

(<https://www.gartner.com/home>)

LICENSED FOR
DISTRIBUTION

Magic Quadrant for Cloud Infrastructure as a Service, Worldwide

Published: 15 June 2017 **ID:** G00315215

Analyst(s): Lydia Leong, Raj Bala, Craig Lowery, Dennis Smith

Summary

The market for cloud IaaS is dominated by two leading service providers. Other service providers have responded by launching new offerings, but customers must carefully manage the risks of adopting less-mature offerings.

Market Definition/Description

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. Cloud compute IaaS constitutes the largest segment of this market (the broader IaaS market also includes cloud storage and cloud printing). Only cloud compute IaaS is evaluated in this Magic Quadrant; it does not cover service providers that exclusively offer cloud storage, platform as a service (PaaS), SaaS, cloud service brokerage (CSB) or any other type of cloud service, nor does it cover the hardware and software vendors that may be used to build cloud infrastructure. Furthermore, this Magic Quadrant is not an evaluation of the broad, generalized cloud computing strategies of the companies profiled.

In the context of this Magic Quadrant, cloud compute IaaS (hereafter referred to simply as "cloud IaaS" or "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. Further information about IaaS is available in "Technology Insight for Cloud Infrastructure as a Service."

This Magic Quadrant evaluation includes not only the cloud IaaS resources themselves, but also the automated management of those resources, management tools delivered as services and cloud software infrastructure services. The last category includes middleware and databases as a service, up to and including some PaaS capabilities. IaaS and PaaS represent a continuum, as discussed in "Technology Insight for Integrated IaaS and PaaS," and many cloud IaaS customers also use PaaS capabilities from the same provider. Fully integrated offerings are referred to in this Magic Quadrant as "integrated IaaS+PaaS."

Even though some businesses may use an application PaaS (aPaaS) in a very IaaS-like manner, we have excluded PaaS providers from this Magic Quadrant, with the exception of those PaaS providers that also have a qualifying IaaS offering. PaaS offerings do not allow customers to obtain raw virtual machines (VMs) that can be loaded with arbitrary operating systems, middleware and applications, which is a requirement for being considered as IaaS. For PaaS providers, see "Magic Quadrant for Enterprise Application Platform as a Service, Worldwide" and "Magic Quadrant for Enterprise Integration Platform as a Service, Worldwide."

We draw a distinction between cloud infrastructure *as a service*, and cloud infrastructure *as an enabling technology*; we call the latter "cloud-enabled system infrastructure" (CESI). In cloud IaaS, the capabilities of a CESI are directly exposed to the customer through self-service. However, other services, including noncloud services, may be delivered on top of a CESI; these cloud-enabled services may include forms of managed hosting, data center outsourcing and other IT outsourcing services. In this Magic Quadrant, we evaluate only cloud IaaS offerings; we do not evaluate cloud-enabled services.

Gartner's clients are mainly enterprises, midmarket businesses and technology companies of all sizes, and the evaluation focuses on typical client requirements. This Magic Quadrant covers all the common use cases for cloud IaaS, including development and testing, production environments (including those supporting mission-critical workloads) for both internal and customer-facing applications, batch computing (including high-performance computing [HPC]), and disaster recovery. It includes not only the hosting of single-application workloads, but also the replacement of traditional enterprise data centers with cloud environments that can support a highly diverse range of workloads. It includes suitability for a wide range of application design patterns, including cloud-native applications, web-era applications and legacy enterprise applications.

Customers typically exhibit a bimodal IT sourcing pattern for cloud IaaS (see "Bimodal IT: How to Be Digitally Agile Without Making a Mess" and "Best Practices for Planning a Cloud Infrastructure-as-a-Service Strategy – Bimodal IT, Not Hybrid Infrastructure"). Most cloud IaaS is bought for Mode 2 agile IT, emphasizing developer productivity and business agility, but an increasing amount of cloud IaaS is being bought for Mode 1 traditional IT, with an emphasis on cost reduction, safety and security. Infrastructure and operations (I&O) leaders typically lead the sourcing for Mode 1 cloud needs. By contrast, sourcing for Mode 2 offerings 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.

This Magic Quadrant strongly emphasizes self-service and automation in a standardized environment. It focuses on the needs of customers whose primary need is self-service cloud IaaS, although this may be supplemented by a small amount of colocation or dedicated servers. In self-service cloud IaaS, the customer retains most of the responsibility for IT operations (even if the customer subsequently chooses to outsource that responsibility via third-party managed services). Third-party managed service providers are covered in the "Magic Quadrant for Public Cloud Infrastructure Managed Service Providers."

Organizations that need significant customization or managed services for a single application, or that are seeking cloud IaaS as a supplement to a traditional hosting solution ("hybrid hosting"), should consult the Market Guide and Magic Quadrants for managed hosting instead ("Market Guide

for Managed Hybrid Cloud Hosting, North America," "Magic Quadrant for Managed Hybrid Cloud Hosting, Europe" and "Magic Quadrant for Cloud-Enabled Managed Hosting, Asia/Pacific"). Organizations that want a fully custom-built solution, or managed services with an underlying CESI, should consult the Magic Quadrants for data center outsourcing and infrastructure utility services ("Magic Quadrant for Data Center Outsourcing and Infrastructure Utility Services, North America," "Magic Quadrant for Data Center Outsourcing and Infrastructure Utility Services, Europe" and "Magic Quadrant for Data Center Outsourcing and Infrastructure Utility Services, Asia/Pacific").

This Magic Quadrant evaluates all industrialized cloud IaaS solutions, whether public cloud (multitenant or mixed-tenancy), community cloud (multitenant, but limited to a particular customer community), or private cloud (fully single-tenant, hosted by the provider or on-premises). It is not merely a Magic Quadrant for public cloud IaaS. To be considered industrialized, a service must be standardized across the customer base; it is insufficient to use a common reference architecture. Although most of the providers in this Magic Quadrant do offer custom private cloud IaaS, we have not considered these nonindustrialized offerings in our evaluations. Organizations that are looking for custom-built, custom-managed private clouds should use our Magic Quadrants for data center outsourcing and infrastructure utility services instead (see above).

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. Many also have an industrialized private cloud offering, 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, and enable the provider to use its engineering investments more effectively. Most of the providers also offer custom private clouds.

Most of the providers have offerings that can serve the needs of midmarket businesses and enterprises, as well as other companies that use technology at scale. A few of the providers primarily target individual developers, small businesses and startups, and lack the features needed by larger organizations, although that does not mean that their customer base is exclusively small businesses.

Most of the providers are oriented toward the needs of Mode 1 traditional IT, especially IT operations organizations, with an emphasis on control, governance and security; many such providers have a "rented virtualization" orientation, and are capable of running both new and legacy applications, but are unlikely to provide transformational benefits. A much smaller number

of providers are oriented toward the needs of Mode 2 agile IT; these providers typically emphasize capabilities for new applications and a DevOps orientation, but are also capable of running legacy applications and being managed in a traditional fashion.

All the providers offer basic cloud IaaS – compute, storage and networking resources as a service. A few of the providers offer additional value-added capabilities as well, notably cloud software infrastructure services – typically middleware and databases as a service – up to and including 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. Only a few providers offer integrated IaaS+PaaS.

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 SSAE 16 audits for their data centers (see Note 1). Some may have security-specific third-party assessments such as ISO 27001 or SOC 2 for their cloud IaaS offerings (see Note 2), both of which provide a relatively high level of assurance that the providers are adhering to generally accepted practices for the security of their systems, but do 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.

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 10%. 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. However, most of the providers do not oversubscribe compute or RAM resources; providers that do not guarantee resource allocations are noted explicitly.

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 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. Some providers either offer a shared-resource pool (SRP) pricing model or are flexible about how they price the service. In the SRP model, customers contract for a certain amount of capacity (in terms of CPU and RAM), but can allocate that capacity to VMs in an arbitrary way, including being able to oversubscribe that capacity voluntarily; additional capacity can usually be purchased on demand by the hour.

Some of the providers are able to offer bare-metal physical servers on a dynamic basis. Due to the longer provisioning times involved for physical equipment (two hours is common), the minimum billing increment for such servers is usually daily, rather than hourly. Providers with a bare-metal option are noted as such.

All the providers offer an option for colocation, unless otherwise noted. Many customers have needs that require a small amount of supplemental colocation in conjunction with their cloud – most frequently for a large-scale database, but sometimes for specialized network equipment, software that cannot be licensed on virtualized servers, or legacy equipment. Colocation is specifically mentioned only when a service provider actively sells colocation as a stand-alone service; a significant number of midmarket customers plan to move into colocation and then gradually migrate into that provider's IaaS offering. If a provider does not offer colocation itself, but can meet such needs via a partner exchange, this is explicitly noted.

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. While some may also offer online sign-up and credit card billing, 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.

Many of the providers have white-label or reseller programs, and some may be willing to license their software. We mention software licensing only when it is a significant portion of the provider's business; other service providers, not enterprises, are usually the licensees. We do not mention channel programs; potential partners should simply assume that all these companies are open to discussing a relationship.

Most of the providers offer optional managed services on IaaS. However, not all 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. Customers that need multilingual support will find it very challenging to source an offering.

All the providers are part of very large corporations or otherwise have a well-established business. However, many of the providers are undergoing significant re-evaluation of their cloud IaaS businesses. Existing and prospective customers should be aware that such providers may make significant changes to the strategy and direction of their cloud IaaS business, including replacing their current offering with a new platform, or exiting this business entirely in favor of partnering with a more successful provider.

In previous years, this Magic Quadrant has provided significant technical detail on the offerings. These detailed evaluations are now published in "Critical Capabilities for Public Cloud Infrastructure as a Service, Worldwide" instead.

The service provider descriptions are accurate as of the time of publication. Our technical evaluation of service features took place between January 2017 and March 2017.

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), 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.

Provider maturity: Cloud IaaS providers vary dramatically in their level of risk – the degree to which a customer can trust them to be secure, reliable, stable businesses. We provide a three-tier maturity model in "Inform Your Cloud Service Choice With Provider Maturity," and for each provider, we list its tier in that maturity model. Tier 1 providers are global megavendors. Tier 2 providers are engaged in a struggle for sustainability, with the largest hoping to break into the top tier. Tier 2 is divided into two categories: Category A is composed of established technology vendors, while Category B consists of cloud-only (or cloud-primary) vendors that have grown enough to be significant. Tier 3 providers are emerging and risky; there are none on this Magic Quadrant. We recommend that customers focus risk assessment and mitigation efforts on Tier 2 providers, which may be undesirably immature, but are more likely to be willing to offer better contractual terms, SLAs and pricing in order to ease customer concerns. Tier 2 providers require attentive vendor management and a potential exit strategy.

Recommended mode: We note whether the vendor's offerings are likely to appeal to Mode 1 safety-and-efficiency-oriented IT, Mode 2 agility-oriented IT or both. We also note whether the offerings are likely to be useful for organizations seeking IT transformation. This recommendation reflects the way that a provider goes to market, provides service and support, and designs its offerings. All such statements are specific to the provider's cloud IaaS offering, not the provider as a whole.

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 use cases it is best used for. For a more detailed explanation of the use cases, see the Recommended Uses section below.

In the list of offerings, we state the basis of each provider's virtualization technology and, if relevant, its cloud management platform (CMP). We also state what APIs it supports – the Amazon Web Services (AWS), OpenStack and vCloud APIs are the three that have broad adoption, but many providers also have their own unique APIs. Note that supporting one of the three common APIs does not provide assurance that a provider's service is compatible with a specific tool that purports to support that API; the completeness and accuracy of API implementations vary considerably. Furthermore, the use of the same underlying CMP or API compatibility does not indicate that two services are interoperable. Specifically, OpenStack-based clouds differ significantly from one another, limiting portability; the marketing hype of "no vendor lock-in" is, practically speaking, untrue.

For many customers, the underlying hypervisor will matter, particularly for those that intend to run commercial software on IaaS. Many independent software vendors (ISVs) support only VMware virtualization, and those vendors that support Xen may support only Citrix XenServer, not open-source Xen (which is often customized by IaaS providers and is likely to be different from the current open-source version). Similarly, some ISVs may support the Kernel-based Virtual Machine (KVM) hypervisor in the form of Red Hat Enterprise Virtualization, whereas many IaaS providers use open-source KVM.

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." We have used those criteria to perform in-depth assessments of several providers: see "In-Depth Assessment of Amazon Web Services," "In-Depth Assessment of Google Cloud Platform," "In-Depth Assessment of SoftLayer, an IBM Company" and "In-Depth Assessment of Microsoft Azure IaaS."

Recommended Uses

For each vendor, we provide recommendations for use. The most typical recommended uses are:

Cloud-native applications. These are applications specifically architected to run in a cloud IaaS environment, using cloud-native principles and design patterns.

E-business hosting. These are e-marketing sites, e-commerce sites, SaaS applications, and similar modern websites and web-based applications. They are usually internet-facing. They are designed to scale out and are resilient to infrastructure failure, but they might not use cloud

transaction processing principles.

General business applications. These are the kinds of general-purpose workloads typically found in the internal data centers of most traditional businesses; the application users are usually located within the business. Many such workloads are small, and they are often not designed to scale out. They are usually architected with the assumption that the underlying infrastructure is reliable, but they are not necessarily mission-critical. Examples include intranet sites, collaboration applications such as Microsoft SharePoint and many business process applications.

Enterprise applications. These are general-purpose workloads that are mission-critical, and they may be complex, performance-sensitive or contain highly sensitive data; they are typical of a modest percentage of the workloads found in the internal data centers of most traditional businesses. They are usually not designed to scale out, and the workloads may demand large VM sizes. They are architected with the assumption that the underlying infrastructure is reliable and capable of high performance.

Development environments. These workloads are related to the development and testing of applications. They are assumed not to require high availability or high performance. However, they are likely to require governance for teams of users.

Batch computing. These workloads include high-performance computing (HPC), big data analytics and other workloads that require large amounts of capacity on demand. They do not require high availability, but may require high performance.

Internet of Things (IoT) applications. IoT applications typically combine the traits of cloud-native applications with the traits of big data applications. They typically require high availability, flexible and scalable capacity, interaction with distributed and mobile client devices, and strong security; many such applications also have significant regulatory compliance requirements.

For all the vendors, the recommended uses are specific to self-managed cloud IaaS. However, many 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.

Magic Quadrant

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



Source: Gartner (June 2017)

Customers that are comparing the 2016 and 2017 Magic Quadrants may notice that the scale of the Magic Quadrant graphic has changed – overall, the Ability to Execute axis has expanded. In a year-to-year comparison, this has the effect of making it look as if the vendors have moved down in execution. For instance, AWS appears to be less far up the Ability to Execute axis than it did in the previous year, despite significant improvements in its business in the past year.

This rescaling reflects Gartner's belief that, as later-adopting customers begin to evaluate cloud IaaS, the overall customer requirements have expanded, and there are still many unmet needs in this market, resulting in greater room for improvement for all vendors.

Customers comparing the 2016 and 2017 Magic Quadrant should also keep in mind that the Magic Quadrant shows the comparative positioning of vendors within the market. Vendors may substantially improve their capabilities from year to year, yet not achieve significant movement in their position, because their position is relative to the overall market. Due to the Ability to Execute scale change, this effect is particularly pronounced this year.

Vendor Strengths and Cautions

Alibaba Cloud

Alibaba Cloud, a subsidiary of Alibaba Group, is a cloud-focused service provider with headquarters in China. It was established in 2009, and initially provided services to Alibaba Group's e-commerce businesses. 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 offers Xen and KVM-virtualized multitenant compute (Elastic Compute Service [ECS]) with compute-independent block storage (cloud disks), object storage (Object Storage Service [OSS]), a CDN service, a Docker-based container service (Alibaba Cloud Container Service) and a variety of PaaS-layer services, including a family of database services (ApsaraDB). More capabilities are offered in the China service than in the international service.

Locations: Alibaba Cloud operates multiple regions in China and additionally has a presence in the U.S., Germany, Australia, Hong Kong, Japan, Singapore and the United Arab Emirates. It has local sales in the U.S. and China. The China service portal, documentation and support are in Mandarin. The international portal, documentation and support are only in English.

Provider maturity: Tier 2B. Alibaba Cloud is a market leader in China, but is a relatively recent entrant to the global market.

Recommended mode: Alibaba Cloud appeals to Mode 2 buyers that seek infrastructure that supports agile workloads.

Recommended uses: Cloud-native applications, batch computing, and development and testing environments for customers that are based in China, or need to locate cloud infrastructure in China.

STRENGTHS

Alibaba Cloud is the current market share leader for cloud IaaS in China, and performs particularly well with Chinese digital businesses and agencies within the Chinese government. Alibaba Group has the financial wherewithal to continue investing in new regions, engineering efforts, and regional sales and marketing for Alibaba Cloud. In China, Alibaba has built an impressive ecosystem consisting of managed service providers and ISVs, and it has begun to attract a global ecosystem to its international offering.

Alibaba Cloud's current offerings, consisting of both public cloud integrated IaaS+PaaS, as well as an on-premises private cloud stack of software and services, demonstrate the vendor's potential to become an alternative to the global hyperscale cloud providers in select regions over time.

Alibaba Cloud not only has a diverse set of capabilities – which today are already comparable to the service portfolios of other hyperscale providers – but also has begun to build out a broad global data center presence.

CAUTIONS

Alibaba Cloud's international offering, with an English-language portal, was launched in mid-2016. It has a limited track record, and does not have the full capabilities or performance of the China offering. Alibaba's international offering has very little in the way of unique differentiation compared to other hyperscale providers. Additionally, Alibaba Cloud's vision seems inextricably tied to that of its global competitors; it takes liberal inspiration from competitors when developing service capabilities and branding.

Alibaba Cloud has rapidly expanded its offering to markets outside of China in the past 18 months, but the company does not have substantial mind share with buyers in those markets, as it is still building the required local talent, industry expertise and go-to-market capabilities. Prospective international customers may also perceive security and regulatory compliance concerns when using a Chinese company, even though Alibaba Cloud has undergone third-party audits. Alibaba Cloud has substantial challenges that it must overcome before it can translate its success in China to markets outside of its home territory.

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 offers Xen-virtualized multitenant and single-tenant compute (Elastic Compute Cloud [EC2]), with multitenant storage, along with extensive additional IaaS and PaaS capabilities, including object storage with an integrated CDN (Amazon Simple Storage Service [S3] and CloudFront), a Docker container service (EC2 Container Service [ECS]), a batch computing service (AWS Batch), event-driven "serverless computing" (Lambda), and an aPaaS-like developer experience (Elastic Beanstalk). It is willing to negotiate large-scale single-tenant and on-premises deals (such as the U.S. intelligence community cloud deal). The AWS Marketplace has an extensive selection of third-party software and services. 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 regions on the East and West Coasts of the U.S., and in Canada, Germany, Ireland, U.K., Australia, India, Japan, Singapore, South Korea and Brazil. It also has one region dedicated to the U.S. federal government. There is a China region operated by Sinnet, which requires 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.

Provider maturity: Tier 1. AWS has been the market share leader in cloud IaaS for over 10 years.

Recommended mode: AWS strongly appeals to Mode 2 buyers, but is also frequently chosen for Mode 1 needs. 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

AWS remains the dominant market leader, not only in IaaS, but also in integrated IaaS+PaaS, with an end-of-2016 revenue run rate of more than \$14 billion. It continues to be the thought leader and the reference point for all competitors, with an accelerating pace of innovation on top of an already rich portfolio of services, and an expanding impact across a range of IT markets. It is the provider most commonly chosen for strategic adoption; many enterprise customers now spend over \$5 million annually, and a few spend over \$100 million. While not the ideal fit for every need, it has become the "safe choice" in this market, appealing to customers that desire the broadest range of capabilities and long-term market leadership.

AWS is the most mature, enterprise-ready provider, with the deepest capabilities for governing a large number of users and resources. Thus, it is the provider not only chosen by customers that value innovation and are implementing digital business projects, but also preferred by customers that are migrating traditional data centers to cloud IaaS. It can readily support mission-critical production applications, as well as the implementation of highly secure and compliant solutions. Implementation, migration and management are significantly eased by AWS's ecosystem of more than 2,000 consulting partners that offer managed and professional services. AWS has the broadest cloud IaaS provider ecosystem of ISVs, which ensures that customers are able to obtain support and licenses for most commercial software, as well as obtain software and SaaS solutions that are preintegrated with AWS.

CAUTIONS

AWS's extensive portfolio of services requires expertise to implement. This is somewhat mitigated by AWS's excellent business-class technical support, accurate documentation, extensive training and certification, and a partner badging system that includes an Audited MSP Partner designation that helps customers choose experienced and capable MSPs. However, customers should be aware that while it's easy to get started, optimal use – especially keeping up with new service innovations and best practices, and managing costs – may challenge even highly agile, expert IT organizations, including AWS partners.

AWS has just begun to adapt to the emergence of meaningful competitors. AWS is perceived as a cost leader, and is the key reference point for pricing in this market, but it is not eager to be the lowest-cost bidder in a competitive situation. Its granular pricing structure is complex; use of third-party cost management tools is highly recommended. Its disciplined approach to contract negotiation and discounts is based almost solely on customer spending and near-term revenue opportunity. Although its baseline enterprise agreement (EA) has recently improved and offers very competitive T&Cs even prior to negotiation, only a few of its services are covered by an SLA.

CenturyLink

CenturyLink, a U.S.-based global communications service provider, acquired Savvis, a web hoster, in 2011. It acquired Tier 3, a pure-play cloud IaaS provider, in November 2013, and merged it into Savvis to create the CenturyLink Technology Solutions business unit, where its cloud efforts reside. CenturyLink acquired managed security services company netAura and hybrid cloud management company ElasticBox in March and June of 2016, respectively. In October 2016, CenturyLink announced its intention to acquire Level 3 Communications, which will expand its network, but not impact its cloud portfolio.

Offerings: CenturyLink Cloud (CLC) is VMware-virtualized; it can be either multitenant or fully single-tenant. CLC also has an option for bare-metal servers. Other capabilities including object storage and database as a service (Relational DB). CenturyLink has discontinued the legacy Savvis offerings, including Cloud Data Center 2 and Cloud Servers. The Marketplace Provider Program provides third-party software. Enterprise-grade support is extra. Managed services are optional. CenturyLink offers an aPaaS (AppFog, a Cloud Foundry derivative), but it is not yet a fully integrated solution.

Locations: CLC is available in multiple data centers across the U.S., along with Canada, the U.K., Germany, Australia and Singapore. CenturyLink has global sales, and business is conducted in local languages, but the service is offered only in English.

Provider maturity: Tier 2A. CenturyLink seems to be re-evaluating the strategic importance of its cloud and hosting offerings.

Recommended mode: CenturyLink's cloud IaaS offerings primarily appeal to Mode 1 buyers, but may meet Mode 2 requirements that are limited to basic cloud IaaS.

Recommended uses: Self-service cloud IaaS in conjunction with managed services, for all applications that run well in a virtualized environment, excluding batch computing.

STRENGTHS

The CenturyLink platform vision is rooted in the ability to deliver the breadth of CenturyLink's capabilities in an API-accessible and composable fashion. CenturyLink has built a solid platform for increasing its own agility and ability to deliver new service offerings. It is, however, increasingly emphasizing hybrid and multicloud capabilities, such as its Cloud Application Manager (which combines a SaaS-based CMP for AWS, Azure and CLC, with optional managed services) and Runner, its Ansible-based hybrid infrastructure management solution that supports both on- and off-premises deployments.

While CenturyLink has historically used cloud IaaS as a means to pursue managed service business, it nevertheless has built a competitive feature set for self-service, and successfully blends the self-service and managed service models across a hybrid solution portfolio. CenturyLink has a track record of successfully delivering enterprise-class solutions, including managed security services. The existing CenturyLink base of managed hosting, colocation and network customers provides it with cross-selling opportunities, and its forthcoming acquisition of Level 3 Communications should further extend its reach.

CAUTIONS

CenturyLink has a solidly capable and well-implemented basic offering, and has been executing successfully on its roadmap, but despite its recent acquisitions, the execution of that roadmap continues to be outpaced by many of its competitors. Although it has stated an intention to integrate its PaaS offering, little progress is evident. CenturyLink is potentially in an uncomfortable "in between" place in the market. On one side are market leaders that have broad portfolios of fully integrated IaaS+PaaS capabilities and managed service provider partners, and that are increasingly capable of attracting risk-averse customers that might have previously chosen a vendor like CenturyLink. On the other side are niche providers that specialize in specific applications and compliance requirements.

The recent signal that CenturyLink intends to inorganically expand global footprint as a communications provider, as well as its exploration of strategic alternatives for its data center assets, creates some uncertainty around the vendor's future cloud ambitions. While cloud-focused acquisitions in 2016 demonstrate CenturyLink's active commitment to delivering cloud services, these acquisitions support a hybrid and multicloud strategy, rather than one focused on CenturyLink's own IaaS and PaaS offerings. Thus, current customers must take care to ensure that CenturyLink will continue to meet their needs in the future.

Fujitsu

Fujitsu is a large, diversified technology company. It first began to offer cloud IaaS in 2010.

Offerings: In 2016, Fujitsu launched Fujitsu Cloud Service K5 IaaS. The K5 IaaS offering is OpenStack-based and KVM-virtualized, with a variety of tenancy models, including public cloud (a multitenant back end with either multitenant or single-tenant compute), hosted private cloud and outsourced private cloud. There is a K5 Cloud Foundry-based aPaaS offering. Fujitsu's earlier offering, Fujitsu Cloud IaaS Trusted Public S5, is still available. S5 is Xen-virtualized and comes in two flavors – a fully multitenant service, and a Dedicated service with single-tenant compute and a multitenant back end. Fujitsu also has legacy regional offerings that use different technology platforms, and carry the Fujitsu Cloud IaaS Private Hosted brand in conjunction with a region name or the "Global" designation.

Locations: K5 is available in data centers in Japan, the U.K., Finland, Germany and Spain. S5 is available in the U.S., Germany, the U.K., Australia, Japan and Singapore. Fujitsu has global sales, and provides support in 34 languages. Both the K5 and S5 portals and documentation are available in English and Japanese.

Provider maturity: Tier 2A. Fujitsu is in the midst of a shift to a new cloud IaaS platform (K5).

Recommended mode: Fujitsu primarily appeals to Mode 1 customers. It may also appeal to Mode 2 customers with digital business initiatives.

Recommended uses: General business applications for customers that need managed services in conjunction with cloud IaaS. Development environments for customers that only need basic cloud IaaS, or want to use IaaS in conjunction with a Cloud Foundry-based PaaS.

STRENGTHS

Fujitsu has a long history in IT services and data center outsourcing. It has a large global sales force, is the leader in IT outsourcing in Asia/Pacific and has a strong European presence. This gives it a large existing base of captive customers into which it can sell cloud services, and it has been successful at extending existing Fujitsu relationships into cloud deals. It has very responsive support and good account management.

Fujitsu has developed a portfolio of cloud IaaS, PaaS and SaaS services, and its vision has expanded to encompass a wider range of digital business capabilities. The K5 platform is intended to offer a consistent service across public cloud, hosted private cloud and outsourced private cloud models. K5 is part of the foundation for Fujitsu's MetaArc digital business platform, which includes PaaS as well as SaaS capabilities, and is intended to appeal to Mode 2 customers.

CAUTIONS

K5 is a new platform, and thus does not have a significant operational track record. Fujitsu intends to expand the service globally throughout 2017. Customers considering Fujitsu must factor the newness of their offering into their decisions. Furthermore, Fujitsu intends to transition S5 customers to the new K5 offering (or, based on the customer's requirements, to other offerings), and thus existing customers must decide whether they wish to make this transition or instead consider all competitive options.

Fujitsu's cloud IaaS capabilities lag significantly behind those of the market leaders. K5 IaaS is a basic cloud IaaS offering; it does not greatly exceed S5 in cloud IaaS capabilities or, in the portal version evaluated by Gartner, quality of user experience. (Fujitsu introduced a new portal after the evaluation period.) K5 IaaS provides a better foundation for the future than S5, but Fujitsu will continue to need to aggressively invest in acquiring and building technology in order to be competitive in this market. Although Fujitsu can sell K5 or S5 on a stand-alone basis, and it can be purchased without the need for a long-term contract, it is most often combined with managed services or a broader outsourcing relationship.

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 (Container Engine) and event-driven "serverless computing" (Google Cloud Functions, in beta). Compute Engine VMs are KVM-virtualized and metered by the minute. Enterprise-grade support is extra. It has a multi-fault-domain SLA. Colocation needs are met via partner exchanges (Google Cloud Interconnect).

Locations: Google groups its IaaS data centers into regions, each of which contains at least two zones (data centers). There are East Coast, West Coast, and central U.S. regions, as well as regions in Belgium, Japan, Singapore 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 and Thai. Documentation is available in English and Japanese.

Provider maturity: Tier 1. Gartner estimates that GCP is a distant third in cloud IaaS and integrated IaaS+PaaS market share. Google has massive investments in infrastructure for Google as a whole.

Recommended mode: GCP primarily appeals to Mode 2 buyers.

Recommended uses: Big data applications, batch computing and cloud-native applications.

STRENGTHS

Google's strategy for GCP centers on commercializing the internal innovative technology capabilities that Google has developed to run its consumer business at scale, and making them available as services that other companies can purchase. GCP is most attractive to cloud-native companies and those that want to "run like Google," but it is now trying to win customers with traditional workloads and IT processes as well. Its Customer Reliability Engineering program (currently offered directly to a limited number of customers, as well as in conjunction with Pivotal) uses a shared-operations approach to teach customers to run operations the way that Google's site reliability engineers do. Google is increasingly positioning itself as an "open" provider, and emphasizing portability; it is involved in many open-source ecosystems, including that of Kubernetes, its open-source container cluster management software.

Google's ability to sell to a broad range of customers has improved significantly over the past year – a visible impact from recent deeper investments in GCP. Google has become much more aggressive in its go-to-market strategy and is increasingly competing for digital business projects in enterprises. Google uses deep discounts and exceptionally flexible contracts to try to win projects from customers that are currently spending significant sums of money with cloud competitors. Gartner clients typically choose GCP as a secondary provider rather than a strategic provider, though GCP is increasingly chosen as a strategic alternative to AWS by customers whose businesses compete with Amazon, and that are more open-source-centric or DevOps-centric, and thus are less well-aligned to Microsoft Azure.

CAUTIONS

GCP has a solid and well-implemented core of fundamental IaaS and PaaS capabilities, but its feature set and scope of services are not as broad as that of the market leaders. GCP is working on trying to achieve feature parity, but is also pursuing innovative and differentiated platform capabilities, such as its BigQuery analytics and Cloud Spanner distributed database, and its feature velocity continues to accelerate. Google is introducing more capabilities and partnerships that are important to enterprise customers, but until recently, it was highly focused on cloud-native applications and DevOps-style operations, not on applications and IT processes from the pre-IaaS era. It still needs to invest deeply in global expansion; GCP has data centers in just five countries, though it intends to enter an additional eight countries during 2017.

Google has only recently begun to build an ecosystem around its IaaS capabilities, which significantly heightens the challenges of adopting GCP. The initial GCP ecosystem has been primarily focused on solution development and GCP's PaaS capabilities. It has few MSP and infrastructure-centric professional services partners, although it has competent technical support staffed by software engineers. Google's acquisition of Orbitera has given it a multicloud marketplace, and will help accelerate its efforts to bring ISVs to its platform. GCP has also been

making aggressive efforts to build an ecosystem of management tools. It is actively recruiting partners, and prospective partner interest is high, but it will take time for such partners to build out capabilities.

IBM

IBM is a large, diversified technology company with a range of cloud-related products and services. In July 2013, it acquired SoftLayer, an independent web hoster with a focus on small or midsize businesses (SMBs), and in January 2014, it shut down its own SmartCloud Enterprise cloud IaaS offering, after migrating its existing customers to SoftLayer. IBM began to absorb the operations of SoftLayer, an IBM company, during 2016, and that process is ongoing.

IBM has two portals for cloud IaaS. The SoftLayer portal contains the full range of services that have previously been sold under the SoftLayer brand. However, in late 2016, IBM began to offer a subset of SoftLayer services through the service catalog in the Bluemix portal. Bluemix was originally IBM's PaaS offering and has since expanded into a broader platform for IBM Cloud; in this context, the SoftLayer services are branded IBM Bluemix infrastructure. The Bluemix portal also contains some infrastructure services that are only available through the Bluemix service catalog. This Magic Quadrant evaluation considers the customer experience through both portals. We use "SoftLayer infrastructure" to refer to all SoftLayer services (whether cloud or noncloud), regardless of which portal is used to provision and manage them.

Offerings: IBM offers both multitenant and single-tenant Citrix-XenServer-virtualized compute (Virtual Servers), as well as paid-by-the-hour nonvirtualized dedicated servers (Bare Metal Servers). It has OpenStack-based object storage with an integrated CDN (via a partnership with Verizon Digital Media Services, formerly EdgeCast); the Bluemix portal additionally offers S3-compatible object storage based on Cleversafe technology. SoftLayer also has noncloud offerings, such as paid-by-the-month dedicated servers (a broader range of configurations than is available per hour) and hosted appliances, but IBM does not make a clear distinction between these offerings and its cloud IaaS capabilities. Bluemix has a Docker-based container service, event-driven "serverless computing" (OpenWhisk), a Cloud Foundry-based aPaaS, and other PaaS capabilities. Managed services are optional. Colocation needs are met via partner exchanges (IBM Direct Link).

Locations: SoftLayer infrastructure is located 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 and Singapore. 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, Cantonese, Mandarin, Korean and Japanese.

Provider maturity: Tier 2A. IBM's cloud infrastructure strategy has shifted over time. It has made multiple forays into the cloud IaaS market, and is currently building a new cloud IaaS offering.

Recommended mode: Before the IBM acquisition, SoftLayer typically sold to Mode 2 customers (specifically startups and gaming companies with a strong interest in bare-metal dedicated hosting). Since the acquisition, IBM has increasingly focused on acquiring Mode 1 customers, but SoftLayer infrastructure better meets the needs of Mode 2 customers (as long as they only require basic cloud IaaS and specifically want bare metal).

Recommended uses: E-business hosting, general business applications and batch computing, 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. Alternatively, IBM outsourcing deals that use bare-metal servers as the hosting platform, where the customer has a need for supplemental basic cloud IaaS. SoftLayer infrastructure may also be used as a component of applications built using the Bluemix PaaS capabilities.

STRENGTHS

IBM is in the midst of a "Next-Generation Infrastructure" (NGI) engineering project, but it has not announced a release date. IBM has set aggressive design goals for the performance and price point of this new cloud infrastructure, which adopts the principles of hyperscale infrastructure architecture. It uses a fully custom IBM integrated system to deliver software-defined infrastructure; it incorporates IBM hardware design innovations as well as custom software. It represents a significant step forward in IBM infrastructure capabilities as well as in IBM's ability to serve the needs of future cloud-native applications, particularly in relation to IBM's broader ambitions in cognitive computing. NGI will eventually be the basis of its future cloud IaaS offerings, as well as the platform for other IBM Cloud offerings, and will be presented as a Bluemix experience. IBM intends to gradually transition its existing infrastructure customers to NGI.

IBM has a strong brand and existing customer relationships across the globe. IBM's base of strategic outsourcing customers help drive cloud-enabled data center outsourcing business into SoftLayer data centers. Its developer ecosystem may help to drive adoption of Bluemix Infrastructure services. IBM intends to make local presence one of its competitive differentiators; since the acquisition, it has taken advantage of SoftLayer's relatively small-scale "pod" architecture to expand the service from three countries to 16. The eventual rollout of its NGI architecture is likely to help IBM evolve beyond its current status as a hosting-scale provider, making it more viable for IBM to match the cost economics of the market leaders.

CAUTIONS

The current offering is SoftLayer infrastructure, not NGI. Other than an early 2015 introduction of new storage options, SoftLayer's feature set has not improved significantly since the IBM acquisition in mid-2013; it is SMB-centric, hosting-oriented and missing many cloud IaaS capabilities required by midmarket and enterprise customers. The details of the future NGI-based cloud IaaS offerings have not been announced. IBM has, throughout its history in the cloud IaaS business, repeatedly encountered engineering challenges that have negatively impacted its time to market. It has discontinued a previous attempt at a new cloud IaaS offering, an OpenStack-based infrastructure that was offered via the Bluemix portal in a 2016 beta. Customers must thus absorb the risk of an uncertain roadmap. This uncertainty also impacts partners, and therefore the potential ecosystem.

The IBM Cloud experience is currently disjointed. Some compute capabilities, such as the IBM Bluemix Container Service and OpenWhisk, reside in Bluemix, but Bluemix is hosted in just three SoftLayer data centers, and is thus not local to most SoftLayer infrastructure. Some SoftLayer infrastructure can be provisioned through the Bluemix portal, but this is not currently an integrated IaaS+PaaS offering, because Bluemix and SoftLayer do not share a single self-service portal and catalog with a consistent CLI and API; do not provide customers with a single integrated low-

latency network context; and do not offer a unified security context that allows the customer self-service visibility and control across the entire environment. Customers can use their IBM ID to sign in to either portal, but using Bluemix capabilities in conjunction with SoftLayer infrastructure is otherwise similar to using Bluemix in conjunction with any other third-party provider's infrastructure.

Interoute

Interoute is a U.K.-based Pan-European communications service provider.

Offerings: Interoute Virtual Data Centre (VDC) is a CloudStack-based offering that can be delivered in the customer's choice of tenancy models and on VMware, Citrix Xen or KVM virtualization. A wide variety of payment models is supported. Its "managed container platform" offers a Rancher-based container management framework on top of VDC infrastructure. Interoute's CloudStore provides a marketplace for third-party software and solutions. Managed services are optional.

Locations: Interoute VDC is located in data centers on the East and West Coasts of the U.S., plus the U.K., France, Germany, Italy, the Netherlands, Spain, Sweden, Switzerland, Turkey, Hong Kong and Singapore. It has global sales. The portal is available only in English. Centralized support is available in English, Dutch, French, German, Italian and Spanish, and Interoute's local technical support can cover most languages spoken in Western and Central Europe. Documentation is available only in English.

Provider maturity: Tier 2A. Interoute's VDC business is complementary to its core networking business.

Recommended mode: Interoute is likely to appeal primarily to Mode 1 customers, but may be a good fit for the needs of Mode 2 customers that value Interoute's unique intersection of networking and cloud IaaS capabilities or that intend to use Rancher.

Recommended uses: E-business hosting, general business applications and development environments for customers that need a broad Pan-European geographic footprint, or for whom tight integration with the WAN is important.

STRENGTHS

Interoute has tightly integrated its cloud services with its global network, which has broad coverage of the major European markets, including especially strong coverage of Central and Eastern Europe. Its "network-attached cloud computing" concept allows customers to easily integrate VDC LAN topologies with Interoute's WAN services, including the ability to use the API to configure networks that span multiple sites. It also provides direct connectivity to AWS and Azure, facilitating multicloud solutions. It is pursuing capabilities that allow the abstraction of distributed compute and network infrastructure. This is useful for customers that have complex distributed applications; Interoute is particularly interested in facilitating microservices and other container-oriented architectures.

Interoute's flexible range of choices for tenancy models, hypervisors, pricing models, and support and service models provides customers with a variety of interoperable options. Interoute's VDC is a capable enterprise-class basic cloud IaaS offering, and it is one of the few such offerings available from a Europe-based, Pan-European provider; customers with European data sovereignty

requirements may find Interoute to be more attractive than a U.S.-based provider. As Europe adopts stricter data privacy requirements, Interoute's footprint and local data center presence will increasingly become a competitive advantage.

CAUTIONS

While Interoute's VDC service supports a differentiated set of network-related capabilities, and it is moving forward with support for enterprise-class container management, Interoute has not ventured into higher-level cloud services, other than communications applications. Instead, Interoute is depending on partners to build and offer higher-level services on the VDC platform. Interoute has begun to build an ecosystem of software partners in its CloudStore, but also needs solution partners. Awareness of its brand is limited outside Europe.

Interoute has entered multiple highly competitive service markets that require substantial investment for success, including cloud IaaS and unified communications. While Interoute's IT service revenue is growing significantly, it has been impacted by the industrywide decline in fixed data revenue. Its limited financial resources make it challenging for Interoute to accelerate its success in the market.

Joyent

Joyent is a U.S.-based small service provider solely focused upon cloud services and software. It was acquired by Samsung Electronics in 2016.

Joyent did not respond to requests for supplemental information. Gartner's analysis is therefore based on previous discussions with Joyent executives, public information, use of Joyent's service, and discussions with Joyent's existing and prospective customers.

Offerings: Joyent's Triton service provides a unified model for running virtual machines and containers. It offers fixed-size, paid-by-the-instance public cloud IaaS, along with private cloud services in a variety of pricing models. It uses Joyent's own SmartOS, an open-source Type 1 hypervisor based on Illumos (an OpenSolaris derivative). SmartOS can run Linux-targeted binaries without modification. Customers have a choice between OS virtualization in a SmartOS Container (with a Docker API) and KVM virtualization on a SmartOS Container for Linux and Windows guests. It also offers Manta, an object storage service with an integrated in-place batch compute service.

Locations: Joyent has data centers in the eastern and western U.S., along with a data center in the Netherlands. It has local sales in the U.S. and U.K. Support is provided in English and Spanish. The portal, documentation and support are in English only.

Provider maturity: Tier 2A. Samsung continues to invest in Joyent in order to fulfill its internal needs, but it lacks a track record in delivering cloud services to businesses.

Recommended mode: Joyent appeals to Mode 2 buyers, particularly developers building applications that use microservice architectures, or CIOs who are primarily interested in portable container-oriented hybrid cloud solutions.

Recommended uses: Cloud-native applications and other applications that use a microservice architecture or are deployed into containers.

STRENGTHS

Joyent is a unique and innovative cloud IaaS provider with deep engineering capabilities, including particularly strong expertise with OS internals. Differentiated capabilities include its unified approach to VMs and containers, and the serverless computing capabilities integrated with its Manta storage service. It is also the progenitor of the open-source Node.js project, which is a key component in many event-driven applications.

Joyent is strongly committed to open-sourcing 100% of its stack. Thus, Joyent's innovations have found their way into open-source projects that customers can utilize to deploy on-premises, private cloud instantiations of Joyent's public cloud IaaS offerings. Customers can combine Joyent's public and private cloud offerings, or their own deployment of Joyent's software (whether open-source or commercially supported by Joyent), for portable hybrid cloud environments.

CAUTIONS

Most of Joyent's customers are startups, technology companies and digital divisions of enterprises; it has limited experience with traditional enterprises. Joyent's customers typically build cloud-native applications, rather than migrating existing applications. Additionally, Joyent offers only a limited set of cloud IaaS capabilities — features such as load balancing are not services, but rely on hosted appliances or container instances. It does not currently have integrated PaaS offerings.

Samsung's primary objective in acquiring Joyent was to obtain and sustain a low-cost hybrid cloud platform, and to own end-to-end capabilities for the mobile and IoT parts of Samsung's business. Samsung's internal priorities will increasingly drive Joyent's priorities. While Samsung's needs currently align well with cloud IaaS buyer priorities in general, there exists a risk that Samsung's internal needs could in the future deviate from and compete with (and likely win against) the needs of Joyent's external customers.

Microsoft

Microsoft is a large and diversified technology vendor that is increasingly focused on delivering its software capabilities via cloud services. Its Azure business was initially strictly PaaS, but 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 offers Hyper-V-virtualized multitenant compute (Virtual Machines), with multitenant storage, along with many additional IaaS and PaaS capabilities, including object storage (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. Enterprise-grade support is extra. It has a multi-fault-domain SLA. Colocation needs are met via partner exchanges (Azure ExpressRoute).

Locations: Microsoft calls Azure data center locations "regions." There are multiple Azure regions in the U.S., Canada, the U.K., Germany, Australia, India, 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, Japanese, Korean and Mandarin. Support and the service portal are available in those languages, plus Czech, Dutch, Hungarian, Polish, Russian, Swedish and Turkish.

Provider maturity: Tier 1. Microsoft's strong commitment to cloud services has been rewarded with significant market success.

Recommended mode: 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 Suite.

Recommended uses: General business applications and development environments, especially those that use Microsoft technologies; migration of virtualized workloads for Microsoft-centric organizations; cloud-native applications (including Internet of Things applications); and batch computing.

STRENGTHS

Microsoft Azure is second in market share, not only in IaaS, but also in integrated IaaS+PaaS. It has sustained a very high growth rate over multiple years, and Gartner estimates its end-of-2016 revenue run rate as approximately \$3 billion. Azure is already a very capable and broad platform, and Microsoft continues to accelerate its new-feature velocity. Microsoft is now launching innovative Azure capabilities of its own, rather than primarily copying competitor capabilities. Microsoft is leveraging its tremendous sales reach and ability to bundle Azure with other Microsoft products and services in order to drive adoption. It is steadily growing the size of Azure customers; many are beginning to spend more than \$500,000 a year, and a few exceed \$5 million in annual spending.

Microsoft is frequently chosen as a strategic cloud provider by customers that are committed to Microsoft technologies or that like Microsoft's overall cloud strategy, which spans IaaS, PaaS, SaaS and on-premises solutions. Furthermore, many customers that are pursuing a multicloud strategy will use Azure for some of their workloads, and Microsoft's on-premises Azure Stack software may potentially attract customers seeking hybrid solutions. Microsoft's increased openness – including immediately supporting Linux in new Azure feature releases (rather than initially supporting Windows only), embracing open-source technologies and working collaboratively with a range of partners in areas of technology innovation – represents a vital and positive strategic shift.

CAUTIONS

While Microsoft Azure is an enterprise-ready platform, Gartner clients report that the service experience feels less enterprise-ready than they expected, given Microsoft's long history as an enterprise vendor. Customers cite issues with technical support, documentation, training and breadth of the ISV partner ecosystem. Microsoft is actively addressing these issues and has made significant improvements over the last year. However, the disorganized and inexperienced ecosystem of managed and professional service partners makes it challenging for customers to

obtain expertise and mitigate risks, resulting in greater reluctance to deploy production applications or conduct data center migrations. Azure Fast Start implementations by Microsoft professional services are inconsistent in quality, and do not always accurately reflect what a customer will need to deploy production applications in Azure.

While Microsoft continues to steadily improve capabilities that help Azure fulfill enterprise needs for security, availability, performance, networking flexibility and user management, not all such functionality is currently implemented with the level of completeness, ease of use or API enablement desired by enterprise customers. Multiple generations of solutions, coupled with unclear guidance on when to use each, create significant complexity in determining the right implementation. Most Azure customers use the portal for manual management, rather than taking a more automated or DevOps approach. DevOps-oriented customers may encounter frustrations with a lack of strong Azure support in some open-source and other third-party tools and software.

NTT Communications

NTT Communications (hereafter "NTT Com"), an NTT Group company, is a Japan-based global communications service provider.

Offerings: NTT Com has three cloud IaaS offerings in general availability. Cloud ⁿ is a fully multitenant, CloudStack-based, KVM-virtualized offering, with an associated object storage offering, a CDN, a MySQL-based database as a service and a Cloud Foundry-based aPaaS. NTT Com Enterprise Cloud 1.0 is a VMware-virtualized, vCloud-API-enabled offering with an SRP pricing model, and it can be either fully multitenant or single-tenant; almost all customers use managed services, but they are optional. The newest offering, launched in March 2016, is Enterprise Cloud 2.0 (simply "Enterprise Cloud" in NTT marketing material), an OpenStack-based multihypervisor and bare-metal offering, but its feature set is similar to Enterprise Cloud 1.0. NTT Com offers a Cloud Foundry-based PaaS on top of Enterprise Cloud 2.0. Managed services are optional, but most Enterprise Cloud 1.0 and 2.0 customers use them.

Locations: Cloud ⁿ is available in multiple data centers in Japan, as well as a U.S. East Coast data center (however, Cloud ⁿ is no longer being sold to new customers in the U.S.). Enterprise Cloud 1.0 is available in data centers on the East and West Coasts of the U.S., plus the U.K., France, Germany, Spain, Australia, Hong Kong, Japan, Malaysia, Singapore and Thailand. Enterprise Cloud 2.0 is available in multiple data centers in Japan, as well as the U.S. (East Coast), U.K., Germany, Singapore, and Australia. NTT Com has a global sales presence. Cloud ⁿ support is available in English and Japanese. Enterprise Cloud 1.0 and 2.0 support is available in English, French, German, Spanish, Cantonese, Hindi, Japanese, Mandarin, Malay and Thai. The portal and documentation for all offerings are available in English and Japanese.

Provider maturity: Tier 2A. NTT Com is in the midst of a shift in its cloud strategy.

Recommended mode: Cloud ⁿ will appeal primarily to Mode 2 customers. Enterprise Cloud 1.0 will appeal primarily to Mode 1 customers. Enterprise Cloud 2.0 will appeal to Mode 2 customers when the OpenStack-based KVM-virtualized option is chosen, and to Mode 1 customers when the hosting-oriented VMware or Hyper-V options are chosen.

Recommended uses for Cloud ⁿ: Development environments and cloud-native applications where Japan-based hosting is desirable.

Recommended uses for Enterprise Cloud 1.0 and 2.0: Development environments and general business applications for customers that need a Pan-Asian footprint and are hosting-focused.

STRENGTHS

NTT Com has a significant base of existing customers, especially in the Asia/Pacific region, to which it can potentially sell cloud services. Other NTT Group companies, such as NTT Data, may also bring NTT Com cloud opportunities, as could NTT Com's partner network. NTT Com also has a long track record in managed hosting and managed security services, and can deliver these solutions in conjunction with Enterprise Cloud 2.0. NTT Com is emphasizing global consultative selling, targeted at solving the digital business challenges that customers face.

NTT Com is using its global network to reduce both the total cost of its cloud solutions and friction in its customer implementations. There are no data transfer charges for the cloud IaaS offerings. NTT Com cloud customers receive a free connection between the offering and NTT Com's Arcstar Universal One network. For NTT Com Enterprise Cloud 1.0 and 2.0, NTT Com has implemented software-defined networking in its data centers and in the interfaces between the offering and the WAN.

CAUTIONS

Enterprise Cloud 2.0 is the unified cloud IaaS platform that was intended to blend Cloud ⁿ and Enterprise Cloud 1.0 capabilities, and still seems like an offering in the midst of a transition. It is a basic cloud IaaS offering with little differentiation, and it is missing many capabilities that are important to customers. Its feature set is similar to Enterprise Cloud 1.0, and it thus represents limited improvement over the previous offering. It is further hampered by poor documentation and a portal that is both slow and difficult to navigate.

NTT Com is in the midst of a strategic shift in cloud IaaS-related offerings. Enterprise Cloud 2.0 is intended to serve as the unified cloud IaaS platform for its cloud efforts. NTT has also begun to emphasize multicloud capabilities, and recently launched NTT Com Cloud Management Platform, a SaaS-based multicloud CMP. These new services have a limited operational track record. The current operational record for Enterprise Cloud 2.0 shows more outages than is typical for competitors.

Oracle

Oracle is a large, diversified technology company with a range of cloud-related products and services. In late 2015, it launched its first public cloud IaaS offering, the Oracle Compute Cloud Service ("Gen 1 Cloud"). In November 2016, it launched its next-generation offering, Oracle Bare Metal Cloud Services (BMC Service, or "Gen 2 Cloud"). In 2016, Oracle purchased Ravello (a cloud service that runs as an overlay on top of third-party clouds as well as Oracle's IaaS), and in 2017, Dyn (a managed DNS provider); neither is in scope for this Magic Quadrant, but are closely related businesses.

Offerings: The Gen 2 service 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, RAC and Exadata) and a Docker-based container service (Oracle Container Cloud Service). The Gen 1 service

offers paid-by-the-hour, Xen-virtualized VMs. Oracle Cloud Machine provides a Gen 1-compatible, on-premises private cloud IaaS offering. Oracle also offers two forms of object storage (Oracle Storage Cloud Service and Oracle BMC Object Storage Service).

Locations: The Gen 2 data centers are grouped into regions, each of which contains at least two availability domains (data centers); there is a western U.S. region and a U.S. East Coast region. The Gen 1 data centers are located in the central and eastern U.S., the U.K., and the Netherlands. Oracle has global sales. The Gen 2 service is available only in English. The Gen 1 service is supported and documented only in English, but the service portal is also available in French, German, Italian, Spanish, Russian, Portuguese, Japanese, Korean and Mandarin.

Provider maturity: Tier 2A. Oracle's cloud IaaS strategy has evolved over time. It has made several previous forays into the market, and it recently introduced a new platform (Gen 2).

Recommended mode: The Gen 2 service will appeal to both Mode 1 and Mode 2 customers, 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. The Gen 1 service will appeal to Mode 1 customers.

Recommended uses for Gen 2: Cloud-native applications or batch computing that requires bare-metal servers, Oracle Databases or other use cases that require bare-metal servers to be provisioned within minutes.

Recommended uses for Gen 1: General business applications or development environments that require basic cloud IaaS capabilities, for customers that are strategically committed to using Oracle solutions.

STRENGTHS

Oracle's Gen 2 offering, despite being branded "Bare Metal Cloud Services," is really both a bare-metal and virtualized cloud IaaS platform. Oracle intends for this platform to be the basis for its future PaaS and SaaS offerings as well. It is being built by a highly experienced engineering team recruited primarily from hyperscale cloud providers. It has well-designed hyperscale cloud architecture, and a thoughtful selection of current and future features. Oracle has a realistic perspective on its late entry into the market, and has a sensible engineering roadmap focused on building a set of core capabilities that will eventually make it attractive for targeted use cases.

Oracle's broader cloud strategy spans IaaS, PaaS and SaaS. It has a strong developer ecosystem as well as a vital anchor — the Oracle Database. Regardless of whether or not Oracle is able to compete successfully for stand-alone cloud IaaS business, its Gen 2 offering will be important for Oracle's PaaS and SaaS businesses. Customers that need to run Oracle RDBMS configurations that require bare metal, but want to do so in a cloud business model, may also be attracted to this offering, even if their infrastructure primarily resides in another cloud provider, once low-latency connectivity to other cloud providers becomes available.

CAUTIONS

The Gen 2 offering is currently a bare-bones "minimum viable product." It contains only the most vitally necessary cloud IaaS compute, storage and networking capabilities. It has a limited operational track record. Most customers are dependent upon direct support from Oracle's

engineering team. Oracle has just begun to build a partner ecosystem. Customers need to have a very high tolerance for risk, along with strong technical acumen. Nevertheless, Oracle field sales and executives are very aggressively promoting the Gen 2 offering. Customers should be cautious of high-pressure sales tactics, understand the reality behind the marketing and not feel obliged to evaluate the offering at this stage of its maturity. Gartner strongly encourages prospective customers to speak with references.

The Gen 1 offering is a basic cloud IaaS offering with little in the way of differentiation, and is primarily purchased as a base for Oracle's PaaS offerings. However, it is consistent with Oracle's Cloud Machine private cloud IaaS offering, and thus may be attractive to customers that are interested in a hybrid cloud solution – but such customers also will be aligned to the legacy Gen 1. Nevertheless, the Gen 2 offering will be Oracle's primary cloud IaaS offering going forward, and Gen 1 customers should factor this into their future planning.

Rackspace

Rackspace is an independent web hoster and managed services provider. Rackspace entered the cloud IaaS market with the 2008 acquisition of Slicehost, and its OpenStack-based offering became generally available in 2012.

Offerings: Rackspace Public Cloud is a fully multitenant, OpenStack-based, Citrix XenServer-virtualized offering; the offering also has OpenStack Ironic-based bare-metal servers (OnMetal Cloud Servers) that are provisioned in approximately five minutes, and paid for per minute. Rackspace also offers three flavors of hosted private cloud: vCloud Director-based and VMware-virtualized; Microsoft Cloud OS-based and Hyper-V virtualized; and OpenStack-based and KVM-virtualized. It also offers a Rackspace-operated OpenStack private cloud on the customer's premises. Private clouds are priced on the basis of dedicated capacity. Rackspace has object storage with an integrated CDN (Cloud Files). Customers must choose either a paid support plan or managed services.

Locations: Rackspace Public Cloud and the hosted private cloud services are offered in data centers in the central and eastern U.S., the U.K., Australia and Hong Kong. However, accounts are region-specific; Europe is a separate region from the rest of the world. Rackspace has sales presence in the countries where it has data centers, along with the Netherlands, Switzerland and Mexico. Support is provided in English only. The portal and documentation are available only in English.

Provider maturity: Tier 2B. Rackspace's cloud strategy has evolved over time, and it is primarily focused on managed services.

Recommended mode: Rackspace appeals to both Mode 1 and Mode 2 customers that value highly responsive customer service.

Recommended uses for Rackspace Public Cloud: Cloud-native applications, requiring a basic cloud IaaS offering that includes large bare-metal servers; cloud IaaS as part of a hybrid hosting solution with DevOps-oriented managed services; hybrid hosting where cloud IaaS is supplementary to a primarily dedicated infrastructure; and development environments where simplicity and ease of use are crucial.

Recommended uses for Rackspace Private Cloud: Private OpenStack environments for development or cloud-native applications; VMware or Hyper-V-based "rented virtualization" for general business applications or development environments; private "Azure-like" (Windows Azure Pack) environments for development; and hybrid environments with AWS, Microsoft Azure or GCP.

STRENGTHS

Rackspace has evolved away from promoting its OpenStack-based public cloud IaaS offering, and is now focused on delivering private cloud IaaS on a variety of technologies. It has also returned to its roots in managed services, which it now offers both on Rackspace's own infrastructure solutions, as well as on the AWS, Azure and GCP platforms. This makes it well-positioned to deliver hybrid and multicloud solutions.

Rackspace is the market leader in industrialized private cloud IaaS. Its industrialized private cloud offerings are thoughtfully constructed, more automated than most competing offerings, and operated in a fashion that allows Rackspace to deliver reliable, well-supported services at economical prices and at scale. It is technology-neutral, with solutions based on VMware, Microsoft and OpenStack platforms. It is one of the only providers to currently support Azure Stack in technical preview, as a proof of concept. It also adds value beyond providing basic infrastructure, such as offering solution templates within its OpenStack-based private cloud. It can help customers learn to operate in a DevOps-oriented fashion.

CAUTIONS

Rackspace delivers basic cloud IaaS offerings, with few differentiated capabilities. This is common in private cloud IaaS, where Rackspace has better-than-average capabilities. However, public cloud IaaS customers expect more in today's market. Rackspace Public Cloud is similar to virtual private server (VPS) solutions, a niche that is increasingly dominated by aggressive competitors such as DigitalOcean, and Rackspace no longer seems to be investing significantly in this service. Rackspace also has not invested significantly in PaaS offerings, and thus it is likely to be dependent upon Azure Stack to deliver private cloud solutions that integrate IaaS and PaaS.

Rackspace was sold to Apollo Global Management, a private-equity investor, in late 2016. Rackspace is no longer a publicly traded company, and therefore is no longer required to be transparent about its financial results. This approach makes it easier for Rackspace to execute on a strategic shift toward delivering cloud managed services on a variety of platforms, but it can make it more difficult for customers to assess vendor-related risks.

Skytap

Skytap is a U.S.-based small independent service provider solely focused on cloud IaaS. It launched its service offering in 2008.

Offerings: Skytap offers VMware-virtualized, paid-by-the-VM, public cloud IaaS, both directly and via a partnership with IBM, where Skytap is resold under the brand IBM Cloud for Skytap Solutions. It has a Docker-based container service (Skytap Container Management). It does not offer any integrated PaaS capabilities. Skytap will assist customers that need private WAN connectivity or colocation options, but only in some locations.

Locations: Skytap's service is available in the U.S. and in Singapore, and, via the IBM partnership, in SoftLayer data centers in the U.K., Australia and China. It has local sales in the U.S. and the U.K. Skytap's service is in English only.

Provider maturity: Tier 2B. Skytap has grown steadily and maintained a consistent focus over time.

Recommended mode: Skytap appeals to Mode 1 customers that need to be able to model complex enterprise on-premises environments for development and testing purposes, or that are interested in duplicating such environments in order to lift-and-shift workloads.

Recommended uses: Developer collaboration, virtual training labs and ISV demo environments.

STRENGTHS

Skytap has a differentiated offering that is focused on modeling complex enterprise environments, especially those with complicated network setups. Skytap allows customers to easily create sandboxes that accurately replicate these environments, thus facilitating development, demos and training in their cloud service, while allowing production to remain on-premises. Skytap also wants to allow customers to lift-and-shift on-premises environments to its cloud, without requiring any changes, thus allowing customers to achieve incremental agility with minimal effort. Skytap currently supports x86-based environments, and has announced support for environments that use IBM's AIX operating system on POWER8.

Skytap composes infrastructure resources and their associated configurations, along with software running on the VMs, into "environments." Skytap's ability to easily snapshot, clone and share these environments facilitates developer collaboration. Skytap has a useful API that customers have successfully integrated into their workflows, including integrating it with learning management systems. Skytap has a responsive support organization that can help customers use its service more productively.

CAUTIONS

Skytap has only recently begun to seek customers that want to run production applications in Skytap's service, and it is primarily focused on traditional applications, not cloud-native applications. Customers for whom agility is a priority or that make frequent changes should take note that a VM stage change (such as starting, stopping, provisioning or removing a VM) will cause a Skytap environment to become read-only for several seconds, preventing any other operations from being performed until the state change is finished.

Skytap is a small venture-capital-backed startup that has successfully carved out a niche in the highly competitive cloud IaaS market, but may have difficulties expanding beyond that niche. Its partnership with IBM is beneficial to its sales reach, and helps it compete against Oracle's Ravello offering (which provides similar capabilities, but with a choice of cloud IaaS platforms), but also gives IBM significant influence over its future. Skytap may be a target for acquisition.

Virtustream

Virtustream, a U.S.-based subsidiary of Dell Technologies, is focused solely on cloud services and software. Virtustream was founded in 2008. It was acquired by EMC in July 2015, and EMC's managed services and some cloud-related assets were moved into Virtustream before EMC was acquired by Dell in September 2016.

Offerings: Virtustream Enterprise Cloud is hypervisor-neutral, but typically supports VMware and KVM. It is offered in both single-tenant and multitenant variants; furthermore, it can support single-tenant compute with a multitenant back end, as well as bare metal. VMs are available by the hour, bare metal is available by the month, and both paid-by-the-VM and SRP models are available. The offering embeds a tool for governance, risk management and compliance (GRC) leveraging capabilities from Virtustream's Viewtrust software. A similar offering, Virtustream Federal Cloud, targets U.S. federal government customers. The Virtustream Storage Cloud offers S3-compatible object storage that can integrate with some EMC storage products. Managed services are optional. Virtustream also offers its CMP, xStream, as software.

Locations: Virtustream has multiple data centers in the eastern and western U.S., the U.K., France, Germany, the Netherlands, Australia and Japan. It has sales presence in the U.S., the U.K., Ireland, Germany, Lithuania, Australia, India and Japan. Virtustream's service portal is provided in English, German, Japanese, Lithuanian, Portuguese and Spanish. Documentation and support are provided in English only.

Provider maturity: Tier 2. Although Virtustream is a strategic asset for a large technology vendor, it is also a highly focused niche provider.

Recommended mode: Virtustream's focus on complex traditional enterprise applications means that it appeals primarily to Mode 1 customers, especially those seeking improved agility.

Recommended uses: Complex workloads, particularly those related to ERP or other enterprise software suites, including applications that may not have been designed to run in virtualized environments.

STRENGTHS

Virtustream occupies a unique niche in the market. Its cloud services are targeted primarily at mission-critical complex enterprise applications, such as ERP suites from SAP, Oracle and Blackbaud. It is differentiated by its application-specific expertise; a platform purpose-built for the availability, performance, security, governance and SLA requirements of such applications; and the micro-VM technology that allows it to charge for resources consumed, rather than resources allocated. Its ability to template and automate the deployment and management of these applications helps improve agility and reduces the risk of application changes.

Virtustream has weathered the challenges of being acquired twice within a short period of time and is thriving under Dell, its new host. It has continued to successfully win large-scale enterprise deals, particularly those focused on SAP and that require managed service capabilities.

Virtustream is expanding the geographies in which it serves customers, increasing investments in R&D, expanding its focus to other complex workloads and leveraging Dell's expansive reach to sell to more customers.

CAUTIONS

Virtustream provides deep and differentiated capabilities in its focus areas, rather than broad general-purpose capabilities. Customers should expect Virtustream to continue to focus on its core strengths, rather than expanding into the broad cloud IaaS market. Virtustream's roadmap is inextricably tied into other Dell entities, such as VMware, EMC and Pivotal, which each have their

own sets of differing, and possibly competing, priorities. Customers should treat Virtustream as a specialized provider for the workloads that suit the strengths and weaknesses of its technology platform.

Although Virtustream supports self-service capabilities, it primarily targets complex, mission-critical applications where it is likely that the customer will purchase professional services assistance for implementation, and managed services on an ongoing basis. Virtustream is a compelling and unique provider for particular enterprise application use cases, but it is better-suited to implementations where an environment will be carefully and consultatively tuned for the needs of particular applications, rather than general-purpose environments where workloads are deployed without oversight. Prospective customers should ensure that they have a clear understanding of roles and responsibilities, and that their expectations match what is actually written in the contract.

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

Interoute

Joyent

Oracle

Skytap

"IBM (SoftLayer)" now appears simply as "IBM."

Dropped

VMware has been dropped, because it has exited the cloud IaaS market. In April 2017, after the completion of the evaluation period for this Magic Quadrant, VMware announced that OVH, a hosting and cloud IaaS provider headquartered in France, would acquire the vCloud Air business. The acquisition closed in May 2017, and excludes vCloud Air in Australia. Customers should expect that the vCloud Air service will continue to be operated in its current data centers, but OVH has not otherwise publicly announced long-term plans for the service. Prospective customers should contact OVH, not VMware.

Because we had completed the evaluation of VMware prior to the announcement of the acquisition, and we believe that both the VMware and vCloud Air-related information remain accurate, we are including the full profile that resulted from our research, even though VMware no longer appears on the Magic Quadrant.

VMware

VMware has historically been a software vendor focused on virtualization technologies. It entered the cloud IaaS market when it launched the VMware vCloud Hybrid Service (vCHS), now renamed vCloud Air, into general availability in September 2013. VMware is a Dell Technologies company.

In April 2017, after the completion of the evaluation period for this Magic Quadrant, VMware announced that OVH, a hosting and cloud IaaS provider headquartered in France, would acquire the vCloud Air business. The acquisition excludes vCloud Air in Australia, which will continue to be owned and operated by VMware. The acquisition closed in May 2017. This Magic Quadrant is a snapshot in time; it evaluates vCloud Air when it was still part of VMware, with no consideration of its expected future as part of OVH. Customers should expect that the vCloud Air service will continue to be operated in its current data centers, but OVH has not otherwise publicly announced long-term plans for the service. Prospective customers should contact OVH, not VMware.

Offerings: vCloud Air is a VMware-virtualized, vCloud-API-enabled offering that comes in three variants: OnDemand (fully multitenant), Virtual Private Cloud (fully multitenant and SRP-priced) and Dedicated Cloud (single-tenant compute with multitenant back end, and SRP-priced with customer-controlled oversubscription). There is also a disaster recovery service. All vCloud Air services share a common portal and are delivered as resource pools out of the same shared hardware.

Locations: vCloud Air is available in multiple data centers in the U.S., as well as in the U.K., Germany, and Australia; the Japan-based service was retired in March 2017. VMware has a global sales presence. Support is available in English, French, German, Portuguese, Spanish, Hindi, Japanese and Mandarin. The portal and documentation are available only in English.

Provider maturity: Tier 2A. Note that the risk profile of vCloud Air will be different under OVH than it was under VMware.

Recommended mode: vCloud Air primarily appeals to Mode 1 customers with existing investments in VMware technology.

Recommended uses: Development environments, general business applications, supplementing existing VMware-virtualized environments, Pivotal Cloud Foundry hosting and disaster recovery for customers seeking a VMware-based solution.

STRENGTHS

VMware is the market share leader and thought leader in virtualization. It has a broad global base of existing customers that are deeply committed to its technologies. Its strategy for vCloud Air is to offer hybrid cloud options to existing VMware customers, enabling its channel partners, reinforcing its position in internal data centers and expanding its total addressable market. It wants to offer customers a consistent experience across VMware-based infrastructure, whether delivered as an on-premises virtualized environment or delivered as a cloud service.

vCloud Air is a vCloud Director-based service (although it has an easier-to-use portal as the primary UI), and takes advantage of VMware's NSX software-defined networking technology; the result is a capable basic cloud IaaS offering. The offering appeals to IT administrators that are comfortable with VMware's technology, and that are interested in supplementing their on-premises VMware-based environments with infrastructure in other geographies, on-demand capacity or disaster recovery.

CAUTIONS

vCloud Air has met with limited commercial success, and VMware has shifted its strategic focus away from this offering. Instead, VMware is focusing on a partnership to deliver VMware Cloud Foundation (VCF) using AWS data centers and infrastructure, as well as its cross-cloud solutions, which deliver capabilities on the cloud IaaS offerings of competitors. It also wants to better enable its vCloud Air Network service providers, which deliver solutions using VMware technology.

vCloud Air customers face an uncertain future, as it is unclear how much will be invested in the future enhancement of this offering.

vCloud Air has limited appeal to the business managers and application development leaders who are typically the key decision makers for cloud IaaS sourcing. VMware administrators in I&O are the most likely champions of vCloud Air within a business, but they often prefer to build internal solutions, and they are also often the people that the business is trying to bypass by going to cloud IaaS. Furthermore, many I&O leaders are interested in cloud-inspired infrastructure, which brings greater automation and efficiency to virtualized infrastructure, rather than true cloud IaaS, which emphasizes self-service and API enablement. I&O leaders may thus be more interested in VCF as a service in conjunction with vRealize Automation than they are in vCloud Air.

Inclusion and Exclusion Criteria

To be included in this 2017 Magic Quadrant, vendors had to demonstrate the following, as of January 2017:

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 15 global providers for the relevant segments (public and industrialized private cloud IaaS, excluding small deployments of one or two VMs), based on Gartner-estimated market share and mind share.

Business capabilities relevant to Gartner clients. They must offer the public cloud IaaS service globally, be able to invoice, offer consolidated billing and be willing to negotiate customized contracts. They must have 24/7 customer support (including phone support). They must offer the contract, service portal, documentation and support in English (either as the service's default language, or as an optional localization).

Technical capabilities relevant to Gartner clients. The public cloud IaaS service must be suitable for supporting production workloads, whether enterprise or cloud-native. Specific service features must include:

- Data centers in at least two metropolitan areas, separated by a minimum of 250 miles, on separate power grids, with SSAE 16, ISO 27001 or equivalent audits

- Real-time provisioning (small Linux VM in five minutes)

- The ability to scale an application beyond the capacity of a single physical server

An allowable VM size of at least eight vCPUs and 64GB of RAM

An SLA 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

Access to a web services API

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; and 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; and 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" (the practice of rebranding or remarketing an

existing offering under a cloud label without offering all the attributes of a cloud service); and 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; and 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; and 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 2016 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 2017)

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 2016 Magic Quadrant. However, we have continued to increase our expectations for the breadth and depth of a provider's vision, particularly with regard to the integration of IaaS and PaaS across a spectrum of capabilities. We believe that a comprehensive vision must encompass the ambition to run any workload, at anytime, anywhere in the world, with the appropriate availability, performance, security and isolation – including the ability to provide self-service for all the necessary compute, storage, network and management capabilities – in cooperation with an ecosystem of supporting partners.

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 2017)

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. They have a track record of successful delivery, significant 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

Some Niche Players may be excellent providers for the use cases in which they specialize, but 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. Providers that specialize in managed services on top of a "good enough" IaaS platform may be in this category. Finally, some Niche Players have weak offerings, or have cloud IaaS businesses with uncertain futures, and should only be chosen with careful attention to managing vendor-related risks.

Context

When people think about "cloud computing," cloud IaaS is often one of the first things that comes to mind. It's the "computing" in cloud computing – on-demand compute, storage and network resources, delivered on-demand, in near-real time, as a service. The market is maturing rapidly; IaaS is on the Slope of Enlightenment on Gartner's "Hype Cycle for Cloud Computing, 2017." However, because the market has consolidated around just two market leaders – Amazon Web Services and Microsoft Azure – many of the other competitors now face significant business challenges, and the customers of those competitors now face significant supplier-related risks.

The stakes involved in this market are increasing because the relevant total addressable market size is increasing; cloud IaaS and PaaS increasingly represent a continuum of integrated services delivered by a single provider, and the leading cloud IaaS providers also have strong PaaS capabilities. IaaS and PaaS represent a spectrum of offerings that balance greater control and customization against greater ease of management and developer productivity. Most customers that adopt the infrastructure resources within a cloud IaaS offering will also adopt associated management services, such as monitoring, and are highly likely to adopt PaaS-level capabilities, such as database as a service, over time.

Consequently, the value proposition of cloud IaaS is no longer simply compute and storage capabilities delivered on-demand, but rather a complete infrastructure platform that delivers both efficiency and agility, combined with unprecedented scalability and global presence. This market direction favors the two incumbent market leaders, and significantly raises the barriers for other vendors trying to gain traction in the market.

Cloud IaaS has broad, mainstream adoption across a wide variety of use cases. While most businesses initially adopted cloud IaaS for Mode 2, agile IT projects, an increasing number of organizations are now migrating Mode 1, safety-and-efficiency-oriented applications – and even entire data centers – to cloud IaaS. Cloud IaaS is increasingly critical not only to digital business, but also to IT modernization and transformation initiatives. Cloud IaaS can now be used for nearly all use cases that can be reasonably hosted on virtualized x86-based servers; the question is no longer "Is cloud IaaS a *viable* solution for my application?" but rather "Is cloud IaaS the *best possible* solution for my application?" Furthermore, cloud IaaS is now a viable alternative to running an internal data center, but it is not the right decision for everyone (see "15 Reasons Not to Migrate Your Data Center to Public Cloud Infrastructure as a Service").

Bimodal IT impacts cloud IaaS sourcing decisions. Mode 2, agile IT organizations typically value cloud IaaS providers that invest deeply in engineering in order to provide a rich suite of features and extensive automation for self-service enablement. Mode 2 adoption is often business-led – driven by business managers who hold the budget, need greater agility and have shorter time frames than IT operations are able to accommodate, and who therefore turn to application developers and enterprise architects for a solution. IT operations organizations typically have a Mode 1 mindset and may initially look for service providers that provide a basic set of IaaS features within a familiar environment that is similar to their existing virtualized infrastructure, but they are likely to rethink this approach if their ultimate goal is IT transformation. Cloud IaaS providers vary in their ability to target these different buying centers. Furthermore, most providers focus on either a Mode 1 or Mode 2 audience, and their feature set and style of service are oriented accordingly, although leading providers offer capabilities attractive to both audiences.

Most organizations now choose one or two long-term strategic partners for cloud IaaS, although they may still use other cloud IaaS providers in a tactical fashion. Most organizations make the choice of which of these cloud IaaS providers to use on a per-project basis, although typically one of the providers is the primary strategic partner and other providers are only used when they are a significantly better fit for the project in question.

Market Overview

Cloud IaaS provides on-demand, near-real-time, self-service access to abstracted, programmatically accessible and highly automated infrastructure resources (at minimum, compute resources, along with associated storage and network resources), on-demand and in near real time. In IaaS, the provider manages the data center facilities, hardware and virtualization, but everything above the hypervisor layer – the operating system, middleware and application – is managed by the customer, or is an add-on managed service from the provider or another third party.

This market is wholly separate and distinct from cloud SaaS, but is increasingly entangled with the PaaS market. Cloud IaaS providers are increasingly offering middleware and other software infrastructure capabilities as a service, as well as services that provision and orchestrate application containers (particularly Docker containers). Customers want to develop, deploy and manage applications efficiently, and will choose the combination of IaaS and PaaS capabilities that best suits their needs – and often, neither customers nor providers will make a definitional distinction

between IaaS and PaaS. To make it easy for applications to span this spectrum of capabilities, an integrated IaaS+PaaS provider needs a single self-service portal and catalog, common identity and access management, an integrated low-latency network context, and an integrated security context.

Cloud IaaS is owned, built and operated by a service provider, but it may be delivered on-premises within a customer's data center or hosted in the provider's data center. It may be "public" (multitenant) or "private" (single-tenant), although, in practice, there is no consistency in the application of these labels to varying degrees of resource isolation, and most hosted offerings use some degree of shared resources in services labeled "private."

Cloud IaaS is not a commoditized service, and even providers with very similar offerings and underlying technologies often have sufficiently different implementations that there is a material difference in availability, performance, security and service features. As a result, risks related to vendor lock-in or application portability need to be thoughtfully managed. (See "Addressing Lock-In Concerns With Public Cloud IaaS" for details.)

What Types of Workload Are Being Placed on Cloud IaaS?

There are four broad categories of customer need in cloud IaaS:

- Digital business enablement

- Mode 2, agile IT projects

- Mode 1, traditional IT data center substitution

- Batch computing

Digital business needs account for the majority of workloads in cloud IaaS. Digital business, however, is not limited to technology companies. Almost every business is being impacted by digital disruption and an increasing number of businesses have "internal startups" or digital business units. (See "Building and Expanding a Digital Business Primer for 2017.") Digital business use cases are very broad, and include digital marketing, e-commerce, e-CRM, SaaS, data services, and Internet of Things applications. These are generally production applications, although cloud IaaS is typically used for the whole application life cycle. Many of these customers have mission-critical needs.

In addition to digital business projects, many organizations have a wide variety of IT projects that they are executing in an agile fashion. Rapid application development, prototyping, experiments and other IT projects that require agility, flexibility and the ability to meet urgent infrastructure needs are frequently executed on cloud IaaS. Although most such Mode 2, agile IT projects are not core to the organization's overall IT portfolio, they may have high visibility and high business impact.

In many organizations, cloud IaaS is gradually replacing or supplementing traditional data center infrastructure. It is typically used very similarly to the organization's internal virtualization environment. Organizations typically begin with development environments or less-mission-critical production applications, but gradually expand to also host mission-critical applications on cloud IaaS. Mode 1, traditional IT organizations typically look to cloud IaaS to deliver cost reductions, but may also be interested in long-term IT transformation. (See "Three Journeys Define Migrating a Data Center to Cloud Infrastructure as a Service" for details.)

The least common need, but one that nevertheless generates significant revenue for the small number of providers that serve this portion of the market, is batch computing. For these customers, IaaS serves as a substitute for traditional HPC or grid computing. Customer needs include rendering, video encoding, genetic sequencing, modeling and simulation, numerical analysis, and data analytics. These customers need to access large amounts of commodity compute at the lowest possible price, with little concern for infrastructure reliability. Some HPC use cases benefit from specialized hardware such as graphics processing units (GPUs) and high-speed interconnects.

Cloud IaaS can now be used to run most workloads, although not every provider can run every type of workload well. Service providers are moving toward infrastructure platforms that can offer physical (nonvirtualized) and virtual resources, priced according to the level of availability, performance, security and isolation that the customer selects. This allows customers to run "cloud native" applications that have been architected with cloud-native principles and design patterns in mind (see "How to Architect and Design Cloud-Native Applications"), as well as to migrate existing business applications from their own virtualized servers in internal data centers into the cloud, without changes. Cloud IaaS is best used to enable new IT capabilities, but it has become a reasonable alternative to an internal data center.

What Key Market Aspects Should Buyers Be Aware Of?

The market has consolidated around two clear leaders. The market consolidated dramatically over the course of 2015. Since 2016, just two providers – AWS and Microsoft Azure – have accounted for almost all of the IaaS-related infrastructure consumption in cloud IaaS, and their dominance is even more thorough if their PaaS-related infrastructure consumption is included as well.

Furthermore, AWS is many times the size of Microsoft Azure, further skewing the market structure. Infrastructure consumption is not equivalent to revenue; these market leaders have less market share in terms of revenue, due to the higher prices charged by competing vendors, as well as the fact that many competitors bundle cloud IaaS with managed services. However, both providers also support robust ecosystems where third parties provide managed and professional services, and therefore their solutions are fully competitive with the providers that deliver bundled managed services. Most customers will choose one of these leaders as their strategic cloud IaaS provider.

The remainder of the market is highly fragmented. Despite the thorough dominance of two market leaders, there are still thousands of service providers that offer cloud IaaS. Some of these are managed hosting providers or local managed service providers, for whom cloud IaaS is simply an infrastructure platform and a means to an end; many such providers are also pivoting to offer their managed services on third-party cloud IaaS offerings. There are also many VPS hosting providers that serve small businesses and have successful cloud VPS offerings; many such providers serve local markets or a single country. However, such providers typically have highly limited capabilities, and most have no supporting ecosystem.

Specialized providers can be the right choice for some use cases. Cloud IaaS providers that have built offerings to serve highly targeted use cases, especially application-specific use cases, can be highly successful within their niche. A deeply differentiated feature set aimed at a particular use case can set a provider apart from the rest of the market, enabling it to win deals even when a customer has chosen one or more strategic providers for their general-purpose workloads.

The market is in a "reboot" phase. Many cloud IaaS providers have recently made strategic shifts in the way that they address this market. Many such providers introduced new or significantly altered cloud IaaS platforms in 2016, or are in the process of doing so during 2017. These new service offerings are unproven, usually have a minimalistic feature set, may have poor operational reliability and lack a supporting ecosystem. Nevertheless, many such providers are aggressively pursuing new customers for these platforms, especially if they have existing customer relationships, and they may be willing to offer generous discounts in order to win customers. Prospective customers of these providers need to carefully manage vendor risks, and to ensure that they speak with reference customers that are similar to them in organization type, IT management style, software development life cycle, workload type, implementation size and tolerance for risk.

The feature gap in the market is growing, not shrinking. It takes considerable time to build the breadth and depth necessary in a competitive feature set, and the market leaders are moving with such velocity that the delta in capabilities between the leaders and everyone else is primarily growing over time, rather than shrinking. It also takes time for cloud IaaS providers to learn to operate and scale their offering, even if they have built large-scale systems in the past. Beware of vendor promises of future features that do not come with firm commitments as to when those features will become generally available.

Most cloud IaaS providers are narrowing their focus. Cloud IaaS providers have increasingly openly acknowledged that they cannot compete directly against the market leaders for public cloud IaaS. Many such providers that have historically have managed hosting businesses have pivoted to offer their managed services on top of market-leading cloud IaaS platforms instead. However, they continue to sell their own cloud IaaS offering in order to meet other customer needs, such as requirements for private cloud IaaS, hybrid infrastructure, VMware-based solutions or data center presence in geographies that are not served by the market leaders. Customers should ensure they are comfortable with their cloud IaaS provider's commitment to the business and level of investment, as well as their potential exit strategy should the provider decide to retire their own cloud IaaS offering.

APIs anchor a partner ecosystem. Programmatic (API) access to infrastructure is crucial, as it enables customers, as well as third parties, to build management tools for their platforms, and to enable applications to take maximum advantage of the infrastructure environment. Providers need to foster rich ecosystems of capabilities. While the leading providers are likely to build a substantial number of capabilities themselves, partners will extend the range of their capabilities, provide overlays for complex heterogeneous multivendor environments, and add "stickiness" to these platforms by offering tight integrations between applications, middleware and infrastructure. Furthermore, cloud IaaS providers that are launching new platforms and hope to "catch up" to the market leaders will be highly dependent upon partners who can supply missing capabilities. Yet the trend is toward proprietary APIs rather than "open" APIs, such as OpenStack.

Many cloud IaaS providers are struggling to build an ecosystem. Cloud IaaS providers need an ecosystem to successfully support a broad array of customer implementations. ISVs need to license and support their software on the cloud IaaS offering. Open-source software needs to be readily available. Tools, especially IT operations management and DevOps tools, must be integrated with the cloud provider's API. Professional services must be available for migration, application

integration and application development on the platform. Managed services need to be delivered in an expert fashion, which requires experience as well as supporting automation tools. Building these ecosystem capabilities takes time. AWS has such a strong ecosystem that Microsoft, despite its overall ecosystem strength, has had challenges building an ecosystem for Azure. Because many potential partners are already occupied with existing integrations and believe the market leaders currently provide enough market opportunity, it is difficult for any additional providers to get an ecosystem foothold.

Vendor relationships are important to some customers, but technical capabilities matter. Some customers that prioritize their relationships with their existing IT vendors may sometimes prefer to obtain cloud IaaS from those vendors, even for use cases where another provider – notably AWS, but sometimes also Azure – can deliver a superior technical solution. However, most customers no longer allow their incumbent vendor relationships to dictate their future cloud strategy, although application platforms and developer relationships are stickier than I&O relationships.

Cloud IaaS is not a commodity. Providers vary significantly in their features, performance, cost and business terms. Although in theory, cloud IaaS has very little lock-in – a VM is just a VM, in the end – in truth, cloud IaaS is not merely a matter of hardware rental, but an entire data center ecosystem as a service. This encompasses the entirety of the ITOM stack, including traditional IT service management capabilities and automation that reduce the burden of operational chores such as patching and backups; DevOps-oriented capabilities; and new forms of automation, analytics and insight (including "smart" infrastructure capabilities) that take advantage of the unique perspective offered by the delivery of integrated compute, storage and networking resources. The more you use those capabilities, the more value you will receive from the offering, but the more you will be tied to that particular service offering. The dynamics of this market resemble a software market, not a traditional IT services market. Providers are in a race to deliver features, and the "winners" are likely to be those that are highly innovative and that have the most resources to invest in the breadth and depth of capabilities development.

Providers' size and scale matter. While scale does impact operational efficiency to some degree, more importantly, it impacts engineering efficiency – the ability to leverage an investment in developers as well as partner capabilities across as large a customer base as possible. Software requires a large upfront investment, but each incremental customer adds comparatively little cost, and software markets tend to become "winner takes all" arenas, where a small number of vendors command dominant market shares. Scale also matters because the ability to deliver a broad range of integrated capabilities will become increasingly crucial. A provider's size, its existing customer relationships and the strength of its brand have an enormous impact on its ability to gain market share and traction, especially on a global basis. Furthermore, the solution ecosystem is rapidly consolidating around a small number of market leaders.

Moving between cloud IaaS providers is challenging. While many customers use multiple cloud IaaS providers, each individual project (or component of a composite application) is typically hosted on a single provider; Gartner refers to this as "multicloud at the point of provisioning." While it is relatively straightforward to move VM images from one cloud to another, truly hybrid multicloud scenarios are rare. "Single pane of glass" management, seamless movement across infrastructure platforms and "cloudbursting" are unlikely to become reality, even between providers using the same underlying

CMP or with use of portable application container technology. Note that the claim that an ecosystem is "open" has nothing to do with actual portability. Due to the high degree of differentiation between providers, the organizations that use cloud IaaS most effectively will embrace cloud-native management, rather than allow the legacy enterprise environment to dictate their choices.

Customers increasingly use third-party management tools for governance, especially multicloud governance. The largest customers use third-party management tools to supplement the native management capabilities of the providers, and these tools are strongly recommended to all customers that intend to make substantive use of cloud IaaS. Management tools can be very helpful for governance functions, and may be designed for single-cloud or multicloud use. If multicloud, they should support integrated cost management, identity and access management, security and compliance reporting, and networking. Management tools cover a wide range of possible functions – from CMPs such as CliQr and Scalr, to cloud service expense management (CSEM) tools such as RightScale Cloud Analytics and Cloudability, to continuous configuration automation tools such as HashiCorp's Terraform. (See "Innovation Insight for Cloud Service Expense Management Tools" for more on CSEM tools.)

Customers are decoupling the choice of cloud IaaS offering from managed service decisions. Customers increasingly choose the cloud platform that is best for their workload, and then seek an MSP to manage it, rather than adopting a "managed cloud" offering from an MSP that can offer only basic IaaS capabilities on its own platform. Customers may also extend existing outsourcing relationships to include management of a third-party cloud IaaS offering. While some Mode 1, traditional IT customers consider it acceptable for an MSP's platform to offer only a basic set of IaaS features, it is generally unacceptable to Mode 2, agile IT customers. Furthermore, such deficiencies have a long-term impact on the quality and cost of the customer's IT operations, which may be strategically unacceptable to Mode 1 customers.

Customers seek third-party managed and professional services for best-in-class cloud IaaS offerings. Some MSPs specialize in cloud-native operations, usually with significant use of DevOps, and can help customers through the transformation process, which may be attractive to both Mode 1 and Mode 2 customers, as well as digital businesses (see "Use Managed and Professional Services to Improve Cloud Operations for Digital Business"). Mode 1 data center migrations also benefit strongly from managed and professional services, even if the approach is not cloud-native (see "Three Journeys Define Migrating a Data Center to Cloud Infrastructure as a Service"). See "How to Choose a Managed Service Provider for a Hyperscale Cloud Provider" for MSP selection guidance, and the "Magic Quadrant for Public Cloud Infrastructure Managed Service Providers" for a market evaluation.

Local sourcing matters to some customers. Customers normally prefer to keep data in-region for reasons of network latency. However, regulatory concerns that require keeping data in-country, as well as revelations about foreign intelligence agencies obtaining access to private data, have heightened the desire of non-U.S.-based customers to purchase cloud IaaS from local providers. (See "The Snowden Effect: Data Location Matters.") Unfortunately, local providers typically lack the scale and capabilities of the global providers, and may focus primarily on small businesses, not enterprises. Furthermore, keeping data local is no guarantee of freedom from either domestic or

foreign surveillance. It is nevertheless possible that the cloud IaaS markets in Europe and Asia will become highly fragmented, which may result in only basic, commodity capabilities being available to customers that cannot use a foreign provider (even when that provider has local presence).

Public cloud IaaS provides adequate security for most workloads. Although many security controls are the responsibility of the customer, not the provider, most major cloud IaaS providers offer a high degree of security on the underlying platform. Transparent encryption of LAN, WAN and storage will become increasingly commonplace as a bundled element of cloud IaaS offerings, as providers react to defend themselves against intrusion from government entities. (See "Take a Risk-Based Approach to Public Cloud IaaS" for guidance.)

Customers do not always save money by using cloud IaaS. Although many customers first investigate using IaaS to achieve cost savings, most customers buy IaaS to achieve greater business agility or to access infrastructure capabilities that they do not have within their own data center. IaaS can drive significant cost savings when customers have short-term, seasonal, disaster recovery or batch-computing needs. It can also be a boon to companies with limited access to capital and to small companies – especially startups – that cannot afford to invest in infrastructure. For larger businesses with existing internal data centers, well-managed virtualized infrastructure, efficient IT operations teams and a high degree of automation, IaaS for steady-state workloads is often no less expensive, and may be more expensive, than an internal private cloud. The less efficient your organization, the more likely you are to save money by using a cloud provider, especially if you take advantage of this opportunity to streamline and automate your operations. The largest-scale providers are continually lowering their prices, and automated managed services will substantially drive down the cost of infrastructure management over time, so cost advantages will continue to accrue to the providers. (See "Can You Save Money Migrating to Cloud IaaS?" for guidance.)

?

China Context

Market Differentiators

China is a unique market for cloud IaaS. Regulation restricts foreign cloud IaaS providers from directly operating their services in China; instead, they must partner with a Chinese operator. Furthermore, the domestic market is on a different trajectory for market maturity than the already-mature Western markets; the China market is not only less mature, but has different adoption patterns. Consequently, China has a distinctive cloud IaaS provider landscape.

This contextualization addresses the needs of organizations that are based outside of China, but intend to use cloud IaaS to host applications in China. Such organizations can use either a non-Chinese cloud IaaS provider with a China region, or a China-based cloud IaaS provider. China-based organizations seeking a domestic China perspective should consult "Market Guide for Cloud Infrastructure as a Service, China" and "Overcome the Three Key Challenges of Adopting Cloud Services in China" instead.

Regulatory Environment

Customers should be aware of the following distinctive features of the regulatory environment:

Foreign cloud providers must partner with a Chinese operator. The 2015 Classification Catalogue of Telecommunications Services officially took effect in March 2016. The 2015 Catalogue creates regulatory restrictions on foreign cloud service providers' direct-operated model in China and forces them to use a China-partner-operated model. This model has been adopted by all foreign cloud IaaS providers, including Amazon Web Services (AWS) and Microsoft Azure.

Foreign cloud providers must segregate their Chinese service. Foreign providers are not allowed to interconnect their China data centers with those outside of China (though their in-China data centers can be connected to each other). Furthermore, foreign providers must give customers separate credentials for the China service, distinct from any credentials the customers may be issued for the global service; most providers handle this restriction by requiring customers to have separate China-specific accounts. China-based cloud service providers are not subject to these restrictions, thus allowing them to offer a seamless global service if they so desire.

China's cybersecurity law mandates data sovereignty. China passed a new cybersecurity law in November 2016, which took effect in June 2017. Even though it is unclear how the rules in this law will be implemented in China, it affects both domestic and foreign firms operating in China, and covers a wide range of activity relating to use of the internet and, more broadly, information and communications (ICT) technologies. One important clause in the law governs data sovereignty; it requires that all personal and critical data that originates in China must be physically stored inside China. Thus, multinational organizations must pay close attention to their cross-border business data flows.

The Multi-Level Protection Scheme (MLPS) provides a regulatory standard for security of critical IT systems. China's MLPS regulation classifies IT systems into five levels of escalating importance, and specifies security controls for each level. The MLPS is applied to government systems as well as other systems considered important to the economy; this includes cloud IaaS providers. Therefore, most cloud IaaS providers will have undergone a third-party audit for a particular MLPS level. MLPS Level 3 is sufficient to allow an offering to be used by Chinese state-owned enterprises.

Public websites require an Internet Content Provider (ICP) number. Public websites that are physically hosted in China may be required to obtain an ICP number. Customers should consult with their legal counsel to determine whether or not their use of cloud IaaS in China requires them to apply for an ICP number. Cloud IaaS providers may deny service to customers who are noncompliant with these regulations.

Competitors sometimes claim that foreign cloud providers are not operating legally in China, but such claims are untrue. Both AWS and Azure have segregated Chinese regions that are operated by partners, fully comply with current regulations for foreign cloud providers, have MLPS Level 3 approval, and can be used by both domestic and international customers.

Foreign vs. Local Competition

Global providers, such as AWS and Azure, have comprehensive integrated IaaS and PaaS (IaaS+PaaS) offerings, but generally do not offer the entire portfolio of their services in China. Furthermore, their China operating partners need to develop experience and maturity with

supporting complex enterprise workloads in these highly sophisticated platforms.

Some of the Chinese local providers, notably Alibaba Cloud, are also pursuing an IaaS+PaaS strategy, and have the advantage of a larger and more geographically dispersed cloud data center footprint within China. However, Chinese local providers also need time to mature their offerings to meet the requirements of enterprise business workloads, especially with regard to security, availability and the ability to support Mode 1 traditional business applications, and not only Mode 2 and digital business workloads.

Considerations for Technology and Service Selection

Global organizations that are deploying applications into cloud infrastructure in China should consider the following factors that are unique to the Chinese market:

- Internet latency challenges
- Security and data privacy challenges
- Availability of managed service providers (MSPs)
- Support for enterprise requirements

These factors are described in greater detail below. Additional information about infrastructure delivery options within China can be found in "Select From Three Infrastructure Options for Global Companies Running Business in China."

Internet Latency Challenges

Most consumers use the internet to access content and applications hosted on public cloud IaaS providers. Therefore, many organizations that use public cloud IaaS depend on the internet for connectivity to their customers. Furthermore, advancements in software-defined WAN (SD-WAN) technology have resulted in more enterprises leveraging internet connectivity to lower the WAN costs associated with cloud connectivity and branch-office connectivity.

However, in China, careful attention must be paid to internet performance. Throughput is impacted by bandwidth bottlenecks in multiple locations, including between a carrier's cross-provincial networks, at cross-carrier peering points within China and at cross-border transit routers at the international points-of-presence of Chinese carriers. Furthermore, China's Great Firewall (GFW) adds sizable additional latency for traffic that transits the GFW.

Prospective customers should assess the cloud IaaS provider's network architecture, and determine how the provider's network capabilities can assist them in reducing the impact of internet bottlenecks and the GFW. Customers should measure network latency as part of their application performance monitoring, and consider use of WAN optimization technology if necessary. Customers should also consider whether private WAN connectivity is a feasible alternative to the internet for their use case.

Because Chinese providers are allowed to directly interconnect all of their data centers, but foreign providers must segregate their Chinese data centers from the rest of their network, customers that use a Chinese provider may be able to lower their cloud IaaS-related network costs by using the

provider's network for transit between China and the rest of the world, rather than purchasing separate international connectivity from a carrier. The use of SD-WAN and VPNs should also be considered.

These internet issues mean that data center proximity to customers is valuable. Customers that can support a distributed application architecture should consider selecting providers with a broader distributed data center footprint and deploying into multiple geographically dispersed data centers in order to move applications closer to customers, and thus improve performance. Other techniques, such as use of a content delivery network (CDN), may also help to address the latency issues. However, foreign providers are restricted from directly offering CDN services in China; they must use a partner.

Security and Data Privacy Challenges

Many foreign customers perceive that deployment into China creates increased risk and challenges in cloud-related security and data privacy. Most large Chinese enterprises do not place critical or sensitive data into cloud IaaS; they may deploy consumer-facing applications with nonsensitive data in cloud IaaS, but store the remaining data in their own data centers or in external hosting. Foreign customers are more likely to simply use cloud IaaS and take additional security precautions, rather than divide their applications between cloud IaaS and hosting.

Gartner believes that the encryption of sensitive data should be considered a mandatory best practice in cloud IaaS, regardless of region; most cloud providers offer encryption capabilities, and customers can also implement encryption themselves. (See "How to Make Cloud IaaS Workloads More Secure Than Your Own Data Center" for details.) More broadly, customers should ensure that they have a comprehensive workload protection strategy; see "Market Guide for Cloud Workload Protection Platforms" for additional guidance.

As is always true for cloud IaaS, security and privacy are shared responsibilities. The provider may offer security controls, and document its operational processes and policies, but it is up to the customer to verify these claims, correctly configure controls (and implement additional controls if necessary) and manage its own security. Customers should ask to see third-party audits, and if feasible, consider doing an on-site audit themselves. Customers working with global providers – that therefore have a Chinese partner that is responsible for operating the offering – should ask for China-specific audits, since there will be differences between what is done in China and the global service.

MSP Availability

Global companies that want cloud managed services should likely use a global MSP to deliver these services. However, not every MSP that supports a given cloud provider will support that provider's Chinese locations, and MSP services that might be available in other regions, such as hybrid colocation or hybrid hosting services, might not be available in China. A handful of global MSPs support Alibaba Cloud, but other domestic Chinese cloud IaaS providers are typically not supported by global MSPs. (See "Magic Quadrant for Public Cloud Infrastructure Managed Service Providers" and "Magic Quadrant for Cloud-Enabled Managed Hosting, Asia/Pacific" for additional guidance.)

Support for Enterprise Requirements

Chinese providers lag the global market in their support for enterprise requirements, especially the requirements of customers with non-cloud-native applications that want to be able to host those applications on cloud IaaS with no or limited modifications. Enterprises often have highly complex and diverse needs, and do not want to significantly refactor applications, and thus require support for a broad array of compute instance and storage configurations, as well as flexible and highly configurable networking. Enterprises also often need their independent software vendors (ISVs) to explicitly support deployments on a particular cloud provider in order to be in compliance with their support agreements.

Many Chinese providers are still maturing and evolving their cloud IaaS offerings. Many such providers were originally consumer-centric online businesses that subsequently entered the cloud IaaS market; their initial customers were startups and digital businesses similar to themselves. These providers are now progressing toward the traditional enterprise computing market. However, many such providers are missing the feature set necessary to provide a cloud IaaS environment that offers all the infrastructure capabilities that enterprises expect. Consequently, customers should carefully assess the ability for a cloud IaaS provider to meet their requirements. (Use the "Evaluation Criteria for Cloud Infrastructure as a Service" for guidance.)

Notable Vendors

Vendors included in this Magic Quadrant Perspective have customers that are successfully using their products and services. Selections are based on analyst opinion and references that validate IT provider claims; however, this is not an exhaustive list or analysis of vendors in this market. Use this perspective as a resource for evaluations, but explore the market further to gauge the ability of each vendor to address your unique business problems and technical concerns. Consider this research as part of your due diligence and in conjunction with discussions with Gartner analysts and other resources.

Three of the cloud IaaS providers in this global Magic Quadrant have Chinese operations: Alibaba Cloud, Amazon Web Services (AWS) and Microsoft Azure. AWS and Azure deliver their Chinese regions in conjunction with Chinese partners, and thus their offerings in China are somewhat different than their global offerings. Alibaba Cloud is China-based and offers a much broader and deeper cloud portfolio within China than it does in its international regions. We explore these differences in the vendor profiles below.

For a list of additional Chinese cloud IaaS providers, consult "Market Guide for Cloud Infrastructure as a Service, China." That Market Guide covers the following additional vendors:

CIB Financial Services (CIB Fintech)

QingCloud

Tencent Cloud

UCloud

Although China Telecom and Huawei are not included in the Market Guide, they too are important Chinese cloud IaaS providers that should be considered when selecting a provider.

Alibaba Cloud

Alibaba Cloud is the current cloud IaaS market share leader in China, and is classified as MLPS Level 4. Its mainland China data centers are located in Beijing, Hangzhou, Qingdao, Shanghai, Shenzhen and Zhangbei. Alibaba Cloud has a broader scope of services within China than it does internationally. Its domestic China service offers an array of integrated IaaS+PaaS capabilities. Its proprietary technology platform, Apsara, is used as the basis of the private cloud solutions that Alibaba Cloud offers within China, thus allowing Alibaba Cloud to also offer hybrid cloud solutions.

Amazon Web Services

AWS opened a region in Beijing in 2014, but originally pursued a route to market in China that did not involve use of a partner; it operated this region while it was in beta. As the regulatory environment evolved, it was eventually obliged to use a partner; since August 2016, its Beijing region, operated by Sinnet, has been in general availability. Sinnet operates and sells the offering in China. AWS provides the technology, guidance and technical support to Sinnet. AWS China has MLPS Level 3 approval.

AWS China (Beijing region) is segregated from the rest of AWS, and customers must obtain China-specific accounts. Unlike AWS's global service, AWS China only accepts business customers; individual users cannot sign up for the service.

AWS's most popular IaaS and PaaS offerings are available within AWS China, but, similar to other AWS regions, not all offerings are available within the region. Global customers should ensure that the services they need are present within the region.

Microsoft Azure

In 2013, Microsoft became the first multinational company to make its public cloud services generally available inside China, although Azure did not become generally available in China until 2014. As regulation requires, Microsoft uses a partner, 21Vianet, which operates and sells Microsoft cloud services (including Azure and Office 365) in China. Azure China has MLPS Level 3 approval and two regions – one in Beijing and one in Shanghai.

Azure China is segregated from the rest of Azure, and customers must obtain China-specific accounts. Azure China can support both business and individual accounts.

Azure's most popular IaaS and PaaS offerings are available within Azure China – comparable to its other Asia/Pacific regions. There are China-specific API endpoints; these are treated as custom endpoints by developer tools such as Visual Studio. Furthermore, Azure Active Directory is somewhat different in China than it is in global Azure. Global customers should ensure that the services they need are present within the region they intend to use, as well as identifying what other changes they may need to make in order to use Azure China.

?

Europe Context

Market Differentiators

According to Gartner's 2017 CIO Survey, respondents in EMEA see cloud services as more strategic than their U.S. counterparts, especially in the areas of differentiation from the competition. In the same survey, Europeans see digital security at a higher potential to change the organization over the

next five years (see "Highlights of the 2017 CEO Survey: CIOs Must Scale Up Digital Business").

Local data center presence is important in Europe. There is a strong cultural preference for data to be hosted in-country. Furthermore, the regulatory environment favors strong data privacy. EU businesses must currently comply with the EU's Data Protection Directive. In 2018, the General Data Protection Regulation (GDPR) will take effect; it protects the rights of EU citizens and their personal data. E-privacy regulation is also making its way through the EU's legislative process.

In Europe, midmarket and small or midsize business (SMB) customers are more reliant on channel partners to help them shape infrastructure strategies and adopt cloud services. Customers typically favor local channel partners, further increasing the importance of a service provider's in-country presence.

Finally, the strength of the U.S. dollar against European currencies has resulted in a price increase for some cloud infrastructure as a service (IaaS) offerings. In some cases, providers state all prices in U.S. dollars, and the price in European countries therefore floats with the foreign exchange markets. In other cases, providers price in local currency, but periodically rebenchmark prices against the U.S. dollar. This could strain IT budgets in the coming years and potentially lead to a slowing down of cloud adoption in comparison to the U.S., although Gartner research does not currently show that the price increases are having a material impact on adoption.

Hyperscale cloud providers (all of whom are based in the U.S.) still dominate overall European market share in cloud IaaS. They have steadily expanded their European footprint into multiple countries. They comply with the EU-U.S. Privacy Shield Framework, and offer EU Model Contract Clauses for customers subject to the Data Protection Directive. However, local providers – who have the advantages of local language support, direct sales and channel-partner presence, and data center presence – continue to be competitive, especially in the SMB market.

Considerations for Technology and Service Selection

Prospective cloud IaaS customers must be aware of data sovereignty and residency requirements for their business operations and customers. Some countries have higher sensitivities to data residency than others. When placing an application in cloud IaaS, ensure that the geographical location of the cloud IaaS data center complies with your data residency requirements. Be conscious of any automated data replication that the provider is performing, and ensure that the geographic scope of that replication is within your data residency requirements.

Ensure that the provider can help you comply with the Data Protection Directive. Also, ensure that the provider will be able to help you comply with the GDPR when it takes effect; some providers currently have more developed policies on being a data processor than others. By 25 May 2018, all providers handling the data of EU citizens will need to have policies in order to assist data controllers to be compliant with GDPR. See "Focus on Five High-Priority Changes to Tackle the EU GDPR" for recommendations.

Notable Vendors

Vendors included in this Magic Quadrant Perspective have customers that are successfully using their products and services. Selections are based on analyst opinion and references that validate IT provider claims; however, this is not an exhaustive list or analysis of vendors in this market. Use this

perspective as a resource for evaluations, but explore the market further to gauge the ability of each vendor to address your unique business problems and technical concerns. Consider this research as part of your due diligence and in conjunction with discussions with Gartner analysts and other resources.

Many of the cloud IaaS providers in the Magic Quadrant have a European presence. Additional evaluations of European providers can be found in "Magic Quadrant for Managed Hybrid Cloud Hosting, Europe" ; many of those providers have their own cloud IaaS platforms.

Some additional European providers of note, listed by the country in which they are headquartered, include:

France: Orange, OVH

Germany: T-Systems, 1&1

Italy: Enter, Sparkle

Netherlands: LeaseWeb, Schuberg Philis

Spain: Gigas, Telefonica

Switzerland: CloudSigma, Exoscale

?

Government (Federal) Context

Market Differentiators

A cloud infrastructure as a service (IaaS) solution designed for U.S. federal government customers ("federal IaaS") is expected to be Federal Risk and Authorization Management Program (FedRAMP (<https://www.fedramp.gov/>))-compliant after proving that they have National Institute of Standards and Technology (NIST)-based security controls to support the targeted levels of data classifications. In practice, however, agencies are also expected to use a risk-based approach to security, and can issue waivers to the FedRAMP process.

FedRAMP Compliance

Cloud IaaS providers have three paths to achieving FedRAMP compliance: Joint Accreditation Board Provisional Authority to Operate (JAB P-ATO), agency-sponsored Authority to Operate (ATO) and security assessment package ("CSP Supplied Package"), assessed by a FedRAMP-accredited Third Party Assessment Organization (3PAO). While all forms of compliance are considered valid for federal IaaS, the JAB approach has been popular to the point where the FedRAMP Program Management Office (PMO) is now capping the number of providers it can support by selecting a set number of business cases. This is because the JAB path involves the Department of Defense (DoD), the General Services Administration (GSA) and the Department of Homeland Security (DHS). Providers have inferred that having the three biggest spenders in the federal market involved would accelerate reach, but it has also created a bottleneck. Although the point of the program is to "do

once and reuse many times" by leveraging reciprocity, in practice, some providers have reported meeting with resistance when leveraging an agency ATO with other agencies, so early due diligence is recommended.

Because FedRAMP is the expected standard in this market, but acquiring an ATO is a difficult, expensive and lengthy process, the number of federal public cloud IaaS providers currently stands at 30 already authorized, with another five in process. Our discussions with vendors that have completed the process suggest an average of 18 months and \$3.5 million to go through the process. The FedRAMP PMO estimates the ongoing cost of maintenance at \$1 million per year. This has led to increasing dissatisfaction with the FedRAMP PMO, particularly for small providers, and it is working on streamlining the process as a result. Because the FedRAMP certification process is lengthy, providers may be flagged as "in-process" of certification, and they should not be discounted, particularly if agencies can act as sponsors to accelerate the process.

Because agencies are expected to make risk-based decisions, some have issued waivers to the FedRAMP requirement if they felt a cloud solution was secure enough for its intended use. This was underscored by the NASA CIO in response to a 7 February 2017 Inspector General (IG) report: "NASA will use FedRAMP-approved cloud services whenever available. Otherwise, NASA will perform an appropriate risk assessment, and may make a risk-based decision to approve the service for use at NASA."

However, as FedRAMP ATOs find their way into an increasing number of RFP requirements, failure to comply will increasingly become a barrier to market competitiveness.

Note that some cloud IaaS providers have a very broad set of solutions, which may include platform as a service (PaaS) capabilities. Providers will normally be specific about which parts of their solutions have been assessed by a 3PAO, and which services are FedRAMP-approved. Federal IaaS solutions normally adhere to International Traffic in Arms Regulations (ITAR) restrictions, and can support adherence to additional public-sector-related compliant practices, such as Health Insurance Portability and Accountability Act (HIPAA) for health-related data and Criminal Justice Information Services (CJIS) for public-safety-related data.

The FedRAMP framework now aligns to the Federal Information Security Management Act (FISMA) Low, Moderate and High levels. The FISMA/FedRAMP "High" baseline, to handle more sensitive information, was introduced in 2016, and the following IaaS providers are authorized at that level: Amazon Web Services (AWS) GovCloud High, CSRA ARC-P Infrastructure-as-a-Service, Microsoft Azure Government and Microsoft Global Foundation Services for Government (GSGO).

Department of Defense and Other Sensitive Requirements

Because of its unique security sensitivity, the Department of Defense (DoD) is adding additional requirements to FedRAMP, called FedRAMP+, and several IaaS providers have received a P-ATO from the Defense Information Systems Agency (DISA). The additional controls are listed in a DoD Cloud Security Requirements Guide (SRG (http://iase.disa.mil/cloud_security/Pages/index.aspx)). In addition, DISA has reviewed and approved a handful of vendors to support Level 5-type data (Controlled Unclassified Information and unclassified National Security Systems) for DoD use.

The DoD cloud security strategy is being flowed down to all DoD providers (including industry and academia) via a Defense Federal Acquisition Regulation Supplement (DFARS) interim rule released in August 2015. The new rule requires vendors to state whether they "anticipate" using cloud computing services in the performance of a particular DoD contract, and imposes security requirements and limitations on access and disclosure of government data and government-related data maintained by the contractor pursuant to a cloud computing services contract. In addition, it mandates providers' compliance with NIST SP 800-171

(<http://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-171.pdf>) (the Special Procedure document from the National Institute of Standards and Technology dealing with protecting controlled unclassified information in nonfederal information systems and organizations), a move that is driving an increasing number of them to evaluate whether leveraging FedRAMP-certified IaaS providers would help them comply.

The Intelligence Community agencies have been driving to an Intelligence Community IT Environment (ICITE) to share data, based in the cloud. In 2014, AWS won a \$600 million contract to deploy and operate an isolated private region for the ICITE community; this region is now in production. Other large agencies have also considered pursuing similar community-cloud deals with leading cloud IaaS providers.

Who Can Use Federal IaaS Solutions?

Federal IaaS is often, but not always, delivered from data centers that are specifically for government customers. When such solutions are hosted in the same data centers as those used for commercial customers, the federal IaaS solutions are physically and logically segregated from the commercial solutions. These federal solutions are operated by U.S.-based personnel.

Use of federal IaaS solutions is normally restricted to U.S. federal government customers, but, in most cases, contractors and other third parties performing work on behalf of government agencies can also use them. Notably, SaaS providers with government customers are often allowed to use these solutions, and federal IaaS providers are actively courting them for this purpose. A pending FedRAMP Tailored process – developed to authorize SaaS vendors at low impact in weeks, rather than months – will accelerate this more.

U.S. state, local and tribal governments are increasingly interested in using federal IaaS solutions, although this varies by provider – it is the provider that decides which customers are permitted in their cloud. Such entities may need to find commercial cloud IaaS solutions instead. Even if they do not want to share the environment, state, local and tribal entities use FedRAMP compliance as a proxy for determining whether a cloud IaaS offering can adequately meet government security requirements; since leading providers also have FedRAMP-approved commercial regions, this is a viable approach for such entities.

Considerations for Technology and Service Selection

Government organizations contemplating the use of cloud IaaS should pay careful attention to bimodal IT requirements. Government IT personnel frequently have a cautious mindset and a tendency to operate only Mode 1 reliable IT, but they increasingly face agile demands that are better-served with Mode 2 agile IT. Government IT organizations often attempt to source cloud IaaS in a Mode 1 fashion, even if the primary need is agility; such an approach is unlikely to fully satisfy users.

Conversely, government IT organizations that are trying to drive Mode 1-oriented cost reductions need to source differently from those whose primary needs are agility and transformation. When selecting an offering, it is vital to keep bimodal requirements in mind.

Government customers should be careful to distinguish between cloud IaaS and more traditional forms of outsourcing. Many so-called "cloud" solutions that are marketed to government entities are simply "cloudwashed" outsourcing; they often come with long-term contracts and relatively inflexible capacity constraints, and lack the automation and industrialization of true cloud IaaS offerings (for more details, see "Don't Be Fooled by Offerings Falsely Masquerading as Cloud Infrastructure as a Service").

Some government customers find it difficult to contract directly with cloud IaaS providers, due to government acquisition rules. There are third-party cloud service brokerages – usually managed service providers (MSPs) and system integrators (SIs), such as Accenture, CSRA, Datapipe and Smartronix – that resell major cloud IaaS providers' solutions, and that typically add significant value; these are probably the best choice for government entities that need a brokered solution. See "Getting Help With Implementing Cloud IaaS" for details on how these brokers can add value, and "Magic Quadrant for Public Cloud Infrastructure Managed Service Providers, Worldwide" for guidance on selecting a broker.

There are other brokers that may add little or no value, but are willing to resell a major cloud IaaS provider's solution while offering much more attractive contract terms, usually with far higher liability caps. Government customers should beware of such resellers, as they frequently lack the financial solvency to fulfill their obligations. Government customers may also want to buy through an 8(a) company in order to fulfill procurement obligations; again, they should be careful, as such resellers often mark up the price without adding value.

The GSA has recognized this lack of maturity in the federal market around cloud procurements, and is in the process of standing up a team dedicated to helping agencies adopt cloud solutions. The intent is for this team to help agencies select the best options for their needs, as well as to support the acquisition process.

Notable Vendors

Vendors included in this Magic Quadrant Perspective have customers that are successfully using their products and services. Selections are based on analyst opinion and references that validate IT provider claims; however, this is not an exhaustive list or analysis of vendors in this market. Use this perspective as a resource for evaluations, but explore the market further to gauge the ability of each vendor to address your unique business problems and technical concerns. Consider this research as part of your due diligence and in conjunction with discussions with Gartner analysts and other resources.

These service providers all have a multitenant cloud IaaS platform that they own and operate. We do not mention any FedRAMP-authorized managed services, private cloud implementations or solutions built using third-party cloud IaaS providers.

Amazon Web Services

AWS's GovCloud (U.S.) is a community cloud dedicated to the U.S. federal government, including contractors, third parties and SaaS providers providing services to federal customers. It is an isolated region located in Oregon. It has a FedRAMP High JAB P-ATO and multiple FedRAMP Moderate agency ATOs. It has Level 2 and 4 authorizations for the DoD SRG, and is pursuing Level 5. It adheres to ITAR regulations and is managed by U.S. Persons only. It can be used for workloads that must adhere to CJIS, IRS 1075, Export Administration Regulation 99 (EAR99) or Family Educational Rights and Privacy Act (FERPA) requirements. Most AWS services are available in GovCloud, although not all services have been assessed by a 3PAO. There is also an agency ATO that covers AWS's U.S. East and West coast regions, although only a subset of AWS's offerings have been assessed by a 3PAO.

AWS has the largest market share in cloud IaaS for government customers. It serves federal customers both in GovCloud and its commercial regions. It also serves a significant customer base of state and local customers from its commercial regions. AWS has a rich ecosystem of partners, and many government customers adopt AWS through an MSP or SI.

CSRA

CSRA's ARC-P Government Community Cloud is a community cloud for the U.S. federal government, including vendors providing services to federal customers. ARC-P originated with Autonomic Resources, which was acquired by CSRA in 2015. ARC-P is located in Equinix data centers in Georgia and Virginia, with additional cloud services delivered from personnel located at CSRA's Integrated Technology Center in Louisiana. ARC-P can also be deployed on-premises. It has a FedRAMP High JAB P-ATO, as well as a Level 2 provisional authorization for the DoD SRG. It adheres to ITAR regulations and is managed by U.S. Persons with a minimum of U.S. Federal Moderate Background Investigation (MBI). It can be used for workloads that must adhere to CJIS and IRS 1075 requirements. ARC-P is also used to host CSRA's FedRAMP High-compliant PaaS offering, ARCWRX, which is based on Red Hat OpenShift.

IBM

SoftLayer Federal Cloud (SFC) provides cloud IaaS and hosting services in data centers dedicated to U.S. federal government workloads. The two data centers are located in Texas and Virginia. SFC obtained a FedRAMP Moderate JAB P-ATO in late 2016. It also has an agency ATO from the Federal Communications Commission. The ATOs are associated with SoftLayer Technologies, an IBM company.

Microsoft

Microsoft Azure Government is a community cloud for U.S. federal, state and local government customers, as well as for qualified partners serving those entities. It has four regions, located in Northern Virginia, Iowa, Texas and Arizona. It has a FedRAMP High JAB P-ATO, and adheres to ITAR regulations. It has Level 2 and 4 authorizations for the DoD SRG, and meets DFARS requirements. Microsoft also has two regions dedicated to the DoD, which support DoD SRG Impact Level 5 workloads, and which have achieved DoD Impact Level 5 provisional authorization. It can be used for workloads that must adhere to CJIS, IRS 1075, EAR99 or FERPA requirements. Many Azure services are available in the Azure Government Cloud, although the newer, higher-performing

compute instance types are not yet available in these regions. The main Microsoft Azure service also has a FedRAMP Moderate JAB P-ATO, although only a subset of Azure services has been assessed by a 3PAO.

Microsoft has been aggressively pursuing government customers, including trying to sign deals at the state level that make it easier for state and local entities to adopt Azure. There are many ways in which government customers can purchase Azure services, which eases what is sometimes a complex procurement process.

QTS

QTS Government Cloud is a community cloud for the U.S. public sector. This cloud service originated as the VMware vCloud Government Service (vCGS), a partnership between VMware and Carpathia. In 2015, QTS acquired Carpathia, which had a long history as a managed hosting provider for government customers. In late 2016, QTS acquired sole rights to vCGS from VMware. During the VMware partnership, it used the same architecture as VMware's own vCloud Air and met the same quality metrics, but QTS now wholly owns and controls the service, and can determine its future. Customers have a choice of two data centers, one in Northern Virginia and one in Arizona. QTS Government Cloud has a FedRAMP JAB P-ATO, and can be used for workloads that must adhere to CJIS requirements.

Verizon

Verizon's Enterprise Cloud: Federal Edition is a community cloud for U.S. federal government customers. It is based on the Terremark Enterprise Cloud, which Verizon acquired in 2011. It is deployed in Northern Virginia, and has an agency ATO from the U.S. Department of Health and Human Services.

Virtustream

Virtustream's Federal Cloud (VFC) is a community cloud for U.S. federal, state and local government customers. It is similar to Virtustream Enterprise Cloud (VEC), with Northern Virginia, Philadelphia-area and San Francisco-based isolated deployments of Virtustream's xStream platform. VFC has a FedRAMP JAB P-ATO, and can be used for workloads that must adhere to CJIS and IRS 1075 requirements. Virtustream has two more isolated nodes in Northern Virginia and Las Vegas, which serve commercial contractors that support DoD customers and adhere to ITAR regulations; these nodes are independent from both VFC and VEC, but do not have a service brand of their own. In addition, Virtustream's Viewtrust solution can be used for continuous monitoring and on-demand compliance reporting.

Evidence

Gartner client inquiries in 2016 and 2017 (currently more than 1,000 cloud IaaS-related inquiries per quarter)

Service provider interviews and product demonstrations in 2016 and 2017

Surveys of more than 75 cloud IaaS providers in 2016 and 2017

Customer references from the service providers in 2016 and 2017

Hands-on trials of service offerings in 2016 and 2017

Performance benchmarks and monitoring of service offerings via Gartner's Cloud Tech Planner service, throughout 2016 and 2017

Public information from sources such as U.S. Securities and Exchange Commission filings, press releases, vendor websites and community support forums

Note 1 SSAE 16

Statement on Standards for Attestation Engagements (SSAE) 16 – that is, Service Organization Control (SOC) 1. See "Market Guide for Organization Security Certification Services."

Note 2 ISO 27001

International Organization for Standardization (ISO)/International Electrotechnical Commission (IEC) 27001. See "Security Research Roundup for ISO 27001 Compliance."

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.

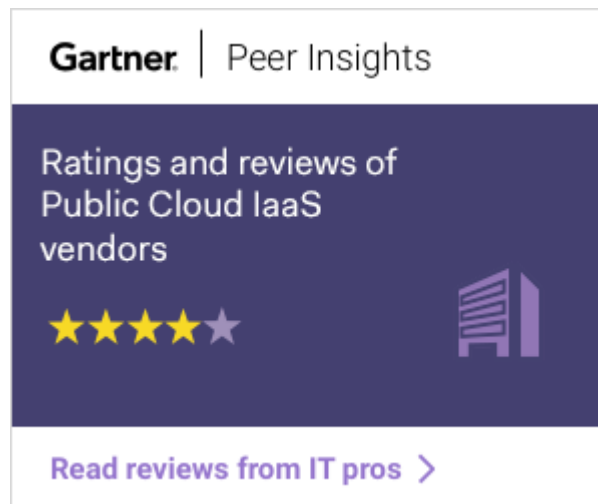
Offering (Product) Strategy: The vendor's approach to product development and delivery that emphasizes differentiation, functionality, methodology and feature sets as they map to current and future requirements.

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.



The banner features the Gartner logo and 'Peer Insights' text at the top. Below, it reads 'Ratings and reviews of Public Cloud IaaS vendors' with a 4.5-star rating (four yellow stars and one grey star) and a server rack icon. A call-to-action button at the bottom says 'Read reviews from IT pros >'.

(https://www.gartner.com/reviews/market/public-cloud-iaas?refval=mq_reprint&campaign=mq_reprint&content=Public_cloud_iaaS)

© 2017 Gartner, Inc. and/or its affiliates. All rights reserved. Gartner is a registered trademark of Gartner, Inc. or its affiliates. This publication may not be reproduced or distributed in any form without Gartner's prior written permission. If you are authorized to access this publication, your use of it is subject to the Usage Guidelines for Gartner Services (/technology/about/policies/usage_guidelines.jsp) posted on gartner.com. 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 and shall have no liability for errors, omissions or inadequacies in such information. This publication consists of the opinions of Gartner's research organization and should not be construed as statements of fact. The opinions expressed herein are subject to change without notice. Gartner provides information technology research and advisory services to a wide range of technology consumers, manufacturers and sellers, and may have client relationships with, and derive revenues from, companies discussed herein. Although Gartner research may include a discussion of related legal issues, Gartner does not provide legal advice or services and its research should not be construed or used as such. Gartner is a public company, and its shareholders may include firms and funds that have financial interests in entities covered in Gartner research. Gartner's Board of Directors may include senior managers of these firms or funds. Gartner research is produced independently by its research organization without input or influence from these firms, funds or their managers. For further information on the independence and integrity of Gartner research, see "Guiding Principles on Independence and Objectivity." (/technology/about/ombudsman/omb_guide2.jsp)"

About (<http://www.gartner.com/technology/about.jsp>)

Careers (<http://www.gartner.com/technology/careers/>)

Newsroom (<http://www.gartner.com/newsroom/>)

Policies (http://www.gartner.com/technology/about/policies/guidelines_ov.jsp)

Privacy (<https://www.gartner.com/privacy>)

Site Index (<http://www.gartner.com/technology/site-index.jsp>)

IT Glossary (<http://www.gartner.com/it-glossary/>)

Contact Gartner (http://www.gartner.com/technology/contact/contact_gartner.jsp)

