

Magic Quadrant for Cloud-Native Application Platforms

4 August 2025 - ID G00824519 - 50 min read

By Tigran Egiazarov, Mukul Saha, [and 1 more](#)

Cloud-native application platforms remove infrastructure management complexity and enable product teams to deliver cloud-native applications and AI agents. This research helps software engineering leaders evaluate cloud-native application platform vendors and find the best fit for their organization.

Market Definition/Description

Gartner defines cloud-native application platforms as those that provide managed application runtime environments for applications and integrated capabilities to manage the life cycle of an application or application component in the cloud environment. They typically enable distributed application deployments and support cloud-native operations — such as elasticity, multitenancy and self-service — without requiring the development team to provision infrastructure or manage containers.

Cloud-native application platforms are designed to facilitate the deployment, runtime execution and management of modern cloud-native or cloud-optimized applications without the need to manage any underlying infrastructure. Also, they are designed to enhance developer productivity, accelerate development and deployment cycles, and increase operational effectiveness by making it easier to scale on demand.

Cloud-native application platforms offer a structured execution environment for applications, effectively hiding the complexities of the underlying infrastructure and computing resources. They also provide vendor-supported versions of application runtimes and frameworks for the commonly used languages (for example, Java, .NET, Node.js, PHP, Python, Go and Ruby). By abstracting the complexities associated with infrastructure

management, cloud-native application platforms enable product teams to deliver faster customer value.

The cloud-native application platforms market reflects the consolidation of technologies across deployment, scalability, security and application observability to streamline software delivery. They are intended to be more than just a platform for running applications; they are essential for businesses aiming to achieve excellence in software engineering, productivity and market responsiveness.

Typical cloud-native application platform benefits include:

- **Operational excellence:** Cloud-native application platforms remove infrastructure management complexities, allowing organizations to focus on innovation and core business goals, thus improving efficiency.
- **Easier to scale:** Cloud-native application platforms ensure applications can scale dynamically to meet demand with minimal manual intervention by using automation and providing seamless performance, even during peak loads, thus enhancing reliability and user experience.

Mandatory Features

The mandatory features for this market include:

- Application runtime services (including language runtime support) for multiple application types including web applications, mobile back ends, microservices, AI/ML models and analytics applications without requiring infrastructure provisioning or creating and maintaining custom container images.
- Automated deployment of cloud-native applications (e.g., integration with DevOps).
- Autoscaling (load balancing, scalability and running of multiple instances).
- Application monitoring and observability — Support for monitoring and observability to improve service-level objectives; gathering production telemetry (logs, metrics, events, traces).
- Fully managed service — Vendor (service provider) handles the maintenance, monitoring, updates and troubleshooting of the cloud-native application platform. This includes support, security, backups and performance optimization. It allows users to focus only on the application that can be deployed on cloud-native application platforms.

Common Features

The common features for this market include:

- Ability to deploy, manage, configure and operate containers at scale.
- Financial management capabilities for effective cost control and optimization.
- AI-assisted runtime environment. This includes intelligent configuration and orchestration of services, along with efficient resource distribution across workloads.
- IDE extensions and development tools to support software engineering teams in building applications for cloud-native application platforms.
- Serverless computing, which eliminates the need to manage application or service instances and their allocated compute resource, automatically scales with application demand and charges based on compute time used, enhancing efficiency and cost-effectiveness. Serverless computing encompasses various models, including functions as a service (FaaS) and serverless container orchestration.
- Automatic updates and security patches to keep cloud-native application platforms up to date and secure without service disruption, reducing vulnerability risks and maintenance effort.
- Polyglot deployment supports multiple programming languages and frameworks.
- Ability to easily integrate with external services such as DBMS, event brokers and caches through standard APIs.
- High availability and disaster recovery, which ensures high availability by automatically spinning off or switching to another application instance or switching to another compute region if the native host region goes down. This includes, but is not limited to, data backup and automatic failover, enhancing application reliability and continuity.

Magic Quadrant

Figure 1: Magic Quadrant for Cloud-Native Application Platforms





Gartner.

Vendor Strengths and Cautions

Alibaba Cloud

Alibaba Cloud is a Leader in this Magic Quadrant. It offers Serverless App Engine (SAE) to support modern web applications and legacy cloud migration; Function Compute to support serverless functions; and Container Compute Service (ACS), a serverless platform to support containers. It also offers other related cloud platform services for API management, AI application development, integration, identity management, observability services and data management.

Alibaba Cloud's operations are mainly in China, and it caters to customers of all sizes across all sectors. The company is investing in its AI capabilities, its AI developer tools, and compute performance optimization within the Alibaba Cloud ecosystem.

Strengths

- **Product or service:** Alibaba Cloud provides a comprehensive platform for modern application development. Developers can build and deploy modern web applications using custom containers or prebuilt templates on Container Compute Service. Alibaba Cloud also enables the creation of AI agents and intelligent applications by integrating Function Compute, SAE and advanced AI services such as Model Studio and Platform for AI.
- **Innovation:** Alibaba Cloud's platform offers innovative features that help customers optimize compute cost and performance. For AI workloads, its serverless GPU delivers pay-as-you-go pricing, idle billing, fine-grained allocation and millisecond-level snapshots. The company has also launched one-click AI application templates with ModelScope and Hugging Face integration, and introduced an AI gateway for unified, multivendor LLM orchestration and API key management.
- **Market understanding:** Alibaba shows awareness of the latest trends with upcoming launches focused on AI-driven platforms, serverless container solutions and advanced security tools. Their offerings address the top needs of software engineers — scalability, seamless AI/ML integration and efficiency — while supporting modernization and greenfield development. With robust AI development tooling, serverless orchestration and real-time processing, Alibaba empowers teams to rapidly build, deploy and scale intelligent AI agents and applications.

Cautions

- **Sales execution:** Alibaba Cloud relies mainly on direct sales rather than sales partners or third-party marketplaces. While this approach helps acquire clients in China, it has limited its ability to reach clients who work with third-party system integrators.
- **Geographic strategy:** Alibaba Cloud primarily operates in China and is mostly focused on the Chinese market. While it has started expanding into the U.S., Europe and South America, customers in these regions should account for increased trade barriers, policy uncertainty and geopolitical risks when evaluating Alibaba Cloud.

- **Offering (product) strategy:** Alibaba Cloud's platform lacks natively provided application connectors for enterprise systems (e.g., Microsoft, Oracle, Salesforce, SAP). These connector integrations primarily rely on third-party solutions like Dapr.

Amazon Web Services

Amazon Web Services (AWS) is a Leader in this Magic Quadrant. It offers AWS Lambda, AWS App Runner, AWS Amplify and AWS Elastic Beanstalk. AWS also offers other related cloud platform services for API management, AI application development, integration, identity management, observability services and data management. These services support serverless functions, native code, containers and cloud-migration use cases.

AWS's operations are geographically diversified and cater to customers of all sizes across all sectors. The company is investing in AI capabilities, continuing to strengthen its AI developer tools and compute performance optimization within the AWS ecosystem.

Strengths

- **Product or service:** AWS's cloud-native application platform offers flexible options for modern application development. Options include AWS Lambda for serverless compute, AWS App Runner for containerized web applications, AWS Amplify for a set of tools and services for front-end web and mobile applications, and AWS Elastic Beanstalk to support cloud migration. AWS enables users to build AI agents and applications by integrating compute offerings with Amazon Bedrock.
- **Innovation:** AWS recently introduced Lambda SnapStart support for Python and .NET, which significantly reduces cold start times for its serverless platforms without additional resource provisioning. In the AI space, AWS is empowering developers to build and scale intelligent applications with offerings like Amazon Bedrock. AWS supports both CPU- and GPU-based workloads, ensuring robust performance and scalability for diverse AI applications.
- **Market understanding:** AWS shows awareness of the latest trends in AI development by supporting major languages and frameworks and enabling deployment of functions, containers or native code. This broad range of compute choices, combined with native integration of services like Amazon SageMaker (for AI/ML training and management) and Amazon Bedrock (for LLM inference), empowers software engineering teams to efficiently build, deploy and manage AI agents.

Cautions

- **Platform complexity:** AWS provides numerous services with overlapping features. Customers need substantial knowledge of AWS platform services in order to select the right service for their needs. For example, deciding between Amazon ECS and AWS Fargate or AWS Lambda can be confusing, as each service is designed for different use cases and operational models, but these distinctions are not readily apparent.
- **Operations:** AWS does not offer formal end-user latency guarantees or negotiate higher uptime SLAs beyond what is publicly stated (e.g., 99.95% for Lambda). Users should verify AWS can meet their uptime and latency requirements.
- **Product strategy:** AWS Lambda is AWS-only technology, so AWS Lambda customers will need to reengineer their applications before replatforming applications to other cloud providers. However, AWS offers an extensive range of cloud platform services that may reduce the need to use multiple clouds.

Cloudflare

Cloudflare is a Challenger in this Magic Quadrant. It offers the Cloudflare Developer Platform, which includes Pages, Workers and Workers AI services within Cloudflare's connectivity cloud. Cloudflare enables users to build serverless applications, which are supported by AI inference services, and deploy them to more than 330 globally distributed points of presence in Cloudflare's serverless infrastructure.

Cloudflare's operations are geographically diversified, and Cloudflare Developer Platform clients tend to be midsize enterprises across various sectors. Cloudflare is working to expand its Developer Platform capabilities to support creation of AI-powered apps.

Strengths

- **Sales execution/pricing:** Cloudflare's innovative CPU-based pricing model ensures customers are only charged for active compute time. For example, when agents are waiting for responses from LLMs or other APIs, the customer is not billed. This pricing model offers a route to cost-efficiency for AI and event-driven workloads.
- **Geographic strategy:** Cloudflare has established a strong market presence by leveraging its vast network of edge locations to deliver low-latency, high-performance compute services across the globe. Its robust security offerings, including DDoS protection, web

application firewall (WAF), bot management, API security, and zero-trust access (SSE/SASE), position it as a trusted leader in secure, scalable internet infrastructure.

- **Innovation:** Cloudflare has rapidly expanded its platform with Workers AI, enabling GPU-powered machine learning and LLMs at the edge. In 3Q25, it will broaden access to its container platform, which is already in production with select enterprises, to support more containerized workloads. Workers is highly optimized for JavaScript, offering near-instant cold starts and efficient edge execution.

Cautions

- **Product or service:** Cloudflare Developer Platform lacks native support for Java and .NET, while offering native support for JavaScript, Python, Rust and WebAssembly. This limits its suitability for organizations seeking to lift and shift legacy enterprise applications or continue building cloud-native applications on these mature, well-known platforms. However, this is being addressed with Cloudflare's container platform, which was released in June 2025 and is currently available in Beta.
- **Offering (product) strategy:** Cloudflare does not yet provide some advanced features and integrations that are commonly offered by the Leaders in this research. The platform lacks a comprehensive marketplace for prebuilt connectors and deep SaaS or legacy system integrations, which may require additional effort from customers.
- **Market understanding:** Cloudflare's focus on edge computing and serverless architectures aligns with modern development trends, but many organizations still rely on legacy systems that are difficult to migrate to edge-native environments. Cloudflare's limited support for legacy workloads may reduce its appeal for organizations that need this support.

Google

Google is a Leader in this Magic Quadrant. It offers Google Cloud Run, a serverless platform to support functions and containers; Firebase to support modern web, mobile and full-stack applications; and Google App Engine to support cloud migration. It also offers numerous services to support API management, AI application development, integration, identity management, observability services and data management. Google supports most languages and back-end frameworks, and Google Cloud buildpacks simplify the creation of containers.

Google's operations are geographically diversified, and it caters to customers of all sizes across all sectors. The company is investing in strengthening its developer tools and optimizing compute performance within its ecosystem.

Strengths

- **Product or service:** Google's cloud-native application platform offers flexible options for modern application development, including Cloud Run for both serverless compute and containerized workloads, and Firebase for modern web, mobile and full-stack applications. These offerings also enable teams to build, run and orchestrate AI agents and applications by integrating with Vertex AI (for model training and inference) and Gemini (for AI developer tools). Google App Engine supports a more traditional three-tier architecture and provides a platform for lift-and-shift migrations.
- **Market understanding:** Google Cloud Run provides developers with flexibility and choice by supporting multiple programming languages and offering tailored buildpack support. This helps avoid vendor lock-in and enables migration of applications across any cloud environment that supports containers.
- **Innovation:** Google continues to deliver platform innovations such as Cloud Run GPU for scalable AI/ML inference; Eventarc for event-driven architectures, with over 150 native event sources; and Log Analytics for powerful, SQL-based analysis of telemetry data at scale. Additionally, Gemini Cloud Assist brings AI-driven assistance to both development and operations, helping customers accelerate innovation and streamline application life cycle management.

Cautions

- **Offering (product) strategy:** Google's workflows are more prescriptive compared to other leading cloud platforms, due to its highly integrated product portfolio. Customers that are used to building complex, customized solutions by combining diverse cloud services may find Google's opinionated approach somewhat restrictive.
- **Operations:** Google does not offer formal end-user latency guarantees or negotiate higher uptime SLAs beyond what is publicly stated (e.g., 99.95% for Cloud Run). Customers and prospects should verify Google can meet their uptime and latency requirements.
- **Geographic strategy:** Google currently lacks dedicated data centers and cloud availability zones within China, which hinders its ability to offer local support and meet

data residency requirements for Chinese customers.

Huawei

Huawei is a Challenger in this Magic Quadrant. It offers Cloud Application Engine (CAE), ServiceStage and FunctionGraph, which provide a platform for modern application development and a runtime environment for both web apps and back-end services. Huawei's cloud-native application platform supports a wide array of languages, back-end frameworks, and deployment of containers and native code.

Huawei's operations are mainly in China, with some operations in Southeast Asia, Latin America and EMEA. Its CAE clients tend to be midsize enterprises across various sectors. Huawei is expanding its service portfolio and integrating AI capabilities into its platform to enhance the developer experience and operational efficiency.

Strengths

- **Innovation:** Huawei has launched WebAssembly-based serverless functions on CDN edge nodes, enabling fast, secure and lightweight edge deployments, and it added native compute for LLMs. Huawei also introduced TaurusDB, a MySQL-compatible, cloud-native database with decoupled compute and storage that supports millions of queries per second for demanding applications.
- **Overall viability:** Huawei continues to reinvest at least 23% of its revenue into research and development. This substantial investment supports the continuous development of its cloud services portfolio. Particular emphasis is placed on enhancing CAE as a serverless application platform, ServiceStage as a platform engineering hub, and FunctionGraph for cloud plus edge hybrid functions through both cloud and CDN-based deployment models.
- **Market understanding:** Recognizing the industry's shift toward AI agent workflows, Huawei has extended its CAE and FunctionGraph services to natively support AI workloads, including those deployed at the edge. This positions Huawei to address emerging customer needs in AI-driven and edge-native application scenarios.

Cautions

- **Product strategy:** Huawei's integrated product portfolio has prescriptive workflows that can be more restrictive than those of other platforms in this research. While Huawei's

opinionated platform benefits organizations seeking efficiency, customers wanting maximum flexibility for custom solutions may find Huawei's approach is less adaptable.

- **Geographic strategy:** Huawei is mostly focused on the Chinese market. While Huawei has started expanding into Europe and South America, economic sanctions have restricted its access to the U.S. and some other markets. Customers should ensure Huawei can support deployments in their region and verify the maturity of their local service support.
- **Operations:** The platform currently has limited visibility and influence within the developer community, which may impact engagement adoption and ecosystem growth. Customers and prospects should assess whether the platform's current level of community support and ecosystem maturity aligns with their needs.

Microsoft

Microsoft is a Leader in this Magic Quadrant. It offers Azure App Service to support cloud migration, Azure Functions and Azure Container Apps to support functions and containers, and Azure Static Web Apps to support multiple web frameworks and languages. Microsoft also offers other related cloud platform services for API management, AI application development, integration, identity management, observability services and data management.

Microsoft's operations are geographically distributed, and it caters to customers of all sizes across all sectors. The company is investing in AI inference capabilities and an AI-native developer experience, and optimization of compute performance and scalability.

Strengths

- **Product or service:** Microsoft provides flexible options for modern application development, including Azure Container Apps for serverless containers and microservices, Azure Functions for serverless event-driven computers, and Azure Static Web Apps to build modern web applications. Microsoft enables users to build AI agents and applications by integrating Azure Container Apps and Azure Functions with Azure AI Foundry.
- **Innovation:** Microsoft continues to deliver innovative features, especially in Azure Container Apps. Its dynamic sessions provide secure access to sandboxed environments for running code separately from other applications, and its serverless GPUs support on-demand AI and ML workloads without infrastructure management. Azure Container Apps also has native integration with Azure AI Foundry for AI model inference.

- **Market understanding:** Microsoft's focus on AI development reflects a strong awareness of current market trends. With Azure AI Foundry for agentic workflows, customers can efficiently build and deploy AI agents and applications.

Cautions

- **Operations:** Microsoft does not offer formal end-user latency guarantees or negotiate higher uptime SLAs beyond what is publicly stated (e.g., 99.95% for Azure Container Apps). Customers and prospects should verify that Microsoft can meet their uptime and latency requirements.
- **Platform complexity:** Microsoft provides numerous services with overlapping features. Customers will need substantial knowledge of Azure platform services in order to select the right service for their needs. For example, deciding between Container Apps and Azure Kubernetes Service (AKS), or running containers on Azure Functions can be confusing, as each service is designed for different use cases and operational models. But these distinctions are not readily apparent.
- **Product strategy:** Azure Functions is optimized to run on Microsoft Azure, so its customers will need to reengineer their applications if they want to replatform applications with other cloud providers. However, Microsoft offers an extensive range of cloud platform services that may reduce the need to use multiple clouds.

Netlify

Netlify is a Niche Player in this Magic Quadrant. It offers Netlify Core for building and deploying web applications, Netlify Create for visual editing of web content and Netlify Connect for unifying application content. Netlify offers a wide variety of front-end web application frameworks and templates, which are part of its composable web platform, to support delivery of modern web applications. Netlify also offers other related cloud platform services such as Blob storage, KV storage and durable edge caching designed to give developers more flexibility across front-end architectures.

Netlify's operations are mainly in the U.S. and Europe. Its clients tend to be small and midsize businesses across various sectors. The company is investing in simplifying its development processes to improve the developer experience.

Strengths

- **Product or service:** Netlify Core delivers AI-enabled deploy assist and durable edge cache, Netlify Connect provides a federated data layer with a scalable GraphQL API, and Netlify Create offers visual/AI-assisted content editing. Netlify also offers website visual editing and it supports and promotes a composable-web approach (formerly Jamstack), a modern front-end architecture style.
- **Market understanding:** Netlify's strategic direction aligns with the evolution of the cloud-native platform market. It plans to continue evolving its platform to enable a wider range of developer personas that manage intelligent, agent-augmented workflows, with an emphasis on Agent Experience (AX) to simplify orchestration and abstract infrastructure complexity. Netlify is investing in observability tools for AI-native applications, which provide insights into agent behavior, traffic and performance.
- **Sales execution/pricing:** Netlify offers flexible integration packages so users can build applications quickly. Netlify uses a tiered pricing model with free Pro and Enterprise plans, so prospective clients can try the platform at low to no cost and then convert to the best tier for their organization's needs. Netlify is moving beyond traditional infrastructure-based pricing to value-based packaged models that can better suit rapid prototyping.

Cautions

- **Offering (product) strategy:** Netlify is not designed to host containerized microservices architectures built in Java or .NET. However, it offers back-end capabilities, such as serverless functions and compatibility with frameworks (e.g., Next.js, Astro or Remix), that support modern back-end use cases for front-end-adjacent applications and AI-native workloads.
- **Marketing strategy:** Platform visibility remains a challenge; the perception persists that Netlify is only suited for static sites or front-end developers, limiting engagement with enterprise platform engineering teams. To overcome this, Netlify is strengthening community engagement through increased open-source contributions, participation in developer events and the publication of technical blogs.
- **Overall viability:** Netlify is a venture-backed vendor in a market with numerous large, well-funded competitors. While smaller than some of its enterprise peers, Netlify has raised over \$200 million in venture funding. It will need to continue growing its market share

and expand into new regions like South America, Asia/Pacific (including China) and the Middle East to remain viable in the market.

Platform.sh (Upsun)

Platform.sh is a Niche Player in this Magic Quadrant. It offers Upsun, a developer-focused platform that supports a wide variety of application frameworks using a Git-based deployment process. Upsun also offers managed services for databases, caching and search with integrated application performance monitoring. Platform.sh supports multicloud deployments across AWS, Azure, Google Cloud, IBM and OVHcloud.

Platform.sh's operations are mainly in North America, Europe and Asia/Pacific. Its clients are typically midsize web agencies and line-of-business departments within large enterprises across a wide range of sectors. The company is investing in financial incentives for deploying projects in low-carbon-emission regions.

Strengths

- **Product or service:** Platform.sh offers unparalleled flexibility, offering customers the choice to use multiple clouds across multiple regions. Upsun is a self-service platform that supports scalable architecture across all major cloud providers. It also enhances multicloud portability by minimizing vendor lock-in.
- **Market understanding:** Platform.sh takes a developer-focused approach by offering a Git-centric solution with a high degree of flexibility. Upsun includes full-stack project cloning of all application code, data and services, which enables customers to replicate environments with improved efficiency. It enables teams to efficiently scale individual applications, optimizing resource use without impacting overall system performance.
- **Marketing strategy:** Platform.sh has positioned itself as a champion of sustainability and cost-effectiveness, which resonates with many organizations. Its carbon footprint modeling reports help customers understand the impact of their deployments. Upsun offers financial incentives for deploying projects in low-carbon-emission regions through its Greener Region Discounts program.

Cautions

- **Innovation:** Platform.sh's product roadmap includes limited AI innovations such as a MCP server to enhance its platform capabilities. The company's main focus is on improving the resilience of its existing platform and its LXC-based container runtime capabilities.

- **Marketing execution:** Platform.sh faces the ongoing challenge of demonstrating that its integrated, end-to-end multicloud platform provides greater value and simplicity than custom solutions built from disparate tools.
- **Geographic strategy:** Platform.sh primarily targets customers in the U.S. and Europe, with some customers in Asia/Pacific. It has limited implementation and hosting support in other regions. Prospective customers should verify if Platform.sh can effectively serve organizations in their region.

Red Hat

Red Hat is a Leader in this Magic Quadrant. It offers Red Hat OpenShift cloud services with a number of different hyperscalers: Red Hat OpenShift Service on AWS (ROSA); Microsoft Azure Red Hat OpenShift (ARO); and Red Hat OpenShift Dedicated on Google Cloud Platform. Red Hat OpenShift is also offered on IBM Cloud, and it can also be deployed as self-managed on-premises software.

Red Hat's operations are geographically diversified, and it caters to customers of all sizes and across all sectors. The company is investing in expanding support for AI workloads, optimizing costs with features like Hosted Control Planes, and broadening regional availability and instance types.

Strengths

- **Offering (product) strategy:** Red Hat offers customers high flexibility on where to run the platform; relatively easy portability across cloud providers and on-premises platforms; and low-friction, integrated services across its own platform. OpenShift provides multiple managed-language runtimes and DevOps tooling to support the deployment of functions, containers and native code.
- **Market understanding:** Red Hat's flexible multicloud strategy shows a recognition that the cloud-native market is rapidly evolving to support AI-powered applications. The continued evolution of Kubernetes into a comprehensive application platform is a key part of its transformation. Red Hat has also recognized the need to improve developer productivity by integrating IBM watsonx Code Assistant with Red Hat Ansible Lightspeed platform.
- **Business model:** Red Hat has forged strong partnerships with AWS, Google Cloud and Microsoft Azure to jointly engineer and operate OpenShift. This allows customers to

purchase directly from the hyperscaler, receive a single bill and leverage cloud committed spend. Red Hat also curates a broad platform ecosystem and marketplace of independent software vendors, system integrators and industry partners where customers can purchase OpenShift.

Cautions

- **Sales strategy:** Red Hat partners with hyperscale cloud providers such as AWS, Google Cloud and Microsoft Azure, but also competes with them. This dynamic encourages customers to seek out the best solution for their needs, choosing between Red Hat's cloud-native application platform and the solutions offered directly by these cloud providers.
- **Product or service:** Red Hat's managed service offerings are not as well-integrated with the hyperscalers' cloud-native services as the hyperscalers' own PaaS offerings. For instance, ARO's integration with Microsoft Azure's cloud services is not as strong as that of Azure App Service, Azure Functions and Azure Container Apps.
- **Marketing execution:** Red Hat faces the ongoing challenge of demonstrating that its integrated, end-to-end multicloud platform provides greater value and simplicity than custom solutions built from disparate tools.

Render

Render is a Niche Player in this Magic Quadrant. It offers a cloud-native application platform that automates the operational role and enables developers to focus on building code rather than setting up cloud infrastructure. Render supports Docker containers and infrastructure as code (IaC) functionality.

Render's operations are mainly in North America, and its clients tend to be small and midsize businesses across various sectors. The company is investing in enterprise priorities such as compliance and certifications, extending its geographic coverage and focusing more on developer experience.

Strengths

- **Market understanding:** Render aligns well with the needs of small and midsize businesses. Its well-integrated platform enables customers to deliver cloud applications and an online experience with autoscaling and high availability. Render's product roadmap shows a clear understanding of the market's direction, including features like AI-

augmented debugging, durable workflow engines for AI agents and GPU compute for advanced applications.

- **Offering (product) strategy:** Render's Blueprint provides native IaC functionality that automatically redeploys any affected services to apply the new configuration. This eliminates the need for users to understand the infrastructure. Customers benefit from this simplified developer experience.
- **Market responsiveness:** Render has delivered new features in the last year that directly address top customer issues. These include unified account management, OpenTelemetry stream to other observability tools and HIPAA-compliant accounts.

Cautions

- **Innovation:** Render lacks AI capabilities that are commonly available in competitors' platforms, such as access to LLMs or a runtime environment that supports AI agents. Render plans to address these feature gaps in its product roadmap.
- **Geographic strategy:** Render has few clients outside the U.S. and Europe, although it plans to expand into Asia/Pacific. While the company currently offers support in English, it provides international coverage and infrastructure capabilities suitable for global deployments.
- **Overall viability:** Render is a venture-backed vendor in a market with numerous large, well-funded competitors. While smaller than some of its enterprise peers, Render has raised over \$157 million in venture funding. It will need to grow its market share, expand into new regions and deliver new capabilities to remain viable.

Salesforce (Heroku)

Salesforce (Heroku) is a Leader in this Magic Quadrant. It offers Heroku, a platform that supports native code deployment using buildpacks and container deployment. Heroku also offers other related cloud platform services for data management, observability and identity management, as well as additional services through Salesforce for integration platform/api management (MuleSoft), AI agents (Agentforce), and data management (Data Cloud).

Heroku operations are geographically diversified, and it has clients of all sizes across various sectors. Heroku continues to focus on cloud-native and polyglot app development while deepening integration with the Salesforce ecosystem to enhance SaaS capabilities.

Strengths

- **Offering (product) strategy:** Heroku simplifies custom AI application deployment through Heroku AI by enabling additional Agentforce use cases via integration, and it improves scalability and reliability with a Kubernetes-based architecture abstracted from users. The addition of .NET and Jupyter Notebook support broadens its appeal to enterprise developers, while tighter integration with Salesforce Data Cloud and Agentforce via AppLink enhances their functionality. Heroku Connect also allows seamless access to real-time customer data.
- **Customer experience:** Heroku has continued to invest in improving its usability, expanding educational content and enhancing support. Developers benefit from updated documentation, new AI-focused tutorials and streamlined onboarding for emerging features like Heroku AI and .NET support.
- **Marketing strategy:** Heroku has expanded its marketing approach to align with its broader Salesforce AI strategy in addition to developer marketing. Its messaging focuses on Heroku's AI innovation, reduced operational complexity and increased developer productivity. This approach strengthens Heroku's credibility with business and technology decision makers.

Cautions

- **Product or service:** Heroku does not have built-in, dedicated support for serverless functions. While developers can mimic serverless behavior using lightweight dynos or external tools, this is not as efficient, scalable or cost-effective as true serverless architectures. Thus, Heroku is better-suited for long-running applications than event-driven workloads.
- **Sales strategy:** While part of Salesforce, Heroku's value is not always effectively integrated or communicated within broader Salesforce deals, leading to missed cross-sell opportunities.
- **Sales execution/pricing:** While pricing for Heroku is simple upfront, its fixed pricing per dyno offers less flexibility. The lack of granular usage-based billing can make optimizing costs for complex microservices architectures harder for teams, particularly compared to platforms with more advanced billing and monitoring.

Vercel

Vercel is a Visionary in this Magic Quadrant. It offers a web-focused platform with support for multiple web frameworks and languages. Vercel's platform offers edge compute, which includes a CDN, WAF, global runtime and caching functionality. Vercel also supports automated deployment and autoscaling capabilities.

Vercel's operations are mainly in the U.S. and Europe, and its clients tend to be small and midsize businesses across various sectors. Vercel is investing in front-end-centric innovations augmented by AI capabilities to improve web application performance, architecture and developer experience, and to increase the stability and scalability of its platform.

Strengths

- **Product or service:** Vercel supports more than 40 web frameworks and uses framework-defined infrastructure for complex front-end deployments. Its embedded AI products, like v0 and the AI SDK, enable customers to develop and deploy AI applications faster. Gartner Peer Insights survey respondents consistently praise its ease of use, reliability and performance, which contribute to improved developer productivity and reduced time to market.
- **Innovation:** Vercel pioneers Fluid compute, which lets multiple invocations share a single function instance, reducing idle compute time and lowering costs. Vercel also plans to invest in a marketplace for AI agent services and develop internal AI agents to optimize the platform and improve the developer experience.
- **Market understanding:** Vercel's framework-defined infrastructure and its deeply integrated AI capabilities show a clear understanding that the cloud-native application platform market is rapidly evolving toward AI-native workflows. Vercel is effectively positioning itself to deliver a high-performance platform for building and running AI-enabled web applications.

Cautions

- **Offering (product) strategy:** Vercel does not support containers, prioritizing front-end and AI gateway technologies. For more complex microservices architectures, primarily written in Java or .NET, Vercel customers should source their back ends separately or rely on their existing back ends, integrating them using Vercel Functions.
- **Marketing execution:** Vercel's ongoing challenge is shifting the belief that powerful infrastructure must be hand-built with low-level tools. Overcoming this skepticism among

organizations with legacy architectures, along with concerns about migration risk, workflow changes and retraining, represent significant hurdles.

- **Customer experience:** Gartner Peer Insights survey respondents indicate Vercel's role-based access control (RBAC) and observability are limited, and the platform lacks robust back-end tools. Customers may also find it difficult to migrate from legacy Java and .NET technology stacks, due to Vercel's strong focus on the JavaScript/Typescript ecosystem.

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

- Alibaba Cloud

Dropped

- **Mia-Platform:** Mia-Platform helps enterprises support developer self-service, streamlining deployment and management of cloud-native and cloud-optimized applications. The platform serves as a unified hub for technology capabilities and processes supporting DevOps. Mia-Platform's operations are mainly in Europe and its clients tend to be large to midsize enterprises across various sectors. The company is investing in improving its developer tools, deployment automation features and runtime performance optimization to allow software engineers to focus on application development. Mia-Platform was not included in this research as it did not meet the criteria for revenue and revenue growth.

Inclusion and Exclusion Criteria

To qualify for inclusion in this Magic Quadrant, each vendor needed to meet the following criteria:

Market Participation

- Meet the market definition of a cloud-native application platform.
- All features applicable to this inclusion criteria and evaluated in this Magic Quadrant and critical capabilities research must be generally available as of **1 April 2025** to all customers and fully documented. Custom development for specific customers does not qualify. “General availability” means the product or service is available on a public-facing price sheet/card for purchase directly by clients. Providers must be able to furnish the link to a pricing page for their cloud-native application platform.
- Sell the solution directly to paying customers without requiring them to engage professional services. The vendor must provide at least first-line support for these capabilities, including the use of bundled open-source and close-source software.
- Demonstrate an active product roadmap, and go-to-market and selling strategies for the solution.
- Have phone, email and/or web customer support. They must offer contract, console/portal, technical documentation and customer support in English (either as the product’s default language or as an optional localization).

Platform Capabilities

Cloud-native application platforms must provide cloud-based-managed application runtime environments for applications and integrated capabilities to manage the life cycle of an application or application component. They must also typically enable distributed application deployments and support cloud-style operations — such as elasticity, multitenancy, and self-service — without requiring infrastructure provisioning or container management. Cloud-native application platforms must be enterprise-grade and aimed at enterprise-class projects by providing high availability and disaster-recovery technical support to customers.

The selected platform must support modern full-stack frameworks using their native runtime environments, without requiring custom wrappers and compatibility layers.

Performance

Size

The vendor must, by 31 May 2025, fulfill one of the following size requirement combinations:

- Platform license and/or subscription revenue of at least \$40 million for its cloud-native application platform over the last 12 months, and at least 100 paying enterprise customer organizations (of at least 1,000 employees) for its cloud-native application platform offering, excluding other related product offerings.

Or

- Platform license and/or subscription revenue of at least \$15 million for its cloud-native application platform over the last 12 months and CAGR growth at least 50% in revenue and/or customer base for its cloud-native application platform, excluding other related product offerings.

The vendor must have direct customers (i.e., not through resellers) within three or more of the following geographies:

- North America
- South America
- Europe
- Middle East and Africa
- Asia/Pacific

In addition, the vendor must rank among the top 20 organizations in the Customer Interest indicator (CII) identified by Gartner for this Magic Quadrant. CII was calculated using a weighted mix of internal and external inputs, including:

- Gartner customer search, inquiry volume and pricing requests
- Frequency of mentions as a competitor to other vendors in the Magic Quadrant for Cloud-Native Application Platforms in reviews for similar use cases on Gartner's Peer Insights forum as of 1 April 2025
- Significant innovations in the market, as noted by major publications, product enhancements or introductions, or industry awards
- Google Trends and web traffic analysis
- Social media followers

Honorable Mentions

Broadcom: Broadcom offers VMware Tanzu Platform, which is a cloud-native application platform sold as a software product that clients can operate themselves on-premises or on the cloud infrastructure of their choice. VMware Tanzu Platform provides buildpack tooling and deployment automation to create and configure containers for deployment into an elastic application runtime based on open-source Cloud Foundry. The platform natively supports multiple languages, including Go, Java, JavaScript, .NET, PHP, Python and Ruby. Broadcom enables users to build AI agents by providing brokered public and private service bindings, including models, MCP servers, enterprise APIs and included vector database, caching, streaming and messaging services. Broadcom did not qualify for this research as it does not offer a managed applications runtime, relying instead on existing public or private IaaS services.

Mia-Platform: Mia-Platform helps enterprises support developer self-service, streamlining deployment and management of cloud-native and cloud-optimized applications. The platform serves as a unified hub for technology capabilities and processes supporting DevOps. Mia-Platform's operations are mainly in Europe and its clients tend to be large to midsize enterprises across various sectors. The company is investing in improving its developer tools, deployment automation features and runtime performance optimization to allow software engineers to focus on application development. Mia-Platform was not included in this research as it did not meet the criteria for revenue and revenue growth.

Clever Cloud: Clever Cloud offers a cloud-native application platform that natively supports and operates more than 15 programming environments, including JavaScript, PHP, .NET, Java, Python, Rust and more, as well as auto-scaling containers. Clever Cloud provides a managed services marketplace for databases, message queues, storage, identity management, AI services, alongside built-in support for continuous delivery, monitoring and cloud migration. These services support native code, containers and cloud migration use cases. Clever Cloud was not included in this research as it did not meet the criteria for revenue and revenue growth within the cloud-native application platform segment.

Fastly: Fastly's Edge Cloud Platform empowers users to build, deliver and secure serverless applications at the edge. In addition to Compute, Fastly offers a range of related cloud services, including application security data storage, edge observability, AI and CDN solutions. These services support both native code and serverless functions, offering flexibility for developers. Fastly was not included in this research as it did not meet the

required criteria for revenue and revenue growth within the cloud-native application platform edge compute segment.

Evaluation Criteria

Ability to Execute

Product/Service: This criterion assesses the core goods and services that compete in or serve the defined market, including current product and service capabilities, quality and feature sets, and whether offered natively or through OEM agreements/partnerships. It seeks to understand the vendor's cloud application runtime capabilities and support for various programming languages and frameworks. Additionally, it evaluates the scalability and availability of the cloud-native application platform, ability to provide monitoring and observability capabilities, as well as strong security and governance capabilities. Last, but not least, this criterion seeks to understand the vendor's ability to provide effective cost-management tooling.

Overall Viability: This criterion assesses the viability of the organization, including its overall financial health, the financial and practical success of the business unit, and the likelihood of continued investment in the product. It seeks comprehensive information about the vendor's financial status, including venture capital funding, profitability, strategies for economic downturn, investment plans for the next 12 months, annual revenue for fiscal year 2024, and projected revenue for FY25 and FY26. Additionally, it evaluates customer and market engagement, focusing on customer retention rates for calendar year 2024 and the first two quarters of 2025, the largest installation by number of concurrent users and any relevant business acquisitions in the last 12 months. Lastly, this criterion examines the organizational structure and workforce, including the number of full-time employees dedicated to the product, and any changes in senior management over the past year.

Sales Execution/Pricing: This criterion assesses the organization's capabilities in all presales activities and the supporting structure, including deal management, pricing and negotiation, presales support and overall effectiveness of the sales channel. It seeks detailed insights into the customer base, such as the top five decision makers in large enterprises who use the product, the current number of customers, segmentation across different sectors and the longevity of relationships with large-enterprise customers. Additionally, it evaluates the pricing models, including variations for different pricing models like pay-as-you-go, long-

term commitment and fixed pricing. Lastly, this criterion examines any free or trial offerings associated with the product, providing an understanding of the customer acquisition strategy and how potential customers can evaluate the product before making a purchase decision.

Market Responsiveness and Track Record: This criterion assesses the organization's ability to respond, change direction, be flexible and achieve competitive success as opportunities develop, competitors act, customer needs evolve and market dynamics change, including the provider's history of responsiveness to changing market demands. It seeks detailed information about the company's active engagement in open-source communities, the percentage of customers and partners contributing to the solution marketplace and the year-over-year growth of the partner marketplace. Additionally, it evaluates the mechanisms used to listen and respond to customer needs, providing specific examples of their effective employment. Lastly, this criterion examines the extent of product customization for different markets and the company's ability to innovate by being early to market with platform capabilities that competitors are only now catching up to.

Marketing Execution: This criterion assesses the clarity, quality, creativity and efficacy of programs designed to deliver the organization's message to influence the market, promote the brand, increase product awareness and establish a positive identification in the minds of customers. It seeks a clear description of how the product is positioned to development teams, the top reasons for developers to use the product, and the key differentiators that set it apart in the market. Additionally, it evaluates the estimated marketing budget for 2025, major marketing initiatives undertaken in the past year, strategies for visibility on search engines and engagement with followers/subscribers on various online and social media channels. Lastly, this criterion examines the physical, virtual/hybrid conferences sponsored or presented at in 2024 and first two quarters of 2025, identifies top competitors and highlights the unique differentiators that set the organization apart from its competitors.

Customer Experience: This criterion assesses the products, services and programs that enable customers to achieve anticipated results, including quality interactions, technical support and account support. It seeks details on training programs for developers, onboarding timelines, success measurement, implementation resources and user training needs. Additionally, it evaluates the customer support structure, dedicated full-time equivalents (FTEs), support availability, SLAs, response times, recent outages and partner involvement, including the options available for enterprises to receive service and support in the event of an emergency system down. Lastly, this criterion examines the organization's

customer success program, retention strategies, user community support, ROI measurement, and the metrics and benchmarks used to gauge the effectiveness of the customer success program.

Operations: This criterion assesses the organization’s ability to meet goals and commitments, focusing on the quality of its structure, skills, experiences, programs and systems. It seeks detailed information about SLAs, including system uptime, upgrade policies, release timing and the growth rate of FTEs devoted to enterprise technical support, as well as subscriber options for update timing. Additionally, it evaluates staff training, partner employee training, operation and support centers worldwide, onboarding speed and formal communication processes with customers. Lastly, this criterion examines the differentiation of hosting strategies and the 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	High
Market Responsiveness/Record	High
Marketing Execution	Low
Customer Experience	Medium
Operations	Medium

Source: Gartner (August 2025)

Completeness of Vision

Market Understanding: This criterion assesses the ability of vendors to understand and translate customer needs into innovative products and services, particularly within the realm of software engineering. It focuses on product development, identification of key offerings and adaptation to evolving client requirements for enhanced product viability. Additionally, it evaluates the proficiency of vendors in monitoring market trends, navigating challenges in the cloud-native application platform market and anticipating technological disruptions to formulate a forward-thinking strategic vision. Lastly, this criterion considers the commitment to corporate responsibility, including plans for carbon neutrality and active customer engagement through initiatives like a customer council program.

Marketing Strategy: This criterion assesses the organization's ability to deliver clear, differentiated messaging that is consistently communicated internally and externalized through social media, advertising, customer programs and positioning statements. It focuses on how well the organization defines its product/service messaging, positioning and go-to-market strategy, and how skillfully it plans its top marketing initiatives for industry understanding and adoption. This criterion also measures how well the organization has identified its target market verticals and client sizes. Additionally, it evaluates the organization's targeting of specific roles for product marketing, outreach to CTOs, CIOs and software engineering leaders, and differentiation of the product value proposition by persona/buyer. Lastly, it considers the identification and targeting of new partners for 2025.

Sales Strategy: This criterion assesses the organization's ability to implement a sound sales strategy using appropriate networks, including direct and indirect sales, marketing, service and communication. It focuses on the organization's understanding of its sales growth strategies for 2025 and 2026, as well as factors shaping the sales pipeline, projected growth of the sales team, planned initiatives for product adoption and the impact of market changes on sales strategy. Additionally, it evaluates potential market expansion in the next 12 months, identifying target countries/regions and industries, and the contribution of indirect sales channel partners to revenue. Lastly, this criterion considers detailed information on pricing strategy, including licensing models, discount offerings and the organization's approach to consumption-based charging.

Offering (Product) Strategy: This criterion assesses an approach to product development and delivery that emphasizes market differentiation, functionality, methodology and features aligned with current and future requirements. It focuses on understanding the cloud-native application platform offering, including its technical abilities, and trained consulting and system integrator partners. Additionally, it evaluates the user base, detailing the number of

developers using free and paid versions, and outlines the vendor's strategy for supporting and contributing to open-source software usage and development. Lastly, this criterion considers the vendor's strategic approach to product development and market positioning, including investment areas, success metrics, methods to avoid commoditization, methods to avoid lock-in, product enhancement strategies and processes for integrating customer feedback into the product roadmap.

Business Model: This criterion assesses the design, logic and execution of the organization's business proposition to achieve continued success. It delves into the vendor's business model changes concerning the cloud-native application platform offering, the model's planned evolution over the next 12 months and the offering's contribution to overall company revenue. Additionally, it evaluates the vendor's partnership strategy, focusing on the percentage of new customers obtained via partners or partner references in the last 12 months.

Vertical/Industry Strategy: This criterion assesses the organization's strategy for directing resources, skills and products to meet the specific needs of individual market segments, including verticals. It seeks a comprehensive understanding of industry-specific go-to-market or technology partnerships, providing insight into strategic alliances and potential synergistic benefits. Additionally, it requires a detailed overview of customer distribution across various verticals, including key customer names and top industry verticals. Lastly, this criterion evaluates major initiatives planned to increase market share in vertical industry segments over the next 12 months, assessing forward-thinking strategies, growth potential and commitment to innovation.

Innovation: This criterion assesses the direct, related, complementary and synergistic allocation of resources, expertise or capital for investment, consolidation, defensive or preemptive purposes. It seeks comprehensive insights into the vendor's innovation strategy, including processes and methodologies, future innovation plans, top differentiating innovations, the proportion of revenue invested in R&D, and strategic partnerships for innovation. Additionally, it requires detailed information on how the vendor differentiates itself in the market with innovative product features and strategic partnerships, providing a clear picture of its competitive edge. Lastly, this criterion evaluates the vendor's commitment to the broader technology community through contributions to open source or open standards related to its product.

Geographic Strategy: This criterion assesses the provider’s strategy to direct resources, skills and offerings to meet the specific needs of geographies outside its native region, either directly or through partners, channels and subsidiaries. It requires a comprehensive overview of the vendor’s differentiated delivery, sales and marketing strategies for various geographies, as well as its top three initiatives aimed at expanding market share beyond its core region. Additionally, it evaluates how the vendor ensures compliance with data sovereignty requirements, the internationalization/localization capabilities of its offering and the number of natural languages supported. Lastly, this criterion considers the vendor’s current and prospective geographic markets, detailing its physical presence, staff count, customers, channel partners and the number of new customers acquired in each region over the past year.

Completeness of Vision Evaluation Criteria

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

Source: Gartner (August 2025)

Quadrant Descriptions

Leaders

Leaders distinguish themselves by offering a platform suitable for strategic adoption and having a clear roadmap. They can serve a broad range of use cases, although they do not excel in all areas, may not necessarily be the best providers for a specific need and may not serve some use cases at all. Leaders in this market have appreciable market share and many referenceable customers.

Challengers

Challengers are well-positioned to serve some current market needs. They deliver a good platform that is targeted at a particular set of use cases, and they have a track record of successful delivery. However, they are not yet adapting to market challenges quickly enough and may lack a broad scope of ambition.

Visionaries

Visionaries have a clear vision of the future and are making significant investments in the development of unique technologies. Their platforms are still emerging, and they have many capabilities in development that are not yet generally available. Although they may have many customers, they might not yet serve a broad range of use cases well, or they may have a limited geographic scope.

Niche Players

The Niche Players in the market for cloud-native application platforms may be excellent providers for particular use cases or in regions in which they operate, but they should ultimately be viewed as specialist providers. They often do not serve a broad range of use cases well or have a broadly ambitious roadmap. Some may have solid leadership positions in markets adjacent to this market, but have developed only limited capabilities in this market.

Context

The cloud-native application platform market exceeded \$3.5 billion revenue in 2024, with worldwide spending growing at a double-digit, year-over-year rate of 16.4%. This market is projected to exceed the \$7 billion revenue mark by 2029, at a five-year CAGR of 15.1% from 2024 through 2029 in constant currency.

For software engineering leaders, the primary objective of using cloud-native application platforms is to streamline software development by leveraging the platforms' capabilities and automation features. These include structured execution environments for applications that effectively conceal the complexities of underlying infrastructure and computing resources.

The core capability of a cloud-native application platform is to provide a cloud-native, runtime environment for application code. The vendors in this Magic Quadrant were assessed based on this core capability. Additional capabilities, such as serverless functions, containers deployment on abstracted infrastructure, AI inference support, integration with AI application development platforms, integration with databases, event brokers, CDNs, edge infrastructure, API gateways, content management services, ERP systems and developer tools were also considered.

Software engineering leaders and their teams should recognize the rapid growth and potential of the cloud-native application platform market and consider investing in these platforms to leverage their benefits for software development and deployment. Furthermore, given the distinct separation between web front-end-focused solutions and comprehensive back-end platforms in the market, adopting a multiplatform strategy can leverage the unique strengths of each platform.

Market Overview

Cloud-native application platforms are not merely environments for running applications from native code. Rather, these cloud-based platforms integrate multiple advanced technologies that enhance the deployment, scalability, security, connectivity and observability of applications and AI agents. Cloud-native application platforms simplify the complexities of managing containers, Kubernetes deployments and virtual machines, and the integration of AI workloads to build and run AI agents and agentic applications. These capabilities drive significant market growth. By addressing challenges in infrastructure

management, these platforms enable more efficient application operations in cloud environments.

The potential to streamline software development and deployment using a single platform is appealing to many buyers. Moreover, cloud-native application platforms simplify infrastructure management, increase operational efficiency, and enable the seamless adoption of AI-driven features, which improves developer experience and productivity. Software engineering teams can focus on developing and deploying innovative software — including AI agents — accelerating their delivery cadence and ensuring applications scale intelligently to meet customer demand.

Cloud-native application platforms support developers in use cases such as:

- **High-volume transactional applications** — Developers use cloud-native application platforms to create and run applications that require high performance, scalability and resilience.
- **API-first shared services** — Developers use cloud-native application platforms to create API-first services, enabling support for microservices architectures.
- **Decoupled web UI/UX** — Developers use cloud-native application platforms to create modern, interactive and responsive user experiences, progressive web apps or embedded mobile components.
- **Cloud migration** — Developers use cloud-native application platforms to migrate legacy stack applications into the cloud without the need to immediately rebuild.
- **AI agents and applications** — Developers use cloud-native application platforms to build AI agents and applications, the autonomous or semiautonomous software entities that use AI techniques to perceive, make decisions, take actions and achieve goals in their digital or physical environments.

Some customers may seek to avoid vendor lock-in by adopting a multiprovider strategy for cloud-native application platforms, but this approach is impractical for most organizations — unless the division is made between front-end and back-end needs.

As a software engineering leader, use our evaluation to understand the performance and strategic vision of cloud-native application platform vendors. Then, identify vendors that closely align with your organization's short-term needs and longer-term strategy.

Use the companion Critical Capabilities for Cloud-Native Application Platforms to determine which products offer the specific capabilities that your organization needs.

Acronym Key and Glossary Terms

CGR	Compound annual growth rate
CDN	Content delivery network
GPU	Graphic processing unit
LXC	Linux containers
MCP	Model Context Protocol
SDK	Software development kit
WAF	Web application firewall

⊕ 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.

[About](#) [Careers](#) [Newsroom](#) [Policies](#) [Site Index](#) [IT Glossary](#) [Gartner Blog Network](#) [Contact](#) [Send Feedback](#)



© 2025 Gartner, Inc. and/or its Affiliates. All Rights Reserved.

[POLICIES](#) [PRIVACY POLICY](#) [TERMS OF USE](#) [OMBUDS](#)

[CONTACT US](#)

© 2025 Gartner, Inc. and/or its affiliates. All rights reserved.

Get The App

