Magic Quadrant for Cloud Infrastructure as a Service, Worldwide

Published 16 July 2019 - ID G00365830 - 43 min read

By Analysts Raj Bala, Bob Gill, Dennis Smith, David Wright

The distinctions among providers are apparent in the market for cloud IaaS in terms of worldwide enterprise adoption, capabilities and service availability. Infrastructure and operations leaders should evaluate providers with broad capabilities and a positive track record for customer success.

Market Definition/Description

This document was revised on 19 July 2019. The document you are viewing is the corrected version. For more information, see the Corrections (http://www.gartner.com/technology/about/policies/current_corrections.jsp) page on gartner.com.

Cloud computing is a style of computing in which scalable and elastic IT-enabled capabilities are delivered as a service using internet technologies. Cloud infrastructure as a service (IaaS) is a type of cloud computing service; it parallels the infrastructure and data center initiatives of IT. In the context of this Magic Quadrant, cloud IaaS is defined as a standardized, highly automated offering, where compute resources, complemented by storage and networking capabilities, are owned by a service provider and offered to the customer on demand. The resources are scalable and elastic in near real time, and metered by use. Self-service interfaces are exposed directly to the customer, including a web-based UI and an API. The resources may be single-tenant or multitenant, and hosted by the service provider or on-premises in the customer's data center. Thus, this Magic Quadrant covers both public and private cloud IaaS offerings.

Enterprise adoption of cloud IaaS often begins with modern workloads, emphasizing developer productivity and business agility. But an increasing amount of cloud IaaS is being bought for traditional IT, with an emphasis on cost reduction, safety and security. Infrastructure and operations (I&O) leaders typically lead the sourcing when cloud IaaS is to be used for traditional IT. By contrast, sourcing for modern workloads is typically driven by enterprise architects, application development leaders and digital business leaders. This Magic Quadrant considers both sourcing patterns and their associated customer behaviors and requirements. Many enterprises require third-party managed service providers in order to effectively utilize cloud IaaS. In those cases, enterprises should use this Magic Quadrant in conjunction with the "Magic Quadrant for Public Cloud Infrastructure Managed Service Providers, Worldwide," which covers third-party managed service providers.

Understanding the Vendor Profiles, Strengths and Cautions

Cloud IaaS providers that target enterprise and midmarket customers generally offer a high-quality service, with excellent availability, good performance, high security and good customer support. Exceptions will be noted in this Magic Quadrant's evaluations of individual providers. Note that, when we say "all providers," we specifically mean "all the evaluated providers included in this Magic Quadrant," not all cloud IaaS providers in general. Keep the following in mind when reading the vendor profiles:

- All the providers have a public cloud IaaS offering. A few also have an industrialized private cloud offering or have made announcements detailing such plans where every customer is on standardized infrastructure and cloud management tools. Although, this may or may not resemble the provider's public cloud service in either architecture or quality. A single architecture and feature set and cross-cloud management, for both public and private cloud IaaS, make it easier for customers to combine and migrate across service models as their needs dictate. They also enable the provider to use its engineering investments more effectively. Most of the providers also offer custom private clouds.

- All of the providers target midmarket businesses and enterprises, as well as other companies that use technology at scale. Some of the providers may also target small businesses and startups. Just because a provider targets a segment, however, does not necessarily mean that it is well suited to that segment's needs. Furthermore, not all providers have the capacity to serve very-large-scale customers, and some have capacity constraints in particular regions.

- All the providers offer basic cloud IaaS — compute, storage and networking resources as a service. They also offer additional value-added capabilities as well, notably cloud software infrastructure services — typically middleware and databases as a service — up to and including platform as a service (PaaS) capabilities. These services, along with IT operations management (ITOM) capabilities as a service (especially DevOps-related services), are a vital differentiator in the market, especially for Mode 2 agile IT buyers. While all providers have both IaaS and PaaS offerings, only some offer integrated IaaS+PaaS; these providers are explicitly noted.

- We consider an offering to be public cloud IaaS if the storage and network elements are shared; the compute can be multitenant, single-tenant or both. Private cloud IaaS uses single-tenant compute and storage, but unless the solution is on the customer's premises, the network is usually still shared.

- All the providers claim to have high security standards. The extent of the security controls provided to customers varies significantly, though. All the providers evaluated can offer solutions that will meet common regulatory compliance needs, unless otherwise noted. All the providers have undergone SOC 1, SOC 2 and SOC 3 audits, as well as SSAE 16, ISO/IEC 27001, ISO/IEC 27017, and ISO/IEC 27018 audits. This provides a relatively high level of assurance that the providers are adhering to generally accepted practices for the security of

their systems, but does not address the extent of controls offered to customers. Security is a shared responsibility; customers need to correctly configure controls and may need to supply additional controls beyond what their provider offers. Furthermore, providers vary in their degree of transparency, although customers typically have access to third-party assessment reports under a nondisclosure agreement.

- In general, monthly compute availability SLAs of 99.95% and higher are the norm, and they are typically higher than availability SLAs for managed hosting. Service credits for outages in a given month are typically capped at 100% of the monthly bill, but some providers have caps as low as 25%. This availability percentage is typically non-negotiable, as it is based on an engineering estimate of the underlying infrastructure reliability. Maintenance windows are normally excluded from the SLA.

- Some providers have a compute availability SLA that requires the customer to use compute capabilities in at least two fault domains (sometimes known as “availability zones” or the like); an SLA violation requires both fault domains to fail. Providers with an SLA of this type are explicitly noted as having a multi-fault-domain SLA.

- Very few of the providers have an SLA for compute or storage performance. None of these providers oversubscribe compute or RAM resources in standard compute instances, but some may have special, less expensive instance types that do, such as “burstable” instances.

- Many providers have additional SLAs covering network availability and performance, customer service responsiveness and other service aspects.

- Infrastructure resources are not normally automatically replicated into multiple data centers, unless otherwise noted; customers are responsible for their own business continuity. Some providers offer optional disaster recovery solutions.

- All providers offer, at minimum, per-hour metering of virtual machines (VMs), and some can offer shorter metering increments, which can be more cost-effective for short-term batch jobs. Providers charge on a per-VM basis, unless otherwise noted.

- Increasingly, providers are able to offer bare-metal physical servers on a dynamic basis, priced by the hour. Providers with a bare-metal option are noted as such.

- All the providers partner with carrier-neutral colocation exchanges. This allows customers to obtain connectivity from a variety of carriers that are located in these facilities. In addition, many customers have needs that require a small amount of supplemental colocation in low-latency proximity with their cloud provider. For instance, they may have a large-scale database, specialized network equipment, or legacy equipment such as a mainframe.

- Some providers offer a software marketplace where software vendors specially license and package their software to run on that provider’s cloud IaaS offering. Marketplace software can be automatically installed with a click, and can be billed through the provider. Some marketplaces also contain other third-party solutions and services.

- All providers offer enterprise-class support with 24/7 customer service, via phone, email and chat, along with an account manager. Most providers include this with their offering. Some offer a lower level of support by default, but allow customers to pay extra for enterprise-class support.
- All the providers will sign contracts with customers, can invoice and can consolidate bills from multiple accounts. All providers also offer online sign-up and credit card billing, as they recognize that enterprise buyers prefer contracts and invoices. Some will sign “zero dollar” contracts that do not commit a customer to a certain volume.
- Unless otherwise noted, all providers will sign the following contract addenda:
  - A U.S. Health Insurance Portability and Accountability Act Business Associate Agreement (HIPAA BAA)
  - An EU Data Protection Directive (95/46/EC) data processing agreement (DPA), which includes the model clauses
  - An EU General Data Protection Regulation (GDPR) DPA
- Some of the providers offer optional managed services on IaaS. Such providers may not offer the same type of managed services on IaaS as they do in their broader managed hosting or data center outsourcing services. Some may have managed service provider (MSP) or system integrator (SI) partners that provide managed and professional services.
- All the evaluated providers offer a portal, documentation, technical support, customer support and contracts in English. Some can provide one or more of these in languages other than English. Most providers can conduct business in local languages, even if all aspects of service are English-only.

The service provider descriptions are accurate as of the time of publication. Our technical evaluation of service features took place between January 2019 and March 2019.

Format of the Vendor Descriptions

When describing each provider, we first summarize the nature of the company and then provide information about its industrialized cloud IaaS offerings in the following format:

Offerings: A list of the industrialized cloud IaaS offerings (both public and private) that are directly offered by the provider. Also included is commentary on the ways in which these offerings deviate from the standard capabilities detailed in the Understanding the Vendor Profiles, Strengths and Cautions section above. We also list related capabilities of interest, such as object storage, content delivery network (CDN), application PaaS (aPaaS) and managed services, but this is not a comprehensive listing of the provider’s offerings.

Locations: Cloud IaaS data center locations by country, languages that the company does business in and languages that technical support can be conducted in.

Adoption profile: The profiles of customers that typically adopt a provider's platform.

Recommended uses: These are the circumstances under which we recommend the provider. These are not the only circumstances in which it may be a useful provider, but these are the scenarios for which, in Gartner's opinion, the provider is best suited. For all the vendors, the recommended uses are specific to self-managed cloud IaaS. However, some of the providers also have managed services, as well as other cloud and noncloud services that may be used in conjunction with cloud IaaS. These include hybrid hosting (customers sometimes blend solutions, such as an entirely self-managed front-end web tier on public cloud IaaS, with managed hosting for the application servers and database), as well as hybrid IaaS/PaaS solutions. Even though we do not evaluate managed services, nonintegrated PaaS and the like in this Magic Quadrant, they are part of a vendor's overall value proposition and we mention them in the context of providing more comprehensive solution recommendations.

For a detailed technical description of public cloud IaaS offerings, along with a use-case-focused technical evaluation, see "Critical Capabilities for Public Cloud Infrastructure as a Service, Worldwide."

We also provide a detailed list of evaluation criteria in "Evaluation Criteria for Cloud Infrastructure as a Service."

Magic Quadrant

Figure 1. Magic Quadrant for Cloud Infrastructure as a Service, Worldwide

Source: Gartner (July 2019)

Vendor Strengths and Cautions

Alibaba Cloud

Alibaba Cloud, also known as Aliyun in Chinese, is a cloud-focused service provider with headquarters in China and is a subsidiary of Alibaba Group. Founded in 2009 to provide platform services to Alibaba Group's e-commerce businesses, Alibaba Cloud now sells to companies around the world. This Magic Quadrant evaluation is focused upon Alibaba Cloud's international business, which is headquartered in Singapore, and our technical assessment was performed using the international service.

Offerings: Alibaba Cloud is integrated IaaS+PaaS. It offers Xen and KVM-virtualized multitenant compute (Elastic Compute Service) with compute-independent block storage (cloud disks). It also offers object storage (Alibaba Object Storage Service), a CDN service, a Docker-based container service (Cloud Container Service), preconfigured private cloud infrastructure (Apsara Stack and ET Brain) and a variety of PaaS-layer services including a family of database services (ApsaraDB).

Locations: Alibaba Cloud operates multiple regions in China and additionally has a presence in the U.S. (East and West Coasts), Germany, Australia, Hong Kong, Indonesia, Japan, India, Malaysia, Singapore, United Arab Emirates and the United Kingdom. It has local sales in the U.S., China, Germany, Australia, Japan, Hong Kong and Singapore. The China service portal, documentation and support are in Mandarin. The international portal, documentation and support are in English, Mandarin and Japanese.

Adoption profile: Alibaba Cloud is the current market share leader for cloud IaaS in China (42%) and performs particularly well with Chinese digital businesses and Chinese public-sector entities. In China, Alibaba offers private cloud infrastructure options for Chinese companies that want a hybrid cloud model. Outside of China, despite announced partnerships in 2018 with SAP and VMware, Alibaba Cloud is most typically used by Mode 2 buyers in the Asia/Pacific region seeking a platform for agile workloads.

Recommended uses: Cloud-first digital business workloads for customers that are (1) based in China or Southeast Asia and wish to leverage Alibaba Cloud's relationship with its parent company, or (2) need to locate cloud infrastructure in China.

Strengths

- Alibaba Cloud has an extensive set of public cloud integrated IaaS+PaaS offerings, comparable in scope to the service portfolios of other globally focused hyperscale providers.
- Alibaba Cloud is the market share leader for cloud IaaS in China, where it has positive customer satisfaction ratings along with strong ties to the Chinese public sector. Chinese businesses view Alibaba Cloud as an enabler of digital transformation and a conduit to begin earning digital revenue.
- Alibaba Cloud has close affiliation with the Alibaba Group of companies, enabling Alibaba Cloud to act as bridge into China for international companies, and a bridge out of China for Chinese companies.

Cautions

- Alibaba Cloud's international offering does not have the full capabilities of the China offering, nor the feature depth of its major global competitors (Amazon Web Services, Microsoft, Google). In many regions, specific services may only be available when using certain compute instance types.

- Alibaba Cloud earns 90% of its revenue in China and has not appreciably grown its enterprise customer base outside of China. Alibaba has limited capabilities in terms of an MSP ecosystem, third-party enterprise software integration and operational tools; and small global field teams limit adoption outside of China.
- Alibaba Cloud's financial losses are increasing and may prevent the company from continuing to invest in necessary expansions to serve international markets as the leading hyperscale provider.

Amazon Web Services

Amazon Web Services (AWS), a subsidiary of Amazon, is a cloud-focused service provider. It pioneered the cloud IaaS market in 2006.

Offerings: AWS is integrated IaaS+PaaS. Its Elastic Compute Cloud (EC2) offers metered-by-the-second multitenant and single-tenant VMs, as well as bare-metal servers. AWS's hypervisors are based on Xen and KVM. There is multitenant block and file storage, along with extensive additional IaaS and PaaS capabilities. These include object storage with an integrated CDN (Amazon Simple Storage Service [S3] and CloudFront), Docker container services (Amazon EC2 Container Service [ECS], ECS for Kubernetes [EKS], and Fargate) and event-driven "serverless computing" (AWS Lambda). VMware offers a VMware Cloud Foundation service within AWS data centers (VMware Cloud on AWS). Enterprise-grade support is extra. It has a multi-fault-domain SLA. Colocation needs are met via partner exchanges (AWS Direct Connect).

Locations: AWS groups its data centers into Regions, each of which contains at least two availability zones (data centers). It has multiple Regions across the U.S., as well as in Canada, France, Germany, Ireland, the U.K., Australia, India, Japan, Singapore, South Korea, Sweden and Brazil. It also has one Region dedicated to the U.S. federal government. There are two China Regions — Beijing (operated by Sinnet) and Ningxia (operated by Ningxia Western Cloud Data Technology [NWCD]) — which require a China-specific AWS account. It has a global sales presence. The portal and documentation are provided in English, Dutch, French, German, Italian, Japanese, Korean, Mandarin, Portuguese and Spanish. The primary languages for support are English, Japanese and Mandarin, but AWS will contractually commit to providing support in a large number of other languages.

Adoption profile: AWS strongly appeals to buyers seeking agile operations, but is also frequently chosen for traditional styles of IT operations. AWS is the provider most commonly chosen for strategic, organizationwide adoption. Transformation efforts are best undertaken in conjunction with an SI.

Recommended uses: All use cases that run well in a virtualized environment. Applications that are potentially challenging to virtualize or run in a multitenant environment — including highly secure applications, strictly compliant or complex enterprise applications (such as SAP business applications) — require special attention to architecture.

Strengths

- Enterprises make larger annual financial commitments and deploy more mission-critical workloads on AWS than with any other hyperscale provider. This speaks to how enterprises perceive AWS as a strategic provider of cloud infrastructure and platform services relative to other providers in the market.

- AWS has a broader range of customer profiles, ranging from startups and small and midsize businesses (SMBs) to large enterprises, than any other provider in this market. Enterprises using AWS benefit from the early adopters, which help to push new technologies into the mainstream, derisking such services and making them easier to consume and manage as a result.

- AWS is the most mature, enterprise-ready provider, with the strongest track record of customer success and the most useful partner ecosystem. Thus, it is the provider chosen by not only customers that value innovation and that are implementing digital business projects, but also preferred by customers that are migrating traditional data centers to cloud IaaS.

Cautions

- AWS makes frequent proclamations about the number of price reductions it has made. Customers interpret these proclamations as being applicable to the company's services broadly, but this is not the case. For instance, the default and most frequently provisioned storage for AWS's compute service has not experienced a price reduction since 2014, despite falling prices in the market for the raw components.

- AWS prioritizes being first to market with respect to delivering new services and capabilities. As a result, it is willing to launch feature-poor services or services without deep cross-platform integration, which it often defers to the future to address. The quest to be first to market sometimes results in services that need years of substantial engineering updates.

- As the ambitions of Amazon's CEO expand into additional markets, the boards of directors for companies in potentially threatened verticals have directed their IT organizations to avoid the use of AWS where possible. This may ultimately limit AWS's success in some verticals, and may impact the associated ecosystem. IT leaders in these verticals should consider a contingency plan for board-level directives.

Google

Google is an internet-centric provider of technology and services. Google has had an aPaaS offering since 2008, but did not enter the cloud IaaS market until Google Compute Engine was launched in June 2012 (with general availability in December 2013).

Offerings: Google Cloud Platform (GCP) combines an IaaS offering (Compute Engine), an aPaaS offering (App Engine) and a range of complementary IaaS and PaaS capabilities, including object storage, a Docker container service (Google Kubernetes Engine [GKE]) and event-driven "serverless computing" (Google Cloud Functions). Google also offers GKE On-Prem (software for on-premises deployment), a container-based offering for enterprise deployments. It has a

multi-fault-domain SLA. Colocation needs are met via partner exchanges (Google Cloud Interconnect).

Locations: Google has multiple regions across the U.S., as well as a presence in Belgium, Japan, Singapore, Germany, the Netherlands, the U.K., India, Australia, Brazil, Canada, Hong Kong, Switzerland and Taiwan. Google has a global sales presence. Support is available in English and Japanese. The portal is available in English, Dutch, French, German, Italian, Polish, Spanish, Turkish, Russian, Portuguese, Korean, Japanese, Mandarin, Cantonese and Thai. Documentation is available in English, German, Japanese and Brazilian Portuguese. Google operates regions with availability zones, but these zones can be separate buildings or separate power, cooling, networking, and control planes.

Adoption profile: GCP initially appealed to Mode 2 buyers with demonstrated strengths associated with big data and other analytics applications, machine learning projects, cloud-native applications, or other applications optimized for cloud-native operations. GCP is also beginning to attract enterprises with traditional Mode 1 workloads such as SAP.

Recommended uses: Big data and other analytics applications, machine learning projects, cloud-native applications, or other applications optimized for cloud-native operations.

Strengths

- Google has leveraged its internal innovative technology capabilities (e.g., automation, containers, networking) by providing a scalable IaaS offering with PaaS capabilities, centered on open-source ecosystems. While catering initially to cloud-native startups, Google is in the process of expanding its reach to enterprise customers.
- Google has differentiated technologies on the forward edge of IT, specifically in analytics and machine learning. This has driven some enterprises to select Google as a strategic cloud provider where they have deployed applications that are anchored by BigQuery.
- Google has innovated programs to assist customers with the process of operations transformation via its Customer Reliability Engineering program. The program uses a shared-operations approach to teach customers to run operations the way that Google's site reliability engineers do. This has the potential to tether Google more closely to enterprise customers.

Cautions

- Google demonstrates an immaturity of process and procedures when dealing with enterprise accounts, which can make the company difficult to transact with at times. This can be attributed to its nascent focus on the enterprise market. The immaturity of process is most pronounced in areas such as contract negotiation, discounting, independent software vendor (ISV) licensing, integration with enterprise systems and support. Google is aggressively targeting these shortcomings.

- Google has a much smaller pool of experienced MSP and infrastructure-centric professional services partners than other vendors in this Magic Quadrant. Its own professional services are still gaining traction in driving customer implementations. Some prospective customers find that these ecosystem limitations heighten migration risk.

- Google's overall enterprise coverage from a field sales and solutions perspective is behind its competitors. Further, enterprises often lament about Google's inability to craft appropriate solutions for enterprise requirements when engaging with solution architects.

IBM

IBM is a large, diversified technology company with a range of cloud-related products and services. The IBM Cloud offering has been built on IBM's July 2013 acquisition of SoftLayer and its previous Bluemix offering.

Offerings: IBM offers both multitenant and single-tenant virtualized compute resources along with bare-metal servers. It has S3-compatible Cloud Object Storage. CDN integration is offered via an Akamai partnership. IBM Cloud offers an Open Container Initiative (OCI)-based container service (IBM Cloud Kubernetes Service), event-driven serverless computing (IBM Cloud Functions), a Cloud Foundry-based aPaaS, and other PaaS capabilities. IBM also offers a Kubernetes- and container-based private cloud offering (IBM Cloud Private). Managed services are optional. Colocation needs are met via partner exchanges (IBM Cloud Direct Link).

Locations: IBM has SoftLayer infrastructure in multiple data centers in the U.S., along with data centers in Canada, Mexico, Brazil, France, Germany, Italy, the U.K., the Netherlands, Norway, Australia, Hong Kong, India, Japan, Korea, Sweden and Singapore. IBM offers formerly Bluemix-branded services in the U.S., the U.K., Germany, Australia and Japan. IBM has a global sales presence. It offers support in the wide range of languages in which IBM does business. The portal and documentation are available in English, French, German, Italian, Portuguese, Spanish, Simplified Chinese, Traditional Chinese, Korean and Japanese.

Adoption profile: IBM appeals to its existing customers who have a strong preference to purchase most of their technology from IBM. These customers primarily have Mode 1 use cases. However, IBM Cloud infrastructure capabilities may also be used to complement other IBM solutions, such as Watson.

Recommended uses: IBM outsourcing deals that use bare-metal servers as the hosting platform, where the customer has a need for supplemental basic cloud IaaS. The infrastructure may also be used as a component of applications built using the IBM Cloud PaaS capabilities. It should also be considered in circumstances that require both API control over scalable infrastructure and bare-metal servers in order to meet requirements for performance, regulatory compliance or software licensing.

Disclaimer: IBM did not respond to requests for supplemental information or for a review of the draft contents of this research. Therefore, Gartner analysis is based on other credible sources, including public information.

Strengths

- IBM has a very large base of customers with critical applications that are beginning to adopt cloud services. IBM is well positioned to partner its service businesses to assist these customers through the cloud transformation journey.
- IBM has a strong brand and existing customer relationships across the globe, and can offer support in local languages, local contracts and billing in local currency. IBM's base of strategic outsourcing customers may help to drive adoption of IBM Cloud infrastructure.
- IBM has transitioned to messaging around hybrid and multicloud computing, which does not directly challenge many of the providers in this Magic Quadrant. Rather, it offers capabilities and tooling to enterprises, allowing them to choose the cloud environment that is best for their specific application requirements, positioning IBM Cloud as a niche or specialty offering during the selection process.

Cautions

- IBM's Next-Generation Infrastructure (NGI) project has produced incremental improvements to the infrastructure services formerly under the SoftLayer umbrella. However, it has not delivered on its fundamental goal — to produce a new set of cloud IaaS offerings based on the principles of hyperscale architecture. Given this, it is unlikely that IBM will become a competitive public cloud IaaS provider.
- Despite having many worldwide data centers, the IBM Cloud experience remains disjointed, as many features are available only in specific locations. This and an unexceptional user experience cause IBM to have a higher level of user dissatisfaction than other vendors in this research.
- IBM has a smaller ecosystem of partners that provide adjacent tooling (e.g., cost management, security, governance). Additionally, it has fewer partnerships (in comparison to other providers) with major software vendors (Microsoft, SAP, Oracle) that could culminate in deployments on the IBM Cloud.

Microsoft

Microsoft is a large and diversified technology vendor that is increasingly focused on delivering its software capabilities via cloud services. Microsoft entered the cloud IaaS market with the launch of Azure Virtual Machines in June 2012 (with general availability in April 2013).

Offerings: Microsoft Azure is integrated IaaS+PaaS. It offers metered-by-the-second Hyper-V-virtualized multitenant compute (Azure Virtual Machines), as well as specialized large instances (such as for SAP HANA). There is multitenant block and file storage, along with many additional IaaS and PaaS capabilities. These include object storage (Azure Blob Storage), a CDN, a Docker-based container service (Azure Container Service), a batch computing service (Azure Batch) and event-driven "serverless computing" (Azure Functions). The Azure Marketplace offers third-party

software and services. Colocation needs are met via partner exchanges (Azure ExpressRoute) such as those from Equinix and CoreSite.

Locations: Microsoft calls Azure data center locations "regions." There are multiple Azure regions in the U.S., Canada, the U.K., Germany, France, Australia, India, Norway, UAE, Switzerland, Japan and Korea, as well as regions in Ireland, the Netherlands, Hong Kong, Singapore and Brazil. There are also six regions for the U.S. federal government; two are dedicated to the Department of Defense. (The two Azure China regions are part of a separate service operated by 21Vianet Group.) Microsoft has global sales. Documentation is available in English, French, German, Italian, Spanish, Portuguese (Brazil and Portugal), Japanese, Korean and Mandarin. Support and the service portal are available in those languages, plus Czech, Dutch, Hungarian, Polish, Russian, Swedish and Turkish.

Adoption profile: Microsoft Azure appeals to both Mode 1 and Mode 2 customers, but for different reasons. Mode 1 customers tend to value the ability to use Azure to extend their infrastructure-oriented Microsoft relationship and investment in Microsoft technologies. Mode 2 customers tend to value Azure's ability to integrate with Microsoft's application development tools and technologies, or are interested in integrated specialized PaaS capabilities, such as the Azure Data Lake, Azure Machine Learning or the Azure IoT solution accelerators.

Recommended uses: All use cases that run well in a virtualized environment, particularly for Microsoft-centric organizations.

Strengths

- Enterprises that are strategically committed to Microsoft technology generally choose Azure as their primary IaaS+PaaS provider. The integrated end-to-end experience for enterprises building .NET applications using Visual Studio and related services while deploying them to Azure is unsurpassed. Microsoft is leveraging its tremendous sales reach and ability to co-sell Azure with other Microsoft products and services in order to drive adoption.
- Azure provides a well-integrated approach to edge computing and Internet of Things (IoT), with offerings that reach from its hyperscale data center out through edge solutions such as Azure Stack and Data Box Edge.
- Microsoft Azure's capabilities have become increasingly innovative and open, where 50% of the workloads are Linux-based along with numerous open-source application stacks. Microsoft has a unique vision for the future that involves bringing in technology partners through native, first-party offerings such as those from VMware, NetApp, Red Hat, Cray and Databricks.

Cautions

- Microsoft Azure's reliability issues continue to be a challenge for customers, largely as a result of Azure's growing pains. Since September 2018, Azure has had multiple service-impacting incidents, including significant outages involving Azure Active Directory. The nature

of many of these outages is such that customers had no controls in order to mitigate the downtime.

- Gartner clients often experience challenges with on-time implementations within budget and that results from Microsoft setting unreasonably high expectations for customers. Much of this stems from modestly improving capabilities of Microsoft's field sales teams to appropriately position and sell Azure within its customer base.
- Enterprises frequently lament the quality of Microsoft technical support (along with the increasing cost of support) and field solution architects. This negatively impacts customer satisfaction, and slows Azure adoption and therefore customer spending.

Oracle

Oracle is a large, diversified technology company with a range of cloud-related products and services. In November 2016, it launched Oracle Cloud Infrastructure (OCI, formerly Oracle Bare Metal Cloud Services). Oracle continues to operate a legacy service, Oracle Cloud Infrastructure Classic, but it is being rapidly deprecated in favor of OCI.

Offerings: OCI offers both paid-by-the-hour, KVM-virtualized VMs as well as bare-metal servers (including a one-click installation and configuration of Oracle Database, Real Application Clusters [RAC] and Exadata) and a Docker- and Kubernetes-based container service (Oracle Container Engine for Kubernetes). Oracle also offers object storage (OCI Object Storage, formerly Oracle Bare Metal Cloud Object Storage). Oracle previously offered Oracle Cloud at Customer, a private cloud IaaS offering, but it is no longer being sold. Colocation needs are met via partner exchanges (Oracle FastConnect).

Locations: The OCI data centers are grouped into regions, some having only one availability zone, such as those in Canada and Japan. Regions with three availability zones are located in the east and west of the U.S., and in Germany and the U.K. Oracle has global sales. OCI's portal documentation is available only in English, but the portal is available in 28 additional languages.

Recommended mode: OCI will appeal to customers with mainly Oracle workloads with Mode 1 styles of operation, especially those with performance needs that are well suited to bare-metal servers, and those that do not need more than very basic cloud IaaS capabilities.

Recommended uses: OCI is best suited for enterprises requiring cloud IaaS for Oracle applications and for applications that require an Oracle Database.

Strengths

- Oracle's cloud strategy is anchored by its applications, database and other middleware, and spans IaaS, PaaS, and SaaS. Oracle's cloud IaaS is primarily an infrastructure foundation for its other businesses. Oracle is mainly targeting customers who want to run Oracle software on cloud IaaS, particularly those who prefer to run on Exadata appliances and bare-metal servers.

- Few companies of Oracle's lineage have been able to design a true hyperscale architecture. OCI stands out for being able to attract top talent from the leading cloud service providers in order to pursue the goal of building well-designed cloud infrastructure.
- Oracle has made good year-over-year progress in terms of new customer growth and of existing customers increasing their usage of OCI. Reference customers also comment as to the positive performance experienced with the platform.

Cautions

- Oracle is unlikely to ever be viewed by the market as a general-purpose provider of integrated IaaS and PaaS offerings. This is due to the dominance of the hyperscale providers, Oracle's late start with OCI, and the polarizing nature of Oracle in the minds of developers who often are the leading influencers for public cloud IaaS.
- Many features developed for OCI will not be extensively used by Oracle's core customer base as the company is building capabilities mainly in response to RFPs, which are often designed around the capabilities of the established hyperscale providers.
- In the past year, Oracle reprioritized its development roadmap in response to customer feedback. Basic management features were implemented, at the expense of a significant delay in introducing other services considered minimally viable in the current cloud IaaS market. This sort of departure from planned roadmap items should be expected, given the nascent nature of OCI's capabilities and the need to dynamically prioritize the roadmap based on deals won.

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 reflect a change in the market and, therefore, changed evaluation criteria, or of a change of focus by that vendor.

No vendors were added or dropped.

Inclusion and Exclusion Criteria

To qualify for inclusion in this Magic Quadrant, vendors need:

- **Market participation.** They must sell public cloud IaaS as a stand-alone service, without the requirement to use any managed services (including guest OS management), or to bundle it with managed hosting, application development, application maintenance, or other forms of outsourcing. They may, optionally, also sell a private version of this offering that uses the same architecture but is single-tenant.

- **Market traction and momentum.** They must be among the top global providers for the relevant segments (public and industrialized private cloud IaaS, excluding small deployments of two or fewer VMs). They must have ISO 27001-audited (or equivalent) data centers on at least three continents. They must have at least one public cloud IaaS offering that meets the following criteria:
  - If the offering has been generally available for more than three years: A minimum of $250 million in 2018 revenue, excluding all managed and professional services; or more than 1,000 customers with at least 100 VMs.
  - If the offering has been generally available for less than three years: A minimum of $10 million in 2018 revenue, excluding all managed and professional services, as well as a growth rate of at least 50% exiting 2018.
- **Business capabilities relevant to Gartner clients.** They must offer the public cloud IaaS service globally (it must be purchasable outside their home region), be able to invoice, offer consolidated billing, and be willing to negotiate customized contracts. They must have 24/7 customer support (including phone support). There must be an option for English-language localization of the contract, service portal, documentation and support.
- **Technical capabilities relevant to Gartner clients.** They must have a public cloud IaaS service that is suitable for supporting mission-critical, large-scale production workloads, whether enterprise or cloud-native. Specific generally available service features must include:
  - Software-defined compute, storage and networking, with access to a web services API for these capabilities.
  - Cloud software infrastructure services facilitating automated management, including at minimum, monitoring, autoscaling and database services.
  - A distributed, continuously available control plane supporting a hyperscale architecture.
  - Real-time provisioning for compute instances (small Linux VM in five minutes, 1,000 Linux VMs in one hour) and a container service that can provision Docker containers in seconds.
  - An allowable VM size of at least 16 virtual CPUs (vCPUs) and 128GB of RAM.
  - A service-level agreement for compute, with a minimum of 99.9% availability.
  - The ability to securely extend the customer's data center network into the cloud environment.
  - The ability to support multiple users and API keys, with role-based access control.

Evaluation Criteria

Ability to Execute

We evaluated vendors' Ability to Execute in this market by using the following criteria:

- *Product/Service:* Service providers were evaluated on the capabilities of their cloud IaaS offering to support all use cases being evaluated. We evaluated the breadth and depth of the feature set, self-service capabilities, automated system management and suitability to run a broad range of workload types. This criterion is important to buyers that want to purchase the most capable, feature-rich service.
- *Overall Viability (Business Unit, Financial, Strategy, Organization):* Providers were evaluated on:
  - The success of their cloud IaaS business, as demonstrated by current revenue and revenue growth since the launch of their service
  - Their financial wherewithal to continue investing in the business and to execute successfully on their roadmaps
  - Commitment to their current offerings, with no plans to execute disruptive platform transitions or migrations in the next two years
  - Their organizational commitment to this business, and its importance to the company's overall strategy

  This criterion is important to buyers that prefer to purchase services from large vendors with ample financial resources, or from vendors that have a position of market leadership and are continuing to invest aggressively in the business. It is also important to buyers that are concerned about their long-term strategic investment in a particular vendor, or who want to avoid potentially disruptive service changes.

- *Sales Execution/Pricing:* Providers were evaluated on their ability to:
  - Address the range of buyers for IaaS, including the different audiences in each mode of bimodal IT
  - Adapt to "frictionless selling" with online sales, immediate trials and proofs of concept
  - Provide consultative sales and solutions engineering
  - Be highly responsive to prospective customers
  - Offer value for money

  This criterion is important to buyers that value a smooth sales experience, the right solution proposals and competitive prices.

- *Market Responsiveness and Track Record:* This market is evolving extremely quickly and the rate of technological innovation is very high. Providers were evaluated on how well they have

historically been able to respond to changing buyer needs and technology developments, rapidly iterate their service offerings, and deliver promised enhancements and services by the expected time. This criterion is important to buyers that value rapid delivery of cutting-edge capabilities.

- *Marketing Execution:* Providers were evaluated on:
  - Their mind share and brand awareness in the market
  - Their ability to convey marketing messages based on their ability to deliver real business value, not empty hype or misleading “cloudwashing” (i.e., the practice of rebranding or remarketing an existing offering under a cloud label without offering all the attributes of a cloud service)
  - The clarity and accuracy of their marketing messages, compared with their actual service offering

  This criterion is important to buyers that prefer to buy from well-known vendors.

- *Customer Experience:* Providers were evaluated on:
  - The quality and responsiveness of their account management and technical support;
  - The ease of use of their self-service functionality
  - The capabilities of their customer portal (additional functionality such as monitoring, reporting and trouble ticketing)
  - The usefulness of their documentation and customer communications
  - The quality of their SLAs
  - The ease of doing business with them
  - Overall customer satisfaction

  This criterion is important to buyers that value the aspects of the vendor relationship and capabilities beyond the IaaS platform itself.

- *Operations:* Providers were evaluated on:
  - Their ability to meet their goals and commitments, including their track record of service delivery
  - The quality of their response to outages
  - Their approach to emergency and scheduled maintenance

- Their ability to meet timelines that are communicated to customers and to the market.

This criterion is important to buyers that want a reliable, predictable service experience.

Our evaluation of a service provider's Ability to Execute remains similar to that of the 2018 version of this Magic Quadrant, with increased expectations across all criteria.

Table 1: Ability to Execute Evaluation Criteria

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

Source: Gartner (June 2019)

Completeness of Vision

We assessed vendors' Completeness of Vision in this market by using the following criteria:

- *Market Understanding:* Providers were evaluated on their understanding of the wants and needs of three different buying constituencies in this market — enterprises, midmarket businesses and digital businesses (whether technology companies or digital business units embedded in nontechnology businesses) — both currently and in the longer term as the use of IaaS matures. This criterion is important to buyers that value a provider's understanding of the market's evolution and broader business trends, which impact a provider's ability to plan a successful long-term strategy.
- *Marketing Strategy:* Providers were evaluated on their ability to articulate their position in the market and their competitive differentiation, and to communicate these messages clearly and consistently, both internally and externally. This criterion is important to buyers that believe that providers should have a clear focus and direction.

- *Sales Strategy:* Providers were evaluated on their understanding of the buying centers for the market, and the way that these different buying centers want to engage with sales, as well as their strategy for adapting their sales force, online channel and partner channels to the IaaS market. This criterion is important to buyers that value a provider's ability to grow its business over the long term.

- *Offering (Product) Strategy:* Providers were evaluated on the breadth, depth, quality and differentiation of their service roadmaps, as relevant to the use cases under evaluation, with an emphasis on self-service, management capabilities (both traditional and DevOps-oriented), and overall feature set, including cloud software infrastructure services. This criterion is important to buyers that want a provider that will lead the market in service capabilities.

- *Business Model:* Providers were evaluated on their overall value proposition and their strategy for providing solutions for the use cases under consideration, not just raw infrastructure elements. This included evaluating how IaaS fits into their broader product portfolio and product strategy. This criterion is important to buyers that view IaaS as part of an integrated set of solutions from a particular provider.

- *Vertical/Industry Strategy:* Providers were evaluated on their ability to offer targeted services for particular vertical markets, such as government, biotechnology, media and entertainment, and retail. This includes sales and marketing to such verticals, their ability to meet specialized compliance needs, and vertical-specific solutions. This criterion is not directly important to most buyers, except to the extent that a provider has a vertical-specific offering that is relevant to them or meets their specific regulatory compliance requirements.

- *Innovation:* Providers were evaluated on the level of investment in the future of their business, and the quality of those investments, whether financial or human capital. This includes aspects such as the deployment of engineering resources, investments in new technology, mergers and acquisitions, and partnerships and alliances. This criterion is important to buyers that care about leading-edge capabilities, and the strength of a provider's ecosystem.

- *Geographic Strategy:* Providers were evaluated on their ability to expand their offering beyond their home region, serving the needs of multinational businesses, as well as adapting their offerings to other geographies. In particular, this included their strategy for international sales and support, as well as their data center footprints and internationalization efforts. This criterion is important to buyers that want to use a global vendor.

Our evaluation of Completeness of Vision remains similar to that of the 2018 version of this Magic Quadrant. However, we have continued to increase our expectations for the breadth and depth of a provider's vision, encompassing both technical capabilities and business alliances that create an ecosystem of supporting partners, with a focus on the traits that make a provider suitable for strategic adoption.

Table 2: Completeness of Vision Evaluation Criteria

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

Source: Gartner (June 2019)

Quadrant Descriptions

Leaders

Leaders distinguish themselves by offering a service suitable for strategic adoption and having an ambitious 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 service that is targeted at a particular set of use cases, and they have a track record of successful delivery. However, they are not adapting to market challenges sufficiently quickly, or do not have a broad scope of ambition.

Visionaries

Visionaries have an ambitious vision of the future, and are making significant investments in the development of unique technologies. Their services are still emerging, and they have many capabilities in development that are not yet generally available. While they may have many customers, they might not yet serve a broad range of use cases well.

Niche Players

The Niche Players in the market for cloud IaaS may be excellent providers for particular use cases or in regions in which they operate, but they should ultimately be viewed as specialist

providers of cloud IaaS. 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 only developed limited capabilities in cloud IaaS.

Context

Enterprises are adopting cloud IaaS for a broad set of use cases. As enterprises become comfortable with the providers and the technology, they increase the criticality of the workloads being deployed. As a result, customers are using cloud IaaS well beyond the initial set of workloads that characterized the market, many of which were workloads with Mode 2 styles of operation.

Some enterprises are also doing entire data center migrations. “We don’t want to be in the real estate business” is a common refrain repeated among Gartner’s enterprise clients in reference to owning and operating their own data centers. Enterprises increasingly view owning data centers as a liability rather than an asset; thus, cloud IaaS becomes an attractive alternative.

The criticality of an enterprise’s usage and its annual spending are increasing. But 2018 and first six months of 2019 have been marred by cloud service provider reliability and availability issues.

Enterprises use a variety of decision criteria when choosing a cloud IaaS provider; provider availability should be at the top of the list.

Market Overview

The market for cloud IaaS is maturing, but revenue is growing unabated. Gartner projects revenue in the cloud IaaS market to increase to $81.5 billion by 2022, up from $41.4 billion in 2019. But most of the enterprise interest and revenue are currently directed toward two providers: AWS and Microsoft. The market views both AWS and Microsoft as being general-purpose providers capable of supporting a broad range of workloads. Google is making steady progress in terms of enterprise adoption, but it remains in a distant third place in terms of overall annual revenue and interest among Gartner’s enterprise clients. All other vendors in this market are forced to focus on regional dominance or niche workloads given the momentum of AWS and Microsoft, and the scale at which they operate. Examples of regional and niche-focused vendors are Alibaba and Oracle. Alibaba dominates the market for cloud IaaS in China, and Oracle is, naturally, mostly focused on Oracle workloads as it attempts to scale in the process of rebooting its cloud endeavors. Lastly, IBM remains in a precarious position due to being slow to improve its cloud IaaS offerings, which are ultimately not competitive with the market leaders.

Evaluation Criteria Definitions

Ability to Execute

Product/Service: Core goods and services offered by the vendor for the defined market. This includes current product/service capabilities, quality, feature sets, skills and so on, whether

offered natively or through OEM agreements/partnerships as defined in the market definition and detailed in the subcriteria.

Overall Viability: Viability includes an assessment of the overall organization's financial health, the financial and practical success of the business unit, and the likelihood that the individual business unit will continue investing in the product, will continue offering the product and will advance the state of the art within the organization's portfolio of products.

Sales Execution/Pricing: The vendor's capabilities in all presales activities and the structure that supports them. This includes deal management, pricing and negotiation, presales support, and the overall effectiveness of the sales channel.

Market Responsiveness/Record: Ability to respond, change direction, be flexible and achieve competitive success as opportunities develop, competitors act, customer needs evolve and market dynamics change. This criterion also considers the vendor's history of responsiveness.

Marketing Execution: The clarity, quality, creativity and efficacy of programs designed to deliver the organization's message to influence the market, promote the brand and business, increase awareness of the products, and establish a positive identification with the product/brand and organization in the minds of buyers. This "mind share" can be driven by a combination of publicity, promotional initiatives, thought leadership, word of mouth and sales activities.

Customer Experience: Relationships, products and services/programs that enable clients to be successful with the products evaluated. Specifically, this includes the ways customers receive technical support or account support. This can also include ancillary tools, customer support programs (and the quality thereof), availability of user groups, service-level agreements and so on.

Operations: The ability of the organization to meet its goals and commitments. Factors include the quality of the organizational structure, including skills, experiences, programs, systems and other vehicles that enable the organization to operate effectively and efficiently on an ongoing basis.

Completeness of Vision

Market Understanding: Ability of the vendor to understand buyers' wants and needs and to translate those into products and services. Vendors that show the highest degree of vision listen to and understand buyers' wants and needs, and can shape or enhance those with their added vision.

Marketing Strategy: A clear, differentiated set of messages consistently communicated throughout the organization and externalized through the website, advertising, customer programs and positioning statements.

Sales Strategy: The strategy for selling products that uses the appropriate network of direct and indirect sales, marketing, service, and communication affiliates that extend the scope and depth of market reach, skills, expertise, technologies, services and the customer base.

Offering (Product) Strategy: The vendor's approach to product development and delivery that emphasizes differentiation, functionality, methodology and feature sets as they map to current and future requirements.

Business Model: The soundness and logic of the vendor's underlying business proposition.

Vertical/Industry Strategy: The vendor's strategy to direct resources, skills and offerings to meet the specific needs of individual market segments, including vertical markets.

Innovation: Direct, related, complementary and synergistic layouts of resources, expertise or capital for investment, consolidation, defensive or pre-emptive purposes.

Geographic Strategy: The vendor's strategy to direct resources, skills and offerings to meet the specific needs of geographies outside the "home" or native geography, either directly or through partners, channels and subsidiaries as appropriate for that geography and market.