

Magic Quadrant for Cloud Infrastructure as a Service, Worldwide

Published: 03 August 2016 **ID:** G00278620

Analyst(s): Lydia Leong, Gregor Petri, Bob Gill, Mike Dorosh

Summary

The market for cloud IaaS has consolidated significantly around two leading service providers. The future of other service providers is increasingly uncertain and customers must carefully manage provider-related risks.

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 cloud storage providers, platform as a service (PaaS) providers, SaaS providers, cloud service brokerages (CSBs) or any other type of cloud service provider, 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.

Cloud IaaS includes not just the 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 PaaS capabilities. However, it does not include full stand-alone PaaS capabilities, such as application PaaS (aPaaS) and integration PaaS (iPaaS).

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 encompasses both single-application workloads and virtual data centers (VDCs) hosting many diverse workloads. It includes suitability for a wide range of application design patterns, including both cloud-native application architectures and enterprise application architectures.

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).

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 Magic Quadrants for managed hosting instead ("Magic Quadrant for Cloud-Enabled Managed 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. 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.

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.

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. 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 "availability sets"); 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 virtual machines (VMs), and some can offer shorter metering increments, which can be more cost-effective for short-term batch jobs. Providers charge on a per-VM basis, unless otherwise noted. Some providers offer either 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.

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.

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 who 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 2016 and April 2016.

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) 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.

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 API. 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, 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 (August 2016)

Vendor Strengths and Cautions

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]), 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). See the In-Depth Assessment for a detailed technical evaluation.

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 Germany, Ireland, Australia, India, Japan, Singapore, South Korea, Brazil and (in preview) China. It also has one region dedicated to the U.S. federal government. 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.

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 has a diverse customer base and the broadest range of use cases, including enterprise and mission-critical applications. It has the largest share of compute capacity in use by paying customers – many times the aggregate size of all other providers in the market. This has enabled it to attract an ecosystem of open-source tools, along with more than a thousand technology partners that have licensed and packaged their software to run on AWS, have integrated their software with AWS capabilities, or deliver add-on services. It also has a tiered and competency-badged network of partners that provide application development expertise, managed services and professional services such as data center migration. That ecosystem, along with AWS's training and certification programs, makes it easier to adopt and operate AWS in a best-practice fashion.

AWS is now a mature provider, yet it remains an agile, innovative thought leader with a broad impact across a range of IT markets. It has the richest array of IaaS and PaaS capabilities. It provides the deepest capabilities for governing a large number of users and resources. It

continues to rapidly expand its service offerings and to offer higher-level solutions. It retains a multiyear competitive advantage over all its competitors, and thus is the common reference point for competitive benchmarking. Although AWS will not be the ideal fit for every need, it has become the "safe choice" in this market, appealing to customers who desire the broadest range of capabilities and long-term market leadership.

CAUTIONS

While it is easy to get started with AWS, optimal use requires expertise. Customers must proactively seek the depth of sales and solution architecture engagement they need to be successful; this may require professional services, not just support. 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. AWS's support offerings are tiered based on the level of support that a customer purchases, rather than on a "relationship" or size-of-spend basis; customers need Business-tier support in order to ensure excellent support.

Organizations that cannot quickly take advantage of new capabilities will not receive the full benefits of AWS's rapid introduction of new services and expanded capabilities for existing services. A customer's best practices may become outdated as better or more cost-effective capabilities are introduced. Less-sophisticated customers may become overwhelmed by the range of possible options; training and third-party assistance are strongly recommended. Furthermore, although many new services are highly successful, services that turn out to be of less interest to customers may not get the same depth of continued investment as more popular services, complicating the question of which services to adopt. Finally, new services are gradually rolled out across regions, so customers outside the U.S. do not receive innovations as quickly, and global customers may not be able to obtain desired capabilities in all of the regions that they use.

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.

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. CenturyLink continues to sell legacy Savvis offerings, such as Cloud Data Center 2 and Cloud Servers, when those solutions are appropriate. The Marketplace Provider Program provides third-party software. Enterprise-grade support is extra. Managed services are optional. CenturyLink offers an aPaaS (AppFog), 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. The legacy Savvis cloud offerings are available in a broad range of data centers globally. CenturyLink has global sales, and business is conducted in local languages, but the service is offered only in English.

Recommended mode: CenturyLink primarily appeals 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

CenturyLink's Platform CenturyLink vision is rooted in the ability to deliver the breadth of CenturyLink's capabilities in an API-accessible and composable fashion. It has built a solid platform for increasing CenturyLink's own agility and ability to deliver new service offerings. CenturyLink is integrating its IaaS and PaaS offerings, and is building out broader service capabilities, including automation-augmented managed services. It is increasingly signaling the desire to be platform-neutral, including potentially managing other providers' cloud services or a customer's on-premises infrastructure, with services such as Runner, its Ansible-based hybrid infrastructure management solution.

While CenturyLink is increasingly focused on using the cloud as a means to enter the data center outsourcing market, it nevertheless has a competitive feature set for self-service, and successfully blends the self-service and managed services 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.

CAUTIONS

CenturyLink has a solidly capable and well-implemented basic offering, and has been executing successfully on its roadmap, but that roadmap is not sufficiently aggressive for the pace of the market. CenturyLink is potentially in an uncomfortable "in between" place in the market. On one side are market leaders that have broad portfolios of capabilities and managed service provider partners, and who are increasingly capable of attracting risk-averse customers that might have previously chosen a vendor like CenturyLink. On the other side are niche providers who specialize in specific applications and compliance requirements.

Cloud leadership turnover, as well as CenturyLink's exploration of strategic alternatives for its data center assets, create additional uncertainty around its future cloud ambitions. It is unclear if CenturyLink is willing to continue investing deeply in cloud IaaS and PaaS offerings beyond the core platform capabilities it needs to underpin Platform CenturyLink and its communications services. Customers should ensure that CLC 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: Fujitsu Cloud IaaS Trusted Public 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 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. Fujitsu has launched Cloud Service K5, a new OpenStack-based, KVM-virtualized offering with a variety of tenancy models. Managed services are optional.

Locations: S5 is available in data centers in the U.S. (West Coast), Germany, the U.K., Australia, Japan and Singapore. Fujitsu has global sales, and provides support in 34 languages; the S5 portal and documentation are available in English, German and Japanese. The regional offerings have their own capabilities and locations, which are different from those of S5. K5 is available only in Japan.

Recommended mode: Fujitsu appeals primarily to Mode 1 customers. In Japan, Fujitsu may also appeal to some customers with Mode 2 digital business initiatives.

Recommended uses: General business applications for customers who need managed services in conjunction with cloud IaaS. Development environments for customers who only need basic cloud IaaS.

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 is developing a portfolio of cloud IaaS, PaaS and SaaS services, and its vision is expanding to encompass a wider range of digital business capabilities. Its new 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 the foundation for Fujitsu's future IaaS and PaaS solutions. It is a new platform, and thus does not have an operational track record. Fujitsu intends to roll out the service globally during 2016. However, customers considering Fujitsu's existing offerings, or who are currently using S5, must factor this new offering into their decisions.

Fujitsu's cloud IaaS capabilities lag significantly behind those of the market leaders. K5 provides a better foundation for the future than S5, and Fujitsu has made acquisitions to accelerate its time to market, but it will continue to need to aggressively invest in acquiring and building technology in order to be competitive in this market. Although Fujitsu can sell its IaaS platform 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 into 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 and a Docker container service (Container Engine). 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). See the In-Depth Assessment for a detailed technical evaluation.

Locations: Google groups its IaaS data centers into regions, each of which contains at least two availability zones (data centers). There are East Coast, West Coast, and central U.S. regions, as well as a European region (located in Belgium) and an Asian region (located in Taiwan). Google has a global sales presence. Support is available in English and Japanese (during business hours). The portal is available in English, French, German, Spanish, Portuguese, Korean, Japanese and Mandarin. Documentation is available only in English.

Recommended mode: GCP appeals to Mode 2 buyers.

Recommended uses: Big data applications and batch computing, as well as cloud-native projects leveraging GCP as a whole.

STRENGTHS

Google's strategy for GCP centers on the concept of allowing other organizations to "run like Google" by exposing Google's internal technology capabilities as services that other companies can purchase. Google has extensive expertise in running technology platforms at scale, thanks to its consumer business. It has pioneered innovative infrastructure-related technologies, from its physical data center designs to its use of OS containers, and has been successful at advancing container-oriented capabilities related to Kubernetes, its open-source container cluster management software. It has a comprehensive vision for, and extensive experience with, how cloud-native applications are developed and managed through the life cycle, and GCP adoption has been driven primarily by cloud-native use cases.

Google is leveraging its expertise and experience with big data in its own consumer business in order to build product strength in technologies such as analytics and machine learning. These capabilities also have synergies with batch computing, where Google distinguishes itself with excellent price/performance value, exceptionally fast VM provisioning and per-minute billing. Evaluate GCP as if it were a specialized cloud platform for projects that play to these strengths. Many GCP customers have other primary cloud IaaS providers, but choose GCP for projects that are especially well-suited to the platform's capabilities.

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. It is still missing key capabilities that are important to established organizations (along with missing some capabilities that are important to startups), such as user management suitable for large organizations, granular and customizable role-based access control (RBAC), complex network topologies equivalent to those in enterprise data centers, and software licensing via a marketplace and license-portability agreements. Although it has been steadily releasing significant new capabilities, the general feature release velocity for GCP has not been as fast as necessary to become a leader in this highly competitive market.

Google is still in the rudimentary stages of learning to engage with enterprise and midmarket customers, especially those that are not technology-centric businesses. Over the course of 2015, Google made only incremental GCP-related progress in this area. Despite a change in leadership at the beginning of 2016, we believe that there is still insufficient forward movement in service features, sales, marketing, globalization and partner ecosystem to make Google broadly attractive as a strategic cloud IaaS provider in 2016.

IBM (SoftLayer)

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. The operations of SoftLayer, an IBM company, are being absorbed into IBM throughout the course of 2016.

Offerings: SoftLayer offers both multitenant and single-tenant Citrix-XenServer-virtualized compute (Virtual Servers). It also offers, as part of its cloud, 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). SoftLayer's primary business is noncloud offerings, such as paid-by-the-month dedicated servers (a broader range of configurations than is available per hour) and hosted appliances, but it does not make a clear distinction between these offerings and its cloud IaaS capabilities. Managed services are available through IBM. There is no support for colocation. IBM's aPaaS (Bluemix) is hosted in SoftLayer data centers, but the offerings are not directly integrated. See the In-Depth Assessment for a detailed technical evaluation.

Locations: SoftLayer has multiple data centers in the U.S., along with data centers in Canada, Mexico, Brazil, France, Germany, Italy, the U.K., Netherlands, Australia, Hong Kong, India, Japan and Singapore. It 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.

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 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 SoftLayer bare-metal servers as the hosting platform, where the customer has a need for supplemental basic cloud IaaS.

STRENGTHS

IBM acquired SoftLayer in order to have a data center platform that could enable IBM's cloud business, where its vision encompasses public and private IaaS, PaaS and SaaS. IBM also believes the shift to cloud will radically transform its business over time, including altering its go-to-market

strategy, with an increasing focus on online sign up and self-service, and SoftLayer represents an important foray into self-service and serving SMB customers. SoftLayer has a long track record as a dedicated hosting provider, and offers the broadest range of bare-metal cloud server configurations along with features that are specific to bare-metal servers.

IBM has a strong brand and existing customer relationships across the globe, and its base of strategic outsourcing customers will help drive a cloud-enabled data center outsourcing business in SoftLayer data centers. 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 14. However, in the near future, IBM will continue to be a hosting-scale provider, which may make it difficult for it to match the cost economics of the market leaders.

CAUTIONS

SoftLayer has limited differentiation beyond the hybrid blending of virtualized and bare-metal capabilities, and a broader geographic presence. 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. SoftLayer uses its own technology and API, which has limited third-party tool support. IBM's strategic direction has changed to integrated IaaS and PaaS, delivered via the Bluemix portal. IBM's new OpenStack-based IaaS options are currently in beta within Bluemix. Existing SoftLayer customers are being transitioned to IBM identities and a new SoftLayer UI within the Bluemix portal, and will have the option to migrate to IBM's new OpenStack-based offerings. We believe that SoftLayer is primarily becoming an enabling data center platform for the rest of IBM's business. It is a viable hosting option for IBM strategic outsourcing customers, but IBM customers looking for cloud IaaS should consider Bluemix's beta alternatives instead.

SoftLayer has historically been strongly focused on self-service dedicated hosting for SMB customers, and it is missing many cloud IaaS capabilities desired by midmarket and enterprise customers. Customers report that SoftLayer services still feel like a small-business experience, particularly with regard to the portal, sales and support. SoftLayer's portal does not distinguish between cloud services and the noncloud, by-the-month hosted hardware and appliances. Provisioning for all services can be initiated via API, but customers need to understand the technical and business implications of choosing the noncloud services, such as less financial flexibility, longer provisioning times and greater operational burdens placed on the customer.

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 launched Azure Infrastructure Services (which include Azure Virtual Machines and Azure Virtual Network) into general availability in April 2013, thus entering the cloud IaaS market.

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) and a CDN. 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). See the In-Depth Assessment for a detailed technical evaluation.

Locations: Microsoft calls Azure data center locations "regions." There are multiple Azure regions in the U.S., Canada, Australia, India and Japan, as well as regions in Ireland, the Netherlands, Hong Kong, Singapore and Brazil. There are also two regions for the U.S. federal government. (The two Azure China regions are part of a separate service operated by 21Vianet Group.) Microsoft has global sales. Azure support is provided in English, Dutch, French, German, Italian, Spanish, Japanese, Korean, Mandarin and Portuguese; support for Azure IaaS capabilities is also available in Hebrew. Technical documentation is available in those languages, as well as Russian. The portal is additionally available in Czech, Hungarian, Polish, Swedish and Turkish.

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, while Mode 2 customers tend to value Azure's ability to integrate with Microsoft's application development tools and technologies. Azure is frequently chosen for strategic adoption by organizations with a strong commitment to Microsoft technologies.

Recommended uses: General business applications and development environments 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 encompasses integrated IaaS and PaaS components that operate and feel like a unified whole. Microsoft has been rapidly rolling out new features and services, including differentiated capabilities. It has a vision of infrastructure and platform services that are not only leading stand-alone offerings, but also seamlessly extend and interoperate with on-premises Microsoft infrastructure (rooted in Hyper-V, Windows Server, Active Directory and System Center), development tools (including Visual Studio and Team Foundation Server [TFS]), middleware and applications, as well as Microsoft's SaaS offerings. Microsoft is also becoming more open and less reliant upon its Windows franchise, and Azure's support for Linux and other open-source technologies is improving quickly.

Microsoft's brand, existing customer relationships, history of running global-class consumer internet properties, deep investments in engineering and innovative roadmap have enabled it to rapidly attain the status of strategic cloud IaaS provider. Microsoft is aggressively pushing Azure into its customer base, and discounting to promote adoption. Azure is growing quickly, and is in second place for market share. Microsoft has pledged to maintain AWS-comparable basic cloud IaaS pricing for the general public; and, on a practical level, customers with Microsoft Enterprise Agreement discounts obtain a price/performance ratio that is comparable to AWS. Although Azure is neither as feature-rich nor mature as AWS, many organizations can now consider it "good enough," and base their vendor decision on factors other than technical capabilities.

CAUTIONS

While Microsoft has met its promised time frames for introducing critical features that help Azure fulfill enterprise needs for security, availability, performance, networking flexibility and user management, not all such functionality is implemented with the level of completeness, ease of use or API enablement that is desired by enterprise customers. These difficulties are exacerbated by disorganized, incomplete and sometimes out-of-date documentation, as well as a support organization that is not always capable of solving complex implementation challenges, a limited number of Azure experts outside of Microsoft (whether consultants or potential employees) and few options for Azure training.

Microsoft is still in the process of building out its Azure ecosystem. It has been aggressively recruiting managed service and professional services partners, but many of these partners lack extensive experience with the Azure platform, which can compromise the quality of the solutions they deliver to customers. Many such partners do not take advantage of cloud-native capabilities, reducing the value their customers receive from Azure. CMP vendors and MSPs report challenges in working with Azure, particularly in the areas of API reliability and secure authentication, which are slowing their ability to deliver solutions.

NTT Communications

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

Offerings: NTT Com has two cloud IaaS offerings in general availability. Cloud ⁿ is a fully multitenant, CloudStack-based, KVM-virtualized offering. Cloud ⁿ has an associated object storage offering, a CDN, a MySQL-based database as a service and a Cloud Foundry-based aPaaS. NTT Com Enterprise Cloud 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. NTT Com launched its Next-Gen Cloud – an OpenStack-based multihypervisor and bare-metal offering – in March 2016, after the technical-evaluation time frame for this Magic Quadrant.

Locations: Cloud ⁿ is available in multiple data centers in Japan, as well as a U.S. East Coast data center. NTT Com Enterprise Cloud 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. NTT Com has a global sales presence. Cloud ⁿ support is available in English and Japanese. NTT Enterprise Cloud support is available in English, French, German, Spanish, Cantonese, Hindi, Japanese, Mandarin, Malay and Thai. The portal and documentation for both offerings are available in English and Japanese.

Recommended mode: Cloud ⁿ will appeal primarily to Mode 2 customers. NTT Com Enterprise Cloud will appeal primarily to Mode 1 customers.

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

Recommended uses for NTT Com Enterprise Cloud: Development environments and general business applications for customers who need a Pan-Asian footprint and want a VMware-virtualized environment in conjunction with managed services.

STRENGTHS

NTT Com has a significant base of existing customers, especially in the Asia/Pacific region, to whom 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. 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, NTT Com has implemented software-defined networking in its data centers and in the interfaces between the offering and the WAN.

CAUTIONS

NTT Com Enterprise Cloud is a basic cloud IaaS offering with little differentiation. While Cloud ⁿ has more features that may be of interest to developers, it is missing capabilities that would make it attractive to enterprise IT operations organizations. There is no integration between the offerings. This makes it difficult for NTT Com to serve customers who need a unified offering, or who have cloud-native use cases that require enterprise-class capabilities.

NTT Com is in the midst of a strategic shift in cloud IaaS-related offerings. Its "Next-Generation Cloud Platform" is intended to serve as the unified cloud IaaS platform for its cloud efforts. It also will launch a Cloud Foundry-based PaaS. It intends to provide a CSB portal that includes not only its own offerings, but third-party clouds as well. These new services will be rolled out globally during 2016, and thus do not have an operational track record. However, customers considering or currently using NTT Com's existing offerings must factor these new options into their decisions.

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.

Recommended mode: Rackspace appeals to both Mode 1 and Mode 2 customers who 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 or Microsoft Azure.

STRENGTHS

Rackspace has successfully pivoted from its "Open Cloud Company," OpenStack-oriented strategy, and returned to its roots as "a company of experts," emphasizing its managed service expertise and superior support experience. Rackspace has a coherent vision of cloud-enabled managed services that utilize automation and a DevOps philosophy, including managing third-party clouds in addition to its own infrastructure. Rackspace offers self-managed cloud IaaS primarily in the context of its private cloud offerings.

Rackspace's 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. While Rackspace continues to participate in the OpenStack community, and has invested significantly in its OpenStack-based private cloud services, it is increasingly technology-neutral, with solid support for VMware and Microsoft private clouds as well. It is now focused on the ways in which it can add value beyond providing basic infrastructure, such as offering solution templates within its OpenStack-based private cloud.

CAUTIONS

Rackspace Public Cloud is a developer-centric offering, and has appealed primarily to small businesses seeking a replacement for low-cost mass-market hosting – cloud virtual private server (VPS) solutions – where it has new and aggressive competitors such as DigitalOcean. Few midmarket or enterprise customers want to use VPSs, though, and in the broader market, although Rackspace now delivers a solid set of basic features, it has not been able to keep up with the pace of innovation of the market leaders. Rackspace will continue to divert investment from its Public Cloud to other areas of its business, rather than try to compete directly for self-managed public cloud IaaS against market-leading providers that can rapidly deliver innovative capabilities at very low cost, or against established IT vendors that have much greater resources and global sales reach.

Rackspace may become an attractive acquisition target as its third-party managed cloud services and private cloud IaaS offerings grow in market success. System integrators and outsourcers may want to acquire their way into a stronger position in these two markets, but an acquisition of Rackspace could compromise its customer service culture. If Rackspace remains independent, it will face increasingly fierce competition from large technology and service providers across all of its cloud-related businesses.

Virtustream

Virtustream, an EMC Federation company, is focused solely on cloud services and software. Virtustream was founded in 2008.

Offerings: Virtustream's service, xStream, 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), including capabilities from Virtustream's 2014 acquisition of ViewTrust Technology. There are variants of xStream targeted specifically at customers who need PCI compliance, and U.S. federal government customers. Managed services are optional. Virtustream also offers xStream as software, and has licensed some of its technology to other service providers, such as IBM. The former EMC cloud storage business now also resides within Virtustream.

Locations: Virtustream has multiple data centers in the eastern and western U.S., Canada, the U.K., France, Germany and the Netherlands. It has sales presence in the U.S., along with London, Walldorf (the German home of SAP), Tokyo, and Dubai sales offices. The service is provided in English only.

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 production applications, especially mission-critical complex enterprise applications, such as ERP suites from SAP, Oracle and Microsoft. Its ability to template and automate the deployment and management of highly complex legacy applications enables customers to achieve agile application development life cycles and reduce the risk of application changes for complex workloads such as those associated with SAP. Virtustream has been successful at winning large-scale enterprise deals, particularly those focused on SAP and that require managed service capabilities. EMC acquired Virtustream in mid-2015, and this acquisition has been beneficial for Virtustream's sales presence, business expansion and engineering investments.

Virtustream has developed its own cloud platform technology, and can offer fully consistent hybrid cloud solutions. Its micro-VM technology enables it to charge for resources consumed, rather than resources allocated, and to offer policy-based service-level management and application

performance SLAs. It has focused on meeting enterprise security and compliance needs, and has some unique capabilities, such as support for Intel's Trusted Execution Technology (TXT) and trust framework. Virtustream typically works closely with customers to determine how best to migrate on-premises applications to its cloud, and it has its own data collection tools that customers run on their on-premises servers to determine what resources they will need on xStream.

CAUTIONS

Virtustream provides deep and differentiated capabilities in its focus areas, rather than broad general-purpose capabilities. It tries to attract ERP systems as well as the constellation of applications that often surround such systems, rather than competing for all workloads. Although the Dell-EMC merger creates some uncertainties about its future direction, customers should expect Virtustream to continue to focus on its core strengths, rather than expanding into the broad cloud IaaS market. Customers should treat Virtustream as a specialized provider for the workloads that suit the strengths and weaknesses of its technology platform.

Although Virtustream supports a solid set of self-service features, 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.

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 an EMC Federation company.

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 is not accepting new customers and will be 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 in English and Japanese.

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 VMware offers 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

VMware has narrowed its focus for vCloud Air to hybrid solutions for VMware's existing customer base. vCloud Air is now merely one of several ways that VMware intends to address the growing shift toward the use of cloud services, such as emphasizing the role of service provider partners that use VMware technology, and delivering management capabilities for both VMware-based and non-VMware-based cloud infrastructure. While VMware continues to invest in the service, and the service continues to grow, VMware is no longer significantly expanding the geographic footprint of vCloud Air, nor investing in the engineering necessary to expand its feature set beyond basic cloud IaaS. VMware has redirected engineering investment into improving vCloud Air's infrastructure capabilities, including incorporating new VMware software features.

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 IT operations 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, while VMware does have a strong channel in its broader business, the vCloud Air channel is not robust; SIs, MSPs and value-added resellers need to learn to sell vCloud Air effectively and deliver value on top.

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

No new vendors were added this year.

Dropped

The market is consolidating, and the increased dominance of just two providers has led to closer relative market share among the other providers in the market. In particular, there is much closer clustering of workload share below the top 10. Therefore, we reduced the number of vendors from 15 to 10, and dropped the following vendors for market share reasons:

CSC. While CSC continues to grow its own cloud IaaS business, it is increasingly focused on CSB services, leveraging its acquisition of ServiceMesh to manage third-party cloud IaaS environments, particularly AWS and Microsoft Azure.

Dimension Data. Dimension Data continues to grow and operate its cloud IaaS businesses, which include both public and private cloud services. It is the only significant global cloud IaaS provider with a presence in South Africa.

Interoute. Interoute continues to invest strongly in its cloud IaaS business, and has demonstrated good growth, especially in its core European markets. It remains a solid choice for customers who need a Pan-European cloud IaaS footprint, or that value the software-defined WAN capabilities that are integrated into Interoute's cloud IaaS offering.

Joyent. Although Joyent continues to grow its cloud IaaS business, it is increasingly focused on also selling Triton, the software that it uses to run its public cloud. Its illumos-based (Solaris-derived) platform, with an integrated Docker service, has unique container-oriented strengths. On 15 June 2016, Samsung announced that it would acquire Joyent.

Verizon. Verizon is paring back its public cloud IaaS efforts and reshaping its cloud IaaS strategy. The former Terremark Enterprise Cloud platform now has a new Verizon Cloud portal experience and is being actively sold under the Verizon Cloud brand. Verizon continues to enhance and evolve that platform, and introduce new cloud capabilities based on that platform. Verizon has retired the second-generation public cloud IaaS platform that became generally available at the end of 2014.

See the 2015 Magic Quadrant and Critical Capabilities for a more detailed examination of these vendors.

Inclusion and Exclusion Criteria

To be included in this 2016 Magic Quadrant, vendors had to demonstrate the following, as of January 2016:

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 10 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

Vendors Considered but Not Included

There are vendors who come up periodically on client inquiries but are not included in this Magic Quadrant. These vendors, and the reasons they are not included, are as follows:

Aliyun. Aliyun is the cloud computing division of Alibaba. In 2015, Aliyun expanded to data centers outside of China, but the service is currently delivered only in Mandarin.

Cisco. In 2014, Cisco acquired Metacloud, an OpenStack-based private cloud IaaS provider. Cisco now offers this service under the Cisco Metapod brand. Cisco is not included in this Magic Quadrant because the inclusion criteria require the provider to offer public cloud IaaS (not just private cloud IaaS), and furthermore, Cisco does not have enough market share to qualify.

Digital Ocean. Digital Ocean is one of many cloud IaaS providers that specialize in serving small businesses that typically use just one or two VMs; other examples include ElasticHosts, Linode and 1&1. These low-cost providers are often used as an alternative to mass-market virtual private server (VPS) hosting. Gartner clients typically have needs beyond what such "cloud VPS" providers can fulfill. Digital Ocean's capabilities do not fulfill Gartner's formal inclusion criteria for this Magic Quadrant.

IBM. While only IBM's SoftLayer offering is included in this Magic Quadrant, IBM also has other IaaS initiatives that were not evaluated in this research. IBM's Bluemix offering – formerly purely a PaaS – has been expanded to include IaaS elements; at the time of this evaluation, however, most of these IaaS elements were not in general availability, and Bluemix as a whole is not a generally available cloud IaaS offering. IBM also offers the Blue Box service, an OpenStack-based private cloud IaaS offering. IBM's go-forward integrated IaaS and PaaS platform will be presented through the Bluemix portal, and will incorporate capabilities from Bluemix, SoftLayer, and Blue Box.

Oracle. At the time of this evaluation, the Oracle Compute Cloud Service was not in general availability, and Oracle did not have another cloud IaaS offering in general availability that meets Gartner's definition of IaaS. Consequently, Oracle also does not have enough market share to qualify for inclusion.

In the evaluations for this Magic Quadrant, we also considered a variety of interesting cloud providers that did not meet the criteria for inclusion. Some of these providers offer solutions that utilize cloud IaaS, or have non-IaaS offerings that might be considered as alternatives. Examples, by use case, include:

Development: Skytap, which has deep self-service "lab management" capabilities and features that support collaboration between developers, using either its own infrastructure or AWS.

Disaster Recovery: Bluelock, Hosting (formerly Hosting.com) and Iland, which offer disaster recovery as a service on their cloud IaaS platforms.

Enterprise applications: NaviSite, whose cloud IaaS platform is suitable for hosting Oracle E-Business Suite and other complex applications.

HPC: CloudSigma, Internap, Cogeco Peer 1 and ProfitBricks, which offer configurations and pricing models that are attractive for HPC.

Even though some businesses may use PaaS 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 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."

Evaluation Criteria

Ability to Execute

Gartner analysts evaluate technology vendors on the quality and efficacy of the processes, systems, methods or procedures that enable IT providers' performance to be competitive, efficient and effective, and to positively affect revenue, retention and reputation. Ultimately, technology providers are judged on their ability to capitalize on their vision, and on their success in doing so.

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 who 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 who 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 who 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 who 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 who 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 who 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 who 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 who want a reliable, predictable service experience.

Our evaluation of a service provider's Ability to Execute remains similar to that of the 2015 Magic Quadrant, with increased expectations across all criteria, as the market consolidates amid explosive growth in customer demand.

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 2016)

Completeness of Vision

Gartner analysts evaluate technology vendors on their ability to articulate logical statements convincingly about current and future market direction, innovation, customer needs and competitive forces, as well as how they map to Gartner's position. Ultimately, technology providers are assessed on their understanding of the ways in which market forces can be exploited to create opportunities.

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 who 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 who 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 who 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, automated ITOM and overall feature set. This criterion is important to buyers who want a provider who 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 who 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.

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 who 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 who want to use a global vendor.

Our evaluation of Completeness of Vision remains similar to that of the 2015 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

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

Source: Gartner (June 2016)

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, 2015." 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 who 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 and PaaS (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.

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 "CEOs and CIOs Must Assume That Every Industry Will Be Digitally Remastered.") Digital business use cases are very broad, and include digital marketing, e-commerce, e-CRM, SaaS and data services. 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?

This phase of the market has been "won." The market consolidated dramatically over the course of 2015. Just two providers – AWS and Microsoft Azure – account 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 who deliver bundled managed services.

The next phase of the market has not yet emerged. It is likely that the next phase of this market will even more tightly integrate IaaS and PaaS capabilities, including an expanded use of container technologies and automated operations management. However, we expect that these changes will be gradually incorporated into existing service offerings; there will not be a sudden shift in the market. We expect the overall competitive environment will not change significantly until 2018 at earliest, and new entrants to the market will have minimal impact before that time. That means that customers can make strategic bets on providers now, and expect that their relative market positions are not likely to change dramatically over the next three years.

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, who have successful cloud VPS offerings; many such providers serve local markets or a single country.

Buyers should be concerned about the long-term viability of non-market-leading offerings. Many service providers are reconsidering their cloud IaaS strategies, and are trying to build a new platform, acquire another provider in order to use their platform instead, specialize in a niche, pivot to become third-party managed service providers on other clouds, or are even deciding to exit the market entirely. Buying from a vendor with a major brand offers little or no protection, and indeed, major IT vendors usually need to experience greater success in order to justify continued investment. For example, within the past year, AT&T and Hewlett Packard Enterprise shut down their public cloud IaaS offerings, Fujitsu and NTT Communications launched new platforms, IBM launched an IaaS offering under its Bluemix brand, and Verizon retired several offerings and withdrew significantly from the public cloud IaaS market.

Buyers should treat new cloud IaaS platforms with caution. Multiple major vendors are launching new cloud IaaS offerings. Some of these vendors are not new to the cloud IaaS market but believe they need a different platform to compete successfully; others are new entrants to the market. Regardless of the size of the vendor, buyers should be cautious of these new offerings. While such offerings may be marketed as being competitive against the market leaders, closer examination will show significant technical deficiencies. It takes considerable time to build the breadth and depth of a feature set, and to learn to operate and scale such offerings. Beware of vendor promises of future features that do not come with firm commitments as to when those features will become generally available.

Vendor relationships are important, but technical capabilities matter. Some customers who 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 who 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 cost-management tools such as RightScale Cloud Analytics and Cloudability, to continuous configuration automation tools such as HashiCorp's Terraform.

Customers are decoupling the choice of cloud IaaS offering from managed services 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.

The software-defined data center is the center of 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. Non-market-leading providers will be particularly challenged to build a strong ecosystem of managed and professional services providers, management tool vendors, and ISVs who package and support their software on the platform.

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.

?

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. U.S. federal government entities are supposed to use FedRAMP-compliant solutions, but, in practice, they do not always do so. However, the Government Accountability Office (GAO) and agency inspector generals (IGs) are increasingly monitoring whether agencies utilize FedRAMP-compliant vendors or not (see "Market Trends: U.S. Federal Cloud Market – Beyond the Hype" and the GAO's September 2015 report, "Agencies Need to Correct Weaknesses and Fully Implement Security Programs" (<http://www.gao.gov/assets/680/672801.pdf>)).

FedRAMP Compliance

Cloud IaaS providers exist in all three forms of 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). All forms of compliance are considered valid for federal IaaS. It is still unclear whether there is any advantage to follow one path versus the other, since the

point of the program is to "do once and reuse many times" by leveraging reciprocity. In practice, some providers have reported meeting with some 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 IaaS providers is limited. Our discussions with vendors that have completed the process suggest an average of 18 months and \$3.5 million to go through the process. This has led to increasing dissatisfaction with the FedRAMP Program Management Office, 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 in the process of certification.

Some providers may not want to undertake the effort and expense of an ATO, but may have solutions that meet the FedRAMP requirements – but these solutions are not considered FedRAMP-compliant and will come under increasing scrutiny. As FedRAMP ATOs find their way into an increasing number of RFP requirements, failure to comply will increasingly become a barrier.

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 HIPAA (Health Insurance Portability and Accountability Act) for health-related data and CJIS (Criminal Justice Information Services) for public-safety-related data.

The current FedRAMP framework aligns to the Federal Information Security Management Act (FISMA) Low and Moderate levels. A FISMA High baseline, to handle more sensitive information, is currently in the pilot phase with four providers.

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 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, Amazon Web Services (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 usually physically and logically segregated from the commercial solutions. These federal solutions are normally 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.

State and local government entities are usually not permitted to use federal IaaS solutions, although this varies by provider – it is the provider that decides which customers are permitted in the community cloud. Such entities may need to find commercial cloud IaaS solutions instead. Increasingly, state and local 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. 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.

Amazon Web Services

AWS's GovCloud (US) 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 and IRS 1075 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.

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. There are two regions, one in Northern Virginia and one in Iowa. It has a FedRAMP High JAB P-ATO. It has Level 2 authorization for the DoD SRG. Microsoft has also announced two regions dedicated to the DoD to support DoD SRG Level 5 workloads, with Level 4 and 5 authorizations expected in 2016. It can be used for workloads that must adhere to CJIS and IRS 1075 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.

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; it is not part of Verizon Cloud. It is deployed in Northern Virginia, and has an agency ATO from the Department of Health and Human Services.

Virtustream

Virtustream's Federal Cloud Services (VFC) is a community cloud for U.S. federal government customers. It is similar to Virtustream's Enterprise Cloud, with Northern Virginia- and San Francisco-based isolated deployments of Virtustream's xStream platform. It has a FedRAMP JAB P-ATO and adheres to ITAR regulations. VFC can be used for workloads that must adhere to CJIS and IRS 1075 requirements. In addition, Virtustream's Viewtrust solution can be used for continuous monitoring and on-demand compliance reporting.

VMware

VMware vCloud Government Service (vCGS), provided by Carpathia, is a partnership between VMware and Carpathia to provide a government community cloud. There are two vCGS data centers, one in Northern Virginia and one in Arizona. vCGS has a FedRAMP JAB P-ATO, and can be used for workloads that must adhere to CJIS requirements. Note that this service uses the same architecture as VMware's own vCloud Air and must meet the same quality metrics, but it is still a distinct service operated by a VMware partner, not VMware itself. Carpathia has a long history as a managed hosting provider for government customers; it was acquired by QTS Realty Trust in 2015.

Evidence

Gartner client inquiries in 2015 and 2016 (currently more than 1,000 cloud IaaS-related inquiries per quarter)

Service provider interviews and product demonstrations in 2015 and 2016

Surveys of more than 75 cloud IaaS providers in 2015 and 2016

Surveys of more than 6,700 end-user organizations in 2015, regarding their use of IaaS and their preferred providers

Customer references from the service providers in 2015 and 2016

Hands-on trials of service offerings in 2015 and 2016

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 "SOC Attestation Might Be Assurance of Security ... or It Might Not."

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.

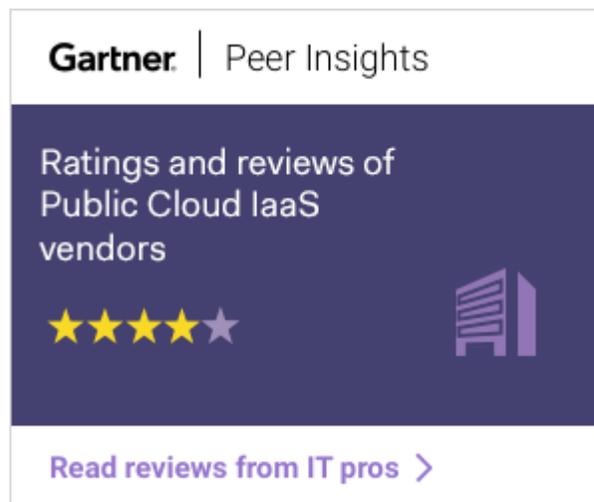
Offering (Product) Strategy: The vendor's approach to product development and delivery that emphasizes differentiation, functionality, methodology and feature sets as they map to current and future requirements.

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.



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

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

© 2016 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 (<http://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)

