluating available purchasing options.

- **Innovation:** ServiceNow aligns its product strategy with generic market shifts, prioritizing incremental feature updates over a differentiated vision. It also declined to disclose its research and development budget for their AILCAP product during the evaluation process. Clients seeking a vendor with a more distinctive innovation roadmap may find it difficult to assess the long-term competitive differentiation of the platform.

Siemens (Mendix)

Siemens is a Leader in this Magic Quadrant. Its core low-code offering is the Mendix platform, which supports AI-augmented development via its Mendix AI Assistant (MAIA) and custom agent creation via Agents Kit. Siemens recently positioned Mendix under the new brand Intelligence Center X, alongside Graph Studio and AI Studio from RapidMiner following the Altair acquisition, with each product sold as a separate SKU. Graph Studio is used for enterprise semantic knowledge graph and contextual grounding, while AI Studio enables teams to develop and deploy machine learning models.

Siemens maintains a geographically diversified presence and serves clients of all sizes across industries. Deployment options include platform as a service (PaaS), hybrid environments, AWS GovCloud, and air-gapped private cloud or on-premises deployments.

Strengths

- **Operations:** Mendix provides enterprise-grade uptime SLA commitments for premium cloud resource packs, maintains a strong uptime record, and offers priority support based on subscription tier. It also delivers broad compliance coverage, including FedRAMP, NIS2, SOC 2, and the ISO 27000 family of standards, and supports an annual major-release cadence alongside ongoing support for Mx11. Clients operating in highly regulated environments or requiring greater sovereignty can benefit from Mendix's flexible deployment and operational capabilities.

- **Product or service:** Mendix integrates with RapidMiner products to deliver strong capabilities in agentic workflow orchestration, application and data integration, and business workflow design. Clients can leverage these capabilities to support complex agentic process development initiatives.

- **Overall viability:** Siemens' ownership of Mendix provides strong organizational stability and long-term backing. Mendix has remained profitable since 2025 and maintains high customer retention driven by deployment flexibility and its focus on sovereignty requirements. Clients can benefit from the platform's long-term viability and Siemens' continued commitment to its products and services.

Cautions

- **Sales strategy:** Mendix projects modest sales team growth with double-digit revenue projection growth. However, the strategy lacks a contingency plan based on the market impact. Clients may struggle to verify dedicated product momentum when making multiyear investments.
- **Marketing execution:** Siemens has recently merged RapidMiner and Mendix into one brand name, with marketing plans still developing for Intelligence Center X. Clients may struggle to assess the vendor's future product direction, services strategy, and alignment with their own goals.
- **Business model:** Customers must purchase third-party products for automation testing and advanced security management, even if they buy the premium tiers. In addition, Siemens continues to evolve its consumption-based pricing model for AI capabilities. Clients may find it difficult to forecast future costs and evaluate long-term pricing structures.

Zoho

Zoho is a Niche Player in this Magic Quadrant. Its AILCAP offering is Zoho Creator sold as a unified, multitenant platform cloud service with on-premises and hybrid options.

Geographically, it serves a global client base and supports it through partners across North America, EMEA, and Asia/Pacific. The platform provides AI-augmented application development capabilities via its embedded AI assistant, Zia, while the embedded Zia Studio enables custom AI agent development. Zoho includes its AI-augmented testing solution, QEngine, with the Creator Enterprise plan, but customers must purchase it separately with lower-tier editions. Additionally, Zoho MCP allows coding agents, such as Claude Code, to access Creator's tools and applications at no additional cost.

Strengths

- **Vertical/industry strategy:** Zoho has a strong partner ecosystem specializing in specific domains. It delivers credible deployments across the manufacturing, education, and

professional services sectors with light configuration and customization. Clients looking for domain-specific solutions will benefit from the robust partner ecosystem.

- **Geographic strategy:** Zoho continues to build a stronger global client and partner ecosystem presence, supported by data centers across Asia, Europe, and North America. It also provides localized UI experiences and a global support network. Clients can benefit from local support capabilities and improved application performance.
- **Product or service:** Zoho Creator's user experience, powered by Zia and its build agent, offers easier AI-augmented application development. Its focus on expanding existing legacy systems with automated processes supports its credibility in a niche space. Development teams can benefit from a simpler learning curve and will be able to build low-code applications quickly.

Cautions

- **Market responsiveness/record:** Zoho's response to AI-augmented development needs is slow compared to leading products in this market. For example, its MCP product capabilities are still evolving with application development using AI coding agents. Clients seeking advanced AI-assisted development capabilities may encounter limitations that restrict certain use cases or slow adoption.
- **Marketing execution:** Zoho's marketing message lacks clear differentiation and aligns closely with broader AI-augmented development narratives. Its legacy modernization strategy focuses more on automating processes through integration with existing systems than on modernizing legacy applications. Clients may find it difficult to understand the scope and limitations of Zoho's modernization capabilities.
- **Sales strategy:** Zoho does not plan to expand its sales organization. Although the company differentiates itself through pricing and its proprietary Zoho LLM, customers that rely on third-party models must license those models separately. This approach can complicate total-cost calculations. Clients may find it difficult to accurately forecast AI-related costs across the platform and associated licensing models.

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

- Globant (GeneXus)
- Huawei

Dropped

- SAP

Inclusion and Exclusion Criteria

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

- Demonstrate a go-to-market strategy and software product offering for their AILCAP that supports cross-industry or general-purpose application development.
    - The AILCAP must be available as a platform for building cross-industry solutions. It must not be a platform for building single vertical/industry applications. It also must not be a product bundled within some other solution or platform.
    - The AILCAP software product must be publicly available, including technical documentation, self-service tooling for development and operations (SDKs), developer community, public pricing, and trial licensing models (usage limitations, freemium, etc.).
    - The AILCAP must support the development and deployment of applications in a distributed environment across org structures of IT and business departments.
    - The go-to-market strategy and software product offering is marketed as a “low-code” platform rather than exclusively as a “no-code” platform.
- Provide AI-augmented low-code capabilities that minimize or eliminate coding and code-maintenance tasks and include all mandatory features identified in the AILCAP market definition to:

    - Develop, version, test, deploy, execute, administer, monitor, and manage applications and their relevant artifacts.
    - Develop AI agents and manage their life cycle, including the ability to build, deploy, and orchestrate autonomous AI agents that execute complex multistep tasks and handle nonlinear reasoning.
    - Support agentic workflow orchestration to automate background tasks and complex end-to-end processes involving multiple application systems, external APIs, and human-in-the-loop interactions.
    - Provide features to design data schema and implement application logic.
    - Support composition of internal and external data sources, including internal databases and extendable connectors to external databases, APIs, and event channels.
    - Create rich application user interfaces for web and mobile channels.
    - Support change management over multiple environments, including development, test, staging and production.
    - Maintain a catalog (a marketplace) for sharing components, modules, connectors, and templates within the client's organization or across different clients.
- Deliver an enterprise-grade AILCAP that provides:
    - High availability and disaster recovery.
    - Secure access to applications.
    - Technical support to customers.
    - Third-party application access to application logic and/or data, via APIs and/or event topics.

In addition to the above market and technical criteria, each participating vendor must meet the following business criteria:

- **Size:** The vendor must, in calendar year 2025, fulfill one of the following size requirement combinations:

    - Generate at least $60 million in AILCAP license and/or subscription revenue during the previous year and support at least 100 paying enterprise customer organizations (with 1,000 or more employees) for its AILCAP offering, excluding revenue from related product offerings.
    - Generate at least $25 million in AILCAP license and/or subscription revenue during the previous year and support at least 5,000 paying enterprise customer organizations for its AILCAP offering, excluding revenue from related product offerings.
    - Maintain an AILCAP developer community of more than 100,000 developers across all customers.
- **International presence:** The vendor must have direct customers (rather than customers obtained through resellers) in three or more of the following regions:
    - North America
    - Europe
    - South America
    - Middle East and Africa
    - China
    - Japan/Asia/Pacific

Client relevance, as determined by analyst expertise and informed by public and proprietary information, is also considered. Evidence may include Gartner customer search, inquiry and pricing requests, frequency of mentions as a competitor to other vendors in this Magic Quadrant, and Gartner Peer Insights data.

Evaluation Criteria

Ability to Execute

Gartner's weightings for Ability to Execute emphasize product and sales execution/pricing criteria.

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

Source: Gartner (September 2026)

Completeness of Vision

Gartner's weightings for Completeness of Vision emphasize market understanding, offering/product strategy, and innovation criteria.

Completeness of Vision Evaluation Criteria	
Evaluation Criteria	***Weighting***
Market Understanding	High
Marketing Strategy	Medium

Evaluation Criteria	*Weighting*
Sales Strategy	Medium
Offering (Product) Strategy	Medium
Business Model	Medium
Vertical/Industry Strategy	Low
Innovation	High

Source: Gartner (September 2026)

Quadrant Descriptions

Leaders

AILCAP Leaders demonstrate strong execution, particularly in terms of business performanc and a strong vision (in terms of product and go-to-market strategies. These vendors stand o in a highly competitive global market and serve a wide range of organizations and applicatio.. use cases with their robust AILCAP offerings.

Challengers

AILCAP Challengers demonstrate strength in execution but lack the vision of Leaders, especially in offering and go-to-market strategies for broader use cases and markets. These vendors have shown strong execution in their respective focus areas and are expanding their customer base. However, they have not demonstrated the market understanding and strategic vision required to extend their offerings beyond their core customers and address a wider range of buyer needs and use cases.

Visionaries

AILCAP Visionaries demonstrate a good strategic vision and a strong understanding of the needs of the AILCAP market. However, these vendors have not yet demonstrated a strong

track record of executing their strategies and may fall behind competitors in the execution of their strategy or product capabilities.

Niche Players

AILCAP Niche Players focus on a specific market area or have a regional geographic footprint. Although they do not demonstrate the same level of Completeness of Vision or Ability to Execute as other evaluated vendors, qualifying for inclusion in this Magic Quadrant remains a significant achievement. Niche Players in the AILCAP market may also excel in adjacent technology markets and, in some cases, represent the most appropriate choice for specific application use cases.

Context

As AI reshapes how software is built, low-code application platforms are evolving based on their ability to embed AI generation, agentic workflows, and the governance capabilities required to control them. Enterprise buyers select low-code platforms based on whether AI-augmented delivery can be trusted, governed, and verified at scale, rather than visual development speed alone. This Magic Quadrant emphasizes providers' abilities to make AI-generated and agent-driven outputs verifiable, governable, and integrable into existing engineering and enterprise estates rather than focusing solely on agent selection or generation features.

Gartner's audience for this market includes enterprise technology leaders, such as CIOs, CTOs, enterprise architects, and software engineering and application leaders who are responsible for teams tasked to build and maintain applications at scale across multiple regions. Although many vendors evaluated in this Magic Quadrant also serve smaller organizations and single-region deployments, the evaluation places greater weight on the multiregional governance, security, and operational requirements that Gartner clients raise most consistently through inquiries.

Software engineering leaders should prioritize the capabilities that separate outcomes rather than those that have become table stakes. Over the past 12 months, AI generation has become broadly available, so vendors received higher evaluations when they delivered AI outputs that were reviewable, testable, and permission-aware by design; provided native, orchestrated, and governed agentic capabilities rather than relying on separate platforms; and

extended cost, security, and life cycle controls to AI consumption. Software engineering leaders should treat verification capabilities, agent governance, and integration with existing toolchains as key selection criteria. They should also validate vendor claims through proofs of concept because the gap between demonstrated AI capabilities and production-ready implementations remains significant.

Market Overview

AILCAPs combine visual, low-code, and professional-code development with AI capabilities that generate applications, automate delivery, and orchestrate autonomous agents. The market has evolved rapidly from classic enterprise low-code platforms, where the primary value centered on accelerating hand-coding through visual modeling, to a category where AI generation and agentic workflows play a central role in application design, development, and operation. This shift has raised enterprise expectations and shifted the basis of competition among vendors.

During the past 12 months, AI-assisted generation of applications, data models, logic and tests has become widely available across the market. Platform differentiation has shifted from code-generation speed to the verification, governance, and security controls surrounding AI output. As AI accelerates output, platforms that enable organizations to review, test and secure that output by design deliver more realized value than platforms that focus on optimizing generation speed alone. Governance has become equally decisive, particularly the ability to keep AI retrieval, generation, and agent actions within each user's existing entitlements and under enterprise policies. Providers that embed these controls directly into their platforms hold a stronger position than providers that leave them to customer configuration or rely on adjacent products.

Agentic capability is the market's most active frontier and its least mature. Vendors increasingly provide agent development, orchestration, and governance capabilities, but maturity and breadth vary significantly. In many cases, agentic tooling resides in a separate AI platform rather than within the low-code platform itself. Interoperability standards such as MCP are beginning to appear in production environments, and support for these standards is becoming an indicator of future alignment. Cost management for AI consumption, agent evaluation and red teaming, and production-grade agent observability remain inconsistent across the market and frequently contribute to buyer dissatisfaction when organizations move from pilot projects to production deployments.

Enterprise buyers should expect meaningful trade-offs. Some vendors lead in enforcing access controls at the data layer, ensuring that security policies apply to every request, including AI retrieval activities, while also providing strong resilience and data-connectivity capabilities. Other vendors differentiate through unified governed development, mobile experiences, or agentic orchestration capabilities, and few vendors lead across all these areas. Deployment flexibility, regional availability, and data residency continue to influence suitability for regulated and sovereignty-sensitive organizations. The vendors best positioned for the next phase of the market treat AI-augmented delivery as a software engineering discipline by applying verification, governance, and life cycle management to AI-generated outputs rather than treating AI as a stand-alone generation feature added to a low-code platform.

⊕ Evaluation Criteria Definitions

Unlock more AI value

With unmatched, proprietary AI insights from Gartner

Explore the AI Hub ↗

© 2026 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.

© 2026 Gartner, Inc. and/or its Affiliates. All Rights Reserved.