

Magic Quadrant for Cloud Infrastructure as a Service

28 May 2014 ID:G00261698

Analyst(s): Lydia Leong, Douglas Toombs, Bob Gill, Gregor Petri, Tiny Haynes

VIEW SUMMARY

The market for cloud compute infrastructure as a service (a virtual data center of compute, storage and network resources delivered as a service) is still maturing and rapidly evolving. Strategic providers must therefore be chosen carefully.

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, software as a service (SaaS) providers, cloud services brokerages 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.

We draw a distinction between cloud infrastructure *as a service*, and cloud infrastructure *as a technology platform*; 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. (See "Technology Overview for Cloud-Enabled System Infrastructure," "Technology Overview for Cloud-Enabled Managed Hosting" and "Don't Be Fooled by Offerings Falsely Masquerading as Cloud Infrastructure as a Service" for more on this distinction.)

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.

This Magic Quadrant primarily evaluates cloud IaaS providers in the context of the fastest-growing need among Gartner clients: the desire to have a "data center in the cloud," where the customer retains most of the IT operations responsibility. 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 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 it may be supplemented by a small amount of colocation or dedicated servers. 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 Managed Hosting, North America," "Magic Quadrant for European Managed Hosting" and "Magic Quadrant for Cloud-Enabled Managed Hosting, Asia/Pacific"). Organizations that do not want self-service, but instead want managed services with an underlying CESI, should consult our 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 only solutions that are delivered in an entirely standardized fashion — specifically, public cloud IaaS, along with private cloud IaaS that uses the same or a highly similar platform. Although most of the providers in this Magic Quadrant do offer custom private cloud IaaS, we have not considered these 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

Learn how
Gartner can
help you succeed

Become a Client now ▶

EVIDENCE

Over 2,000 Gartner client inquiries in 2013 and 2014
Service provider interviews and product demonstrations in 2013 and 2014
Surveys of more than 75 cloud IaaS providers in 2013 and 2014
Customer references from the service providers in 2013 and 2014
Hands-on trials of service offerings in 2013 and 2014
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.

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 of 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. All the providers also offer custom private clouds, unless otherwise noted.

Most of the providers are oriented toward the needs of traditional IT operations, with an emphasis on control, governance and security, and the ability to run both new applications and legacy workloads. The providers that are oriented toward the needs of developers are noted as such; these providers typically emphasize easy access to infrastructure for individuals who are building new applications. Some developer-oriented offerings target small or midsize businesses (SMBs) and startups, and lack the features needed by enterprises and midmarket entities.

Most of the providers have resilient infrastructure, achieved through redundant infrastructure in conjunction with virtual machine (VM) clustering, or the ability to rapidly detect VM failure and immediately restart it on different hardware. They are thus able to offer very high SLAs for infrastructure availability — sometimes as high as 99.999% (sometimes expressed as a 100% SLA with a 10-minute exclusion). Offerings without VM clustering or fast VM restart — which provide higher levels of infrastructure availability than can be expected from a single physical server — are noted as lacking autorestart.

Most of the providers have maintenance windows that result in downtime of the control plane (including the GUI and API), and may require infrastructure downtime. Some offerings can utilize live migration of VMs, largely eliminating the need for downtime to perform host maintenance, but this does not eliminate maintenance windows in general.

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. Storage performance varies considerably between providers. Some providers can offer tiered storage with a solid-state drive (SSD) option.

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

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

All providers offer, at minimum, per-hour metering of VMs, and some can offer shorter metering increments, which can be more cost-effective for short-term batch jobs. Providers charge on a per-VM basis unless otherwise noted. Some providers 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 allow customers to choose arbitrary-size VMs — any combination of virtual CPUs (vCPUs), RAM and VM storage, subject to some limits. Providers that do not allow this are explicitly noted as offering fixed-size VMs. Some providers with fixed-size VMs have a very limited range of VM sizes, while others have a wide variety of sizes and suit a broad range of use cases. Some providers who offer arbitrary-size VMs may enforce a maximum ratio between vCPUs and RAM.

Most of the providers can resize an existing VM without needing to reprovision it; those that cannot are explicitly noted as offering nonresizable VMs. Some of the providers can resize an existing VM without needing to reboot it.

Most of the providers can provision a basic Linux VM within 5 minutes (although this will increase with large OS images, and Windows VMs typically take somewhat longer). Those that cannot are noted as having slow provisioning. Most providers can do simultaneous provisioning of multiple VMs; for example, provisioning 20 VMs will finish about as quickly as one VM. Those that cannot are noted as such, and the degradation can be significant (some providers take over an hour to provision 20 VMs). Nonsimultaneous provisioning has a major negative impact in disaster recovery, instant high-scalability and batch-computing scenarios.

Some of the providers are able to offer an option for single-tenant VMs within a public cloud IaaS offering, on a fully dynamic basis, where a customer can choose to place a VM on a host that is temporarily physically dedicated to just that customer, without the customer needing to buy a VM that is so large that it consumes the whole physical host. These VMs are typically more expensive than VMs on shared hosts. Providers that have this option are noted as such.

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

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.

significant number of midmarket customers plan to move into colocation and then gradually migrate into that provider's IaaS offering.

Typically, the storage associated with an individual VM is persistent. However, some providers have ephemeral storage, where the storage exists only during the life of the VM; if the VM goes away unexpectedly (for instance, due to hardware failure), all data in that storage is lost. Ephemeral storage is always noted explicitly.

All the providers offer VM-independent block storage unless otherwise noted. A few providers allow storage volumes to be mounted on multiple VMs simultaneously, although customers must correctly architect their solutions to ensure data integrity (just as they would with a traditional storage array).

All the providers offer object-based cloud storage, unless otherwise noted. In many cases, this service is integrated with a content delivery network (CDN).

All the providers have a private WAN that connects their data centers, unless otherwise noted. They offer an option for private network connectivity (usually in the form of Multiprotocol Label Switching [MPLS] or Ethernet purchased from the customer's choice of carrier), between their cloud environment and the customer's premises, unless otherwise noted. Providers for which we state "third-party connectivity is via partner exchanges" are ones where private connectivity is obtained via cross-connect in the data centers of select partners, such as Equinix; this also meets the needs of customers who require colocated equipment. Some carriers may also have special products for connecting to specific cloud providers — for example, AT&T NetBond and Verizon Secure Cloud Interconnect.

Most of the providers support the use of Internet-based IPsec VPN (virtual private networking). All the providers allow customers to have VMs with only private Internet Protocol (IP) addresses (no public Internet connectivity), and also allow customers to use their own IP address ranges, unless otherwise noted. Some providers may enforce secure access to management consoles, restricting access to VPNs or private connectivity.

All the providers claim to have high security standards. The extent of the security controls provided to customers varies significantly, though. All providers offer multifactor authentication (MFA), unless otherwise noted. Most providers offer additional security services. 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), and some may have ISO 27001 audits for their cloud IaaS offering (see Note 2); audits should not be taken as indications of security.

Most providers offer a firewall (intrusion detection system/intrusion prevention system) as part of their offering, although a few offer only access control lists (ACLs) and a few offer no self-service network security at all; this will always be explicitly noted. All providers provide distributed denial of service (DDoS) attack mitigation, unless otherwise noted.

All the providers offer customers a self-service ability to create complex network topologies with multiple network segments and multiple virtual network interface cards (NICs), unless otherwise noted.

All the providers allow customers to bring their own VM images, unless otherwise noted. This allows a customer to create snapshots of existing VMs within their own internal data center, and then directly import them via self-service into the provider's cloud, rather than having to start from the provider's own VM image library. This also allows the import of VM appliances and other prepackaged VM images from independent software vendors (ISVs). All providers allow storage snapshots and have the ability to turn the snapshot into a VM image, unless otherwise noted. All the providers have the ability to do bulk import and export of data on physical media, unless otherwise noted.

Some providers allow customers to create provisioning templates that group multiple resources, including compute, storage and network elements, and allow them to be provisioned as a unit.

Some of the providers also have post-provisioning hooks, allowing customers to run scripts after VM provisioning is complete but before the VM is available for login.

All the providers offer self-service monitoring as an option, unless otherwise noted. A few offer trigger-based autoscaling, which allows provisioning-related actions to be taken based on a monitored event. Some providers offer schedule-based autoscaling, which allows provisioning-related actions to be executed at a particular time.

All the providers offer self-service, front-end load balancing, unless otherwise noted. All also offer back-end load balancing (used to distribute load across the middle and back-end tiers of an application), unless otherwise noted.

All the providers offer a portal and self-service mechanism that is designed for multiple users and that offers hierarchical administration and role-based access control (RBAC). However, the degree of RBAC granularity varies greatly. From most to least control, RBAC can be per element, tag, group or account. Unless stated otherwise, a provider's RBAC applies across the account. Providers typically predefine some roles; the ability to have customer-defined roles offers more control, and is noted where available. We strongly recommend that customers that need these features, but that want to use a provider that does not have strong support for them, evaluate a third-party management tool, such as Dell Cloud Manager (formerly Enstratius), RightScale or ServiceMesh (acquired by CSC).

All providers log events such as resource provisioning and deprovisioning, VM start and stop, and account changes, and allow customers self-service access to those logs for at least 60 days, unless otherwise noted.

All providers, unless otherwise noted, offer the ability to place metadata tags on provisioned resources, and to run reports based on them, which is useful for internal showback or chargeback. Some providers also offer cost control measures such as quotas (limits on what a user can provision) and leases (time-limited provisioning of resources).

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.

All the providers evaluated are believed to be financially stable, with business plans that are adequately funded. Customers should not need to worry about them going out of business. However, many of the smaller providers are likely to be potential acquisition targets; an acquisition can cause significant changes in the strategy and direction of a business, and may result in a service

transition period if the merged companies consolidate their platforms. Furthermore, small IaaS businesses within large vendors may be subject to turmoil if the vendor changes its overall cloud strategy.

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. (See "Infrastructure as a Service in the Cloud Services Value Chain" for details.)

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 services 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.

Format of the Vendor Descriptions

When describing each provider, we first briefly summarize the nature of the company and then provide information about its public cloud IaaS offerings (and any single-tenant offerings that are otherwise identical), in the following format:

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

Compute, storage, network and security notes: Notes on the offering, including any missing core functionality or significant features, compared with the standard functionality discussed above.

Other notes: We list other capabilities of note, including important missing capabilities. We specifically note other cloud-related services, such as cloud storage (which all providers have, unless otherwise noted), as well as the availability of managed services, even though those service offerings are not specifically evaluated in the context of this Magic Quadrant, because they are capabilities frequently requested by customers in conjunction with cloud IaaS. (See "Market Insight: Customers Need Hybrid Cloud Compute Infrastructure as a Service" for details.)

In the compute notes, 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, neither the use of the same underlying CMP nor API compatibility indicates 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 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).

Services that use VMware's virtualization technologies are labeled as follows:

vCloud Datacenter Service. This service has been certified to meet VMware's globally consistent service definitions, security and regulatory compliance requirements, and requirements for availability and high performance. It is based on a prescriptive architecture intended to maximize portability between providers of vCloud Datacenter Service and a business's own VMware-virtualized data center infrastructure. Only 10 providers worldwide have such a service and most of them do not yet have a significant customer base on this platform. These providers also meet the requirements for being vCloud Powered.

vCloud Powered. These providers are part of VMware's service provider partner program. The service is based on VMware's vSphere and vCloud Director (vCD), exposes the vCloud API, and supports the Open Virtualization Format (OVF) for image upload and download. Unless otherwise stated, these providers expose the vCD UI to customers. Because the vCD features exposed can be customized by the service provider, and the service provider typically needs to provide an array of features not included in vCD (such as monitoring), there is still significant differentiation between vCloud Powered providers. In a vCloud Powered offering with the vCD UI exposed, vCD is used to drive self-service management and provide a service catalog. vCD is a key part of VMware's strategy for driving adoption of hybrid internal-external cloud IaaS, and facilitates interoperability between VMware-virtualized infrastructures, regardless of whether they are internal to a business or offered by a service provider. vCD provides the capability to manage very complex infrastructure needs, but also requires a greater investment in training and setup time from an IT administrator in order to facilitate easier self-service for users.

vCloud Express. vCloud Express is a VMware-defined offering aimed at developers and small businesses, with online sign-up, credit card payment, self-service and by-the-hour service.

VMware-virtualized. This service uses VMware's hypervisor, but is not a vCloud Datacenter, vCloud Powered or vCloud Express service. Many such offerings are high-quality services from early, market-leading innovators; these providers typically entered the market before vCD became available and have elected to continue to develop their own technology.

We summarize all of the provider descriptions, including a comparison of their capabilities against our baseline expectation of capabilities, in tabular format in "Toolkit: Comparison Matrix for Cloud Infrastructure-as-a-Service Providers, 2014."

We provide a detailed list of evaluation criteria in "Evaluation Criteria for Cloud Infrastructure as a Service." Our "Critical Capabilities for Public Cloud Infrastructure as a Service" provides a use-case-focused technical evaluation of the public cloud IaaS offerings of the included providers.

Recommended Uses

For each vendor, we also 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 transaction processing (TP) principles.

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 TP 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.

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



Vendor Strengths and Cautions

Amazon Web Services

AWS, a subsidiary of Amazon.com, is a cloud-focused service provider with a very pure vision of highly automated, cost-effective IT capabilities, delivered in a flexible, on-demand manner.

Locations: AWS has groups of data centers, which it calls "regions," on the East and West Coasts of the

U.S., and in Ireland, Japan, Singapore, Australia, Brazil, and (in preview) China. It also has one region dedicated to the U.S. federal government. It has global sales. Support is provided in English, Japanese and Portuguese. Technical account managers can also provide support in German, Spanish, Hindi, Korean and Mandarin. The portal and documentation are available in English, Dutch, French, German, Portuguese, Spanish, Japanese, Korean and Mandarin.

Compute: Elastic Compute Cloud (EC2) offers multitenant, fixed-size and nonresizable, Xen-virtualized VMs without autorestart. Single-tenant VMs are available via Dedicated Instances. There are special options for HPC, including graphics processing units (GPUs). AWS does not have any formal private cloud offerings, though it is willing to negotiate such deals (such as its deal for the U.S. intelligence community cloud).

Storage: VM storage is ephemeral. Persistence requires VM-independent block storage (Elastic Block Store). There is an option for SSDs, as well as storage performance guarantees (Provisioned IOPS). Object-based storage (Simple Storage Service [S3]) is integrated with a CDN (CloudFront), there is an option for long-term archive storage (Glacier), and AWS offers its own cloud storage gateway appliance.

Network: AWS offers a full range of networking options. Complex networking and IPsec VPN is done via Amazon Virtual Private Cloud (VPC). Third-party connectivity is via partner exchanges (AWS Direct Connect).

Security: RBAC is per-element, with customer-defined roles and exceptional control over permissions. AWS has obtained many security and compliance-related certifications and audits.

Other notes: Enterprise-grade support is extra. The SLA is multi-fault-domain, but does not have any exclusion for maintenance; AWS also offers continuous availability on its portal and API. Notable capabilities include orchestration (CloudFormation and OpsWorks), autoscaling, database as a service (Relational Database Service [RDS]), Hadoop as a service (Elastic MapReduce), data warehousing as a service (Redshift), and desktop as a service (WorkSpaces). AWS does not offer colocation; a partner exchange must be used instead. We provide purchasing guidance in "What Managers Need to Know About Amazon Web Services" and a detailed technical evaluation in "Amazon Web Services (AWS): In-Depth Assessment."

Recommended uses: All use cases that run well in a virtualized environment, although highly secure applications, strictly compliant or enterprise applications (especially complex ones 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 is the overwhelming market share leader, with more than five times the cloud IaaS compute capacity in use than the aggregate total of the other 14 providers in this Magic Quadrant. It is a thought leader; it is extraordinarily innovative, exceptionally agile, and very responsive to the market. It has the richest array of IaaS features and PaaS-like capabilities, and continues to rapidly expand its service offerings. It is the provider most commonly chosen for strategic adoption.

AWS has a very large technology partner ecosystem. Many software vendors have specially licensed and packaged their software to run on EC2, either independently or via the AWS Marketplace, which eases deployment and eliminates some of the challenges associated with licensing software to run in the cloud. Its API is supported by many third parties that provide associated management tools, and many open-source and commercial CMPs are compatible with its API. Although AWS is not directly involved in the hybrid on-premises cloud business, it has partners that can offer such solutions.

AWS has an extensive network of partners that can offer assistance with adopting its platform. It is increasingly targeting mainstream enterprises and the midmarket via go-to-market partnerships with SIs (such as Capgemini, Cognizant and Wipro) that provide application development expertise, managed services and professional services such as data center migration, although such providers do not necessarily deliver solutions that are optimal for the cloud environment. There are also many consulting and managed services partners that offer a cloud-native approach, such as Datapipe and 2nd Watch.

AWS has multiple "availability zones" (AZs) within its regions. These AZs are effectively multiple data centers in close proximity to one another. AWS's services are designed to make it easier to run applications across multiple AZs; customers are responsible for architecting their applications for high availability. However, new capabilities are rolled out incrementally, region by region, so newer capabilities are not necessarily available in every region.

Cautions

AWS is a price leader, but it charges separately for optional items that are sometimes bundled with competitive offerings. This increases the complexity of understanding and auditing bills. Prospective customers should be careful to model the costs accurately, especially for network-related charges, and to compare the costs of reserved and unreserved capacity, as well as AWS's "spot pricing" market.

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; the quality of support differs materially between tiers. AWS does not include enterprise-grade support by default; customers will need to buy Business tier support for this. Its Enterprise tier support offers a dedicated technical account manager and other "platinum" capabilities, providing a higher degree of support than most of its competitors offer without managed services, but it carries up to a 10% premium on the customer's overall AWS spend.

AWS has field sales, solutions engineering and professional services organizations, but the rapid growth of AWS's business means that sales capacity is insufficient to consistently satisfy prospective customers who need consultative sales. For better terms and conditions, customers should sign an Enterprise Agreement, which is typically a zero-dollar contract. Invoicing is available on request.

AWS is beginning to face significant competition — from Microsoft in the traditional business market, and from Google in the cloud-native market. So far, it has responded aggressively to price drops by competitors on commodity resources. However, although it is continuously reducing its prices, it does not commodity-price services where it has superior capabilities. AWS currently has a multiyear competitive advantage, but is no longer the only fast-moving, innovative, global-class provider in the

market.

CenturyLink

CenturyLink acquired Savvis, a Web hoster with a long track record of leadership in the hosting market, 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. It has a broad suite of offerings, including colocation, managed hosting, data center outsourcing, and PaaS (via its acquisition of AppFog). Its public cloud IaaS offering, CenturyLink Cloud, uses the newly acquired Tier 3 platform. The former Savvis public and private cloud IaaS offerings, which were under the Symphony brand, and are now known as Cloud Data Center and Cloud Servers, are no longer being actively marketed to new customers, although CenturyLink will still provide these solutions if appropriate.

Locations: CenturyLink Cloud (CLC) is available in multiple data centers across the U.S., along with Canada, Germany and the U.K. CenturyLink's private 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.

Compute: CLC is a multitenant, paid-by-the-VM, VMware-virtualized offering. The legacy Cloud Data Center 2 service (CDC2, the successor to Savvis Symphony VPDC) is a multitenant, paid-by-the-VM, VMware-virtualized offering. The Cloud Servers service (originally offered under the Symphony Open and Dedicated brands) is VMware-virtualized, with varying pricing models, and does not have an API. All offerings lack simultaneous provisioning.

Storage: CLC has persistent local storage with an SSD option, as well as VM-independent block storage that is integrated with rolling backups and disaster recovery options. The other offerings only have local storage.

Network: There is no back-end load balancing. None of the offerings support the ability to use customer-supplied IP addresses, nor do they fully support self-service complex network topologies. Most of the CLC data centers are not yet privately connected via the CenturyLink network.

Security: CLC's RBAC is per element. CDC2's RBAC is per group. CLC does not support MFA. CLC can support PCI requirements, but other compliance requirements require noncloud elements or managed services.

Other notes: CLC's Blueprints can be used to script the provisioning of complex, multi-data-center infrastructure configurations. CLC supports leases. Managed services for the guest OS layer are optional in CLC and CDC2. The full range of managed services are available in CenturyLink's private cloud IaaS offerings. CenturyLink also offers database as a service.

Recommended uses: General business applications, enterprise applications, development environments, cloud-native applications, e-business hosting.

Strengths

CenturyLink has a compelling and distinctive vision of application-fluent infrastructure that spans network, compute and storage capabilities. With the acquisition of Tier 3, it has obtained an IaaS offering that combines an excellent, differentiated set of features on a well-engineered platform with an easy-to-use self-service portal. It is one of the few services with both cloud-native capabilities that are attractive to developers and the governance and management features needed by large enterprises.

CenturyLink has a well-established track record of delivering enterprise cloud services for production and mission-critical needs, including addressing the needs of customers with extensive security and compliance requirements. However, many of CenturyLink's capabilities are delivered as managed services, rather than true cloud services.

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. The existing CenturyLink base of managed hosting, colocation and network customers provides it with cross-selling opportunities.

Cautions

While CenturyLink has an ambitious vision and road map, it will be challenged to muster the resources and agility of the leading providers in this market. CenturyLink will need to focus on midmarket and enterprise customers who need tailored solutions and managed services.

CenturyLink is consolidating all of its IaaS platforms. Previously, it had multiple types of single-tenant and multitenant IaaS. This means that existing customers will eventually have to migrate to the new CLC offering; the older offerings are still being supported but no longer being improved. However, CLC is currently only multitenant, so CenturyLink continues to sell its previous single-tenant offerings, and in some cases, may also sell its other legacy offerings. Current and prospective customers who do not intend to use CLC should consider what they will do over the long term.

CenturyLink is still learning to deliver solutions on the CLC offering, which creates challenges for sales, solution engineering, installation and delivery. Customers who are not simply using self-service on the standard multitenant CLC offering should take extra care to understand how CenturyLink intends to architect and deliver their solution. Furthermore, over the past two years, Savvis customers have reported customer service challenges.

CSC

CSC is a large, traditional IT outsourcer with a broad range of data center outsourcing capabilities.

Locations: CSC has multiple cloud data centers in the U.S., as well as in Canada, Brazil, Germany, Luxembourg, the Netherlands, Switzerland, the U.K., Australia, Malaysia and Singapore. It has global sales. Support is provided in English, French, German, Italian, Spanish and Mandarin. The portal and documentation are available only in English.

Compute: CSC offers a vCloud Datacenter Service, a VCE Vblock-based cloud IaaS architecture in different tenancy models — public multitenant in a CSC data center (CloudCompute), single-tenant compute with a multitenant back-end (BizCloud Virtual Private Edition [VPE]), and private single-tenant

in a CSC data center or in the customer's own data center (BizCloud) — and optional managed services. It offers both paid-by-the-VM and SRP pricing, but requires annual contracts. Because features are introduced into BizCloud before being rolled into CloudCompute, the latter contains a subset of BizCloud features; furthermore, it lacks autorestart. While customers can access vCloud Director if they prefer, CSC has built its own, more user-friendly portal.

Storage: VM-independent block storage can be SSD-accelerated. CSC does not offer object-based cloud storage.

Network: CSC has the full range of networking options.

Security: RBAC is per element. There is no DDoS mitigation.

Other notes: Managed services are optional. CSC also has significant additional software capabilities supporting IT operations management, along with trigger-based and schedule-based autoscaling, and quotas and leases for resource management. It also offers database as a service (CloudDB) and Hadoop as a service (via its acquisition of Infochimps).

Recommended uses: Cloud-enabled data center transformation for all workloads that run well in a virtualized environment, excluding batch computing.

Strengths

Unlike most other traditional data center outsourcers, CSC has fully embraced the highly standardized, highly automated cloud model, successfully blending the benefits of a true cloud service into an enterprise-ready offering. It has a solid platform that is attractive to traditional IT operations organizations that still want to retain control, but need to offer greater agility to the business and are willing to embrace data center transformation.

CSC has a strong road map focused on bringing enterprise-class IT operations management (ITOM) tools, including automated managed services, to cloud IaaS. It is trying to integrate not just traditional ITOM tools, but also DevOps tools; for example, it offers CA LISA Release Automation (formerly Nolio) as part of its platform. Its November 2013 acquisition of ServiceMesh has brought significant additional capabilities, including multicloud management. It is also building infrastructure utility services for specific applications on top of the platform.

CSC is one of the few providers to have a standardized architecture across both public and private cloud offerings, as well as a single rate card across all these offerings — although the pricing is the same, the minimum commitments vary. CSC's pricing for infrastructure resources is very competitive.

CSC has developed a portfolio of cloud-related professional services, including Smart Start, a proof-of-concept program intended to help a customer achieve a "quick win" in moving an application onto IaaS, and then methodically migrating other workloads over time. In general, CSC is generous about offering trials to prospective customers.

Cautions

Since acquiring ServiceMesh, CSC has reduced its focus on developing its own cloud IaaS offering, in favor of becoming a cloud services brokerage that can capably manage multiple clouds, including AWS and Microsoft Azure. CSC now leads sales with ServiceMesh's Agility Platform capabilities, and then helps the customer choose the appropriate IaaS offering, which might or might not be CSC's own. This may impact CSC's willingness to invest in its own cloud IaaS platform.

Cloud computing is driving a radical reinvention of the way in which CSC delivers services, including significantly broadening the range of companies that CSC targets with its offerings. The cloud division is run as its own business unit, which gives it greater agility but also sometimes brings it into conflict with its slower-moving and more conservative parent company. At present, however, CSC is investing heavily in its cloud business.

Prospective CSC customers should be careful to understand the distinction between CSC's outsourcing business and its cloud offerings; in particular, they should be sure to understand what is and isn't within the scope of a particular tier of managed services. Existing CSC outsourcing customers report challenges in getting CSC to engage in cloud opportunities.

CSC's focus has been on enabling traditional IT operations organizations to make the transition to cloud infrastructure, and this remains true as it expands to managing multiple cloud IaaS offerings. CSC is trying to increase its traction with developer audiences, but will find this challenging, given the positioning of the broader CSC brand.

Dimension Data

Dimension Data, an NTT Group company, is a large SI and value-added reseller. It entered the cloud IaaS market by acquiring OpSource in 2011.

Locations: Dimension Data has data centers on the East and West Coasts of the U.S., plus the U.K., Netherlands, Australia, Hong Kong, Japan, Brazil and South Africa. Local-language sales and support is provided in 51 countries, with cloud-specialized support provided from its regional service centers. The portal is available only in English. Documentation is available in English and Japanese.

Compute: Dimension Data offers paid-by-the-VM public and private cloud IaaS. The former is VMware-virtualized; the latter can be Hyper-V or VMware-virtualized. Provisioning is nonsimultaneous.

Storage: VM-independent block storage has tiered performance, including SSD-accelerated tiers.

Network: Dimension Data has a full range of networking options. All traffic between its data centers is WAN-optimized.

Security: RBAC permissions are across the whole account, but customers can define their own roles.

Other notes: Monitoring is not self-service. There are two tiers of optional managed services, which include improved SLAs.

Recommended uses: E-business hosting, cloud-native applications and general business applications.

Strengths

Dimension Data's Managed Cloud Platform (MCP) is a single unified architecture across its public and private cloud offerings; it is one of the few providers to provide such an architecture. It is pursuing a federated model, whereby service provider partners offer MCP-based services via resale or white label, but Dimension Data provides one consistent, unified service globally. It also provides sales and marketing enablement to its OneCloud Alliance members.

OpSource had a long history as a SaaS hoster, and Dimension Data has retained these capabilities. Its rich suite of offerings for that market includes not only infrastructure, but also an on-demand billing platform, custom application management and help desk support. However, its previously excellent SLAs have been significantly weakened in the past year.

Dimension Data has launched Cloud Software, a set of partnerships with ISVs. It offers Dimension Data-tested and -licensed software from those ISVs, on demand, priced by the hour. However, there is little software available in this model.

Cautions

While Dimension Data's offering is VMware-virtualized, it is not vCloud Powered. Instead, Dimension Data is doing extensive software development of its own, in hopes of driving a faster pace of innovation and controlling its costs better. Although it has consistently rolled out incremental improvements in a continuous delivery model, it has not been able to keep up with the pace of innovation in the overall market. It is now focusing primarily on deals that include managed services.

Although Dimension Data has a capable basic cloud IaaS offering, it needs value-added capabilities, whether delivered through its own services or via an ecosystem of partners. It has tried to differentiate based on its network model, but such capabilities are now "table stakes."

Dimension Data is owned by NTT Group. While NTT has deliberately chosen Dimension Data to be its most agile business, with minimal interference from the parent, Dimension Data's future ability to move quickly is likely to depend on continued support and noninterference.

Fujitsu

Fujitsu is a large diversified technology company. It has a range of cloud IaaS offerings, including Fujitsu Cloud IaaS Trusted Public S5 (formerly the Fujitsu Global Cloud Platform), multiple regional offerings based on a global reference architecture (Fujitsu Cloud IaaS Private Hosted, formerly known as Fujitsu Local Cloud Platform), and multiple private cloud offerings. Although Fujitsu has received vCloud Datacenter Service Provider partner status, it has not yet launched this offering.

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.

Compute: S5 is a multitenant, fixed-size and nonresizable, paid-by-the-VM, Xen-virtualized offering; it is also available in a single-tenant version (S5 Dedicated). There is no autorestart. Provisioning is nonsimultaneous.

Storage: Block storage is persistent and VM-independent. Although S5 has storage snapshots, snapshots cannot be used as VM images. Customers cannot import their own VM images.

Network: Although S5 supports private connectivity and private-IP-only VMs, it cannot use customer-provided IP addresses. It does not fully support complex network topologies.

Security: RBAC is per element. There is no DDoS mitigation. Fujitsu's cloud IaaS offerings cannot meet common compliance requirements.

Other notes: There is no metadata tagging of assets, but user quotas are supported. Managed services are optional.

Recommended uses: General business applications, and test and development.

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 a vendor with a broad product portfolio of hardware and software, and it is developing most of its own technology for its cloud offerings, across IaaS, PaaS and SaaS. It is using its existing technology, including its ITOM software, to accelerate the pace of its cloud business. Fujitsu's ServerView Resource Orchestrator Cloud Edition software is used to provide the visual designer used in its IaaS user interfaces, as well as significant depth of portal features.

Fujitsu's previous strategy of allowing its regions to pursue their own cloud strategies has enabled certain regions, such as Australia, to develop offerings tailored to the needs of their local markets, at a faster pace than Fujitsu has been able to do so as a global entity. Fujitsu has begun an "offering standardization process" intended to bring the regional offerings in line with the global ones, but it may be a few years before this is accomplished.

Fujitsu rolls out S5 features first in Japan, then extends them to its other regions. Furthermore, Fujitsu in Japan offers additional cloud capabilities — Japan-based organizations or projects targeted at the Japanese market should investigate what capabilities are specifically available in Japan, such as object-based storage, database as a service and Hadoop as a service.

Cautions

Fujitsu's cloud IaaS capabilities lag significantly behind those of the market leaders, and the gap is increasing, not narrowing, although Fujitsu is actively investing in new capabilities and has been highly responsive to customer requests. Customers are likely to be able to meet a broader range of needs by combining multiple Fujitsu offerings. The multitenant S5 offering is significantly more expensive than competing offerings.

Fujitsu's previous strategy of allowing regional control means that development efforts have been

fragmented across the globe. As such, the private hosted service offerings may differ in each region, making it difficult for Fujitsu to capitalize fully on engineering resources and achieve economies of scale, although Fujitsu has recently strengthened its global cloud strategy and management.

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. Fujitsu is not focused on self-managed cloud IaaS.

GoGrid

GoGrid is a small, independent cloud-IaaS-focused provider.

Locations: GoGrid's data centers are on the East and West Coasts of the U.S., and in the Netherlands. It has local sales in its San Francisco and Amsterdam offices. The service is provided only in English.

Compute: GoGrid offers fixed-size, paid-by-the-VM, Xen-virtualized IaaS in both multitenant and single-tenant variants. RAM is resizable, but CPU is not.

Storage: Customers cannot import their own VM images. Object-based cloud storage is integrated with a CDN (via a partnership with EdgeCast).

Network: Customers cannot bring their own IP addresses. There is no back-end load balancing.

Security: There is no MFA. Changes to network perimeter security configurations are not logged.

Other notes: RBAC permissions are whole-account. There is no metadata tagging of assets. Managed services are optional.

Recommended uses: Cloud-native applications and e-business hosting, with an emphasis on big data.

Strengths

GoGrid has repositioned itself as a provider of "open data services." Its "1-Button Deploy" enables quick provisioning of complex deployments of Hadoop and NoSQL technologies. GoGrid is engaging in active partnerships with relevant software vendors, building validated architectures with partners such as Cloudera, DataStax, Hortonworks and MongoDB. It is focusing its infrastructure options on those suitable for running big data workloads, including VMs with large amounts of RAM and high I/O performance.

GoGrid has broad and aggressive SLAs that include 100% availability and cover compute, storage and networking elements. It is one of the few providers that has a standard architecture across its public and private cloud offerings.

The GoGrid Exchange enables software vendors to license and package their software to run on GoGrid, easing deployment and eliminating some of the headaches associated with licensing software to run in the cloud. It has a unique multipartner compensation model that enables partners to build on top of each other's software stacks.

Cautions

GoGrid's feature set is limited, compared with the current state of the market. GoGrid's software is developed entirely in-house, which has created significant long-term challenges in terms of its ability to compete against providers that can devote significant resources to research and development, but its shift to focus on big data solutions should allow it to target engineering investments where they will be most effective.

GoGrid has its own API, which is supported by a limited number of third-party tools. GoGrid is pursuing a strategy of broader interoperability, and its future success will depend on ensuring that the company can partake in one or more of the emerging platform ecosystems.

Google

Google is an Internet-centric provider of technology and services. Google Cloud Platform combines an IaaS offering (Google Compute Engine [GCE]), an application PaaS offering (Google App Engine) and a range of complementary services. Google has been operating App Engine since 2008, but did not enter the IaaS market until the general-availability launch of GCE in December 2013.

Locations: Google groups its GCE data centers into "regions," each of which contains at least two availability zones. There is a central U.S. region, a European region (located in Belgium), and an Asia region (located in Taiwan). It has global sales. Support is available in English and Japanese. The portal is available in English, Spanish, Portuguese, Japanese and Mandarin. Documentation is available in those languages, plus French, German, Russian, Cantonese and Korean.

Compute: GCE offers multitenant, fixed-size and nonresizable, KVM-virtualized VMs, metered by the minute. Provisioning is exceptionally fast (typically under 1 minute).

Storage: VM storage is persistent, and there is also VM-independent block storage. All block storage is encrypted.

Network: Third-party private connectivity is not supported. Customers cannot bring their own private IP addresses (although this need may possibly be addressed by GCE's Advanced Routing features). There is no back-end load balancing.

Security: RBAC permissions apply to the whole account.

Other notes: Enterprise-grade support is extra. The SLA is multi-fault-domain, and excludes maintenance in Europe (where Transparent Maintenance is not available yet). Google cannot address the need for colocated equipment. Autoscaling capabilities (Replica Pools) are in preview (beta).

Recommended uses: Cloud-native applications and batch computing, as well as projects leveraging Google Cloud Platform as a whole.

Strengths

Google's strategy for Google Cloud Platform centers on the concept of allowing other organizations to "run like Google" by taking Google's highly innovative internal technology capabilities and exposing

them as services that other companies can purchase. Consequently, although Google is a late entrant to the IaaS market, it is primarily productizing existing capabilities, rather than having to engineer those capabilities from scratch. It will therefore be able to advance its offering more rapidly than most competitors.

Google has a comprehensive vision for, and extensive experience with, how cloud-native applications are developed and managed through the life cycle. It has a fluid notion of the boundaries between IaaS and PaaS that will, over time, enable customers to choose their trade-offs between control and automated management.

Google, as an overall corporate entity, has a very broad global data center footprint, an enormous amount of existing infrastructure capacity, and its own high-capacity global network. It is extremely efficient at operations. GCE represents little incremental cost to Google, which means it can price aggressively and emphasize the value of high performance for the money. Over time, though, Google will differentiate itself with platform and manageability features, not prices.

Cautions

GCE is brand-new and consequently lacks an operational track record, though it has suffered several small operational glitches since it entered general availability. Its feature set is nascent. However, Google's road map is innovative and aggressive, and it is rapidly introducing new capabilities. Its short-term focus is on better enabling new cloud-native applications, with less attention being paid to capabilities needed for other workloads. It has no hybrid cloud strategy.

Although Google has significant appeal to technology-centric businesses, it is still learning to engage with enterprise and midmarket customers, and needs to expand its sales, solutions engineering and support capabilities. Furthermore, Google needs to earn the trust of businesses. Google also lacks many capabilities important to businesses that want to migrate legacy workloads to the cloud.

Google needs to build an ecosystem around GCE (and, more broadly, Google Cloud Platform), and its partner program is nascent. However, Google is perceived as a likely future market leader, and cloud management tool vendors and other software companies, along with system integrators, are rapidly adding support for GCE.

HP

HP is a large diversified technology company with a range of cloud-related products and services. Its only true cloud IaaS offering is the HP Public Cloud, although it has some cloud-enabled infrastructure services, such as the HP Helion Managed Virtual Private Cloud (formerly HP Enterprise Cloud Services). It also has a Cloud Foundry-based PaaS in private beta.

Location: HP's Public Cloud data centers are in the eastern and western U.S. Although it has global sales, the service is offered in English only.

Compute: HP Public Cloud is a multitenant, fixed-size and nonresizable, KVM-virtualized, OpenStack-based offering.

Storage: VM storage is ephemeral. There is an option for persistent, VM-independent block storage. There is object-based storage with an integrated CDN (via a partnership with Akamai).

Network: The load-balancing service is currently in private beta.

Security: RBAC permissions are group-based. Common compliance requirements cannot be supported. Audit logs for noncompute elements are available only via a support request. There is no MFA.

Other notes: The SLA is multi-fault-domain, but does not have any exclusion for maintenance. The monitoring service is in private beta.

Recommended uses: Development environments for existing HP customers, or those who specifically want to explore an OpenStack-based cloud offering.

Strengths

HP has an ambitious and comprehensive vision of interoperable public and private cloud infrastructure, integrating HP's hardware capabilities, ITOM tools, outsourcing capabilities and partner ecosystem.

HP has a large global sales force and marketing budget, and has been willing to offer extraordinary discounts to prospective customers in order to bring them onto its service. However, HP Public Cloud has data centers only in the U.S., which limits its attractiveness to global customers, and such discounts may not be sustainable.

Cautions

Although HP has invested significant engineering resources into its Public Cloud over a multiyear period, the service, which entered general availability in December 2012, still has a nascent feature set, and important capabilities remain in a lengthy beta period. HP is focused on hybrid cloud, including private cloud enablement and cloud-enabled managed services, and its cloud IaaS business is simply a component of that overall strategy, not a business that HP is aggressively pursuing on a stand-alone basis.

Although HP's public cloud offering is based on OpenStack, as is HP Cloud OS (the platform for HP Converged Cloud and part of the HP CloudSystem CMP), there is limited interoperability between the Public Cloud and private clouds based on HP CloudSystem. HP Public Cloud's inability to export VM images on a self-service basis means that customers cannot directly transfer workloads from this service to their private cloud, even if their private cloud is based on OpenStack.

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 SMBs, and in January 2014, shut down its own SmartCloud Enterprise cloud IaaS offering, after migrating its existing customers to SoftLayer. IBM's related non-IaaS portfolio includes IBM Cloud Managed Services (cloud-enabled managed services, formerly called SmartCloud Enterprise+), as well as Codename: BlueMix (a Cloud Foundry-based PaaS, currently in beta).

Locations: SoftLayer has multiple data centers in the U.S., along with data centers in the Netherlands and Singapore. It has global sales. It offers support in English, Dutch, German, and Spanish. The portal and documentation are available in English only.

Compute: SoftLayer's cloud IaaS offering was formerly known as CloudLayer, but SoftLayer has stopped distinguishing this offering from its dedicated hosting options. SoftLayer's IaaS is paid-by-the-VM, Citrix-Xen-virtualized, and available as a single-tenant or multitenant offering. There is an option for single-tenant VMs within SoftLayer's multitenant public cloud. There is also an option for "bare metal" (nonvirtualized) servers, by the hour.

Storage: Storage is persistent and VM-independent block storage is available. There is OpenStack object-based storage with an integrated CDN (via a partnership with Internap).

Network: SoftLayer supports a full range of networking options.

Security: SoftLayer supports only grouping, not full metadata tagging, and RBAC permissions are per-group.

Other notes: SoftLayer has an extensive dedicated server business, although these bare-metal custom configurations require monthly commitments. Managed services are available through IBM. There is no support for colocation.

Recommended uses: E-business hosting, general business applications, large-scale use cases such as gaming where bare metal is desirable, and infrastructure for IBM outsourcing deals.

Strengths

IBM's cloud vision encompasses public and private clouds at both the system and application infrastructure levels. IBM's portfolio of products and services has IaaS, PaaS and SaaS elements, along with hardware, ITOM software, middleware and database software, and comprehensive consulting and outsourcing capabilities. SoftLayer will be the common hosting platform for IBM's services in the future.

IBM has a strong vision for how 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. SoftLayer is one of its first forays into self-service and serving SMB customers. 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 on top of SoftLayer.

IBM intends to make local presence one of its competitive differentiators, taking advantage of SoftLayer's relatively small-scale "pod" architecture to expand the service from three countries to 15, over the course of 2014. 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 hyperscale market leaders.

SoftLayer has a particularly clean service composition, with a range of options that can be added on a per-instance, paid-by-the-hour basis, including aspects such as the type of monitoring and the automated response to a failure detected by monitoring. Instances can be VMs or dedicated servers, and SoftLayer emphasizes that bare-metal physical servers can be managed with the same elasticity as VMs, and that single-tenant and multitenant infrastructure can be seamlessly mixed within the same solution.

Cautions

Although IBM has set out a broad vision for cloud services, it lacks integrated IaaS and PaaS capabilities; IBM's PaaS (Codename: BlueMix) is simply hosted on SoftLayer. It has not articulated a vision for how SoftLayer itself will be differentiated, beyond the availability of bare-metal servers and future in-country presence. We believe it is more likely that customers who pursue strategic adoption of IBM cloud services will do so for the managed services and other higher-level capabilities, than for the core SoftLayer capabilities.

SoftLayer has historically been strongly focused on self-service for SMB customers. The acquisition by IBM has brought enterprise sales capabilities, including channel capabilities, to the SoftLayer business, as well as the higher-level IBM managed services. However, customers report that SoftLayer services still feel like a small-business experience, not an enterprise experience, particularly with regard to the portal, sales and support. We recommend that customers engage with SoftLayer through an IBM sales channel, rather than through SoftLayer's own direct sales mechanisms.

IBM is "cloudwashing" SoftLayer's dedicated server business, applying a "cloud IaaS" label to what is actually simply dedicated hosting. While SoftLayer has significant automation, as well as a customer-accessible API, for this portion of its business, customers should be sure to understand the technical and business differences between SoftLayer's dedicated hosting (customized bare-metal servers by the month) and cloud IaaS (fully automated, by-the-hour compute in the form of VMs or bare metal).

Although IBM has an OpenStack-based strategy for its private cloud business, SoftLayer's compute infrastructure is not OpenStack-based. Customers who want OpenStack or other CMPs on SoftLayer must obtain them as custom, hosted, managed services. SoftLayer has a proprietary API that has not yet gained widespread third-party tool support, although IBM is working to build the necessary ecosystem, and has also begun a project called Jumpgate, to provide an OpenStack-to-SoftLayer API translation gateway.

Joyent

Joyent is a small, independent service provider that focuses solely on cloud services and software.

Joyent did not respond to requests for supplemental information or to review the draft contents of this document. Gartner's analysis is therefore based on Joyent's previous-year responses, public information, use of Joyent's service, and discussions with Joyent's existing and prospective customers.

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

Compute: Joyent offers fixed-size, paid-by-the-VM public cloud IaaS (SmartMachines), and private cloud services in a variety of pricing models. The host OS is Joyent's own SmartOS, an open-source derivative of OpenSolaris, managed by the illumos community. Customers have a choice between OS virtualization in a SmartOS Container and KVM virtualization on a SmartOS Container for Linux and Windows guests.

Storage: VM storage is persistent, but there is no VM-independent block storage. There is an SSD option. Joyent's S3-compatible object-based cloud storage (Manta) has a unique architecture designed for batch jobs that require high-performance access to large amounts of storage, with an in-place batch compute service separate from Joyent's main compute service.

Network: Joyent has a full range of networking options.

Security: Joyent has a single-account model, although an account can have multiple API keys. Its RBAC is whole-account.

Other notes: Enterprise-grade support is extra. Joyent's Content Delivery Cloud, which utilizes Riverbed Stingray Traffic Manager, offers functionality akin to a private CDN.

Recommended uses: Cloud-native applications and e-commerce sites where visibility into application performance is crucial; batch computing on large datasets.

Strengths

Joyent has a unique vision for cloud IaaS and is exceptionally innovative from a technology perspective. It is developing an integrated technology stack and its infrastructure offerings verge on the platform space. It is making deep investments in fundamental technologies, including its own SmartOS operating system, based on illumos Solaris and its x86 hypervisor. Joyent's cloud uses SmartOS Containers (originally known as Solaris Zones). KVM runs natively within a container, thus providing additional security, resource control and resource visibility within the virtualization layer.

Joyent has a pure focus on new, cloud-native applications, including mobile applications, as well as, with its newly-launched Manta service, big data applications and massively parallel data analytics. Joyent is the sponsor of Node.js. It offers commercial support for Node.js, as well as proprietary tools focused on Node.js operations within its platform.

Joyent places strong emphasis on application performance and takes a holistic approach to its delivery, including integrating network-based acceleration. It has particularly deep portal-based performance analytics, which use the DTrace framework for application instrumentation. It has an excellent cost-performance ratio.

Joyent intends to derive its future revenue by offering cloud services directly and selling its SmartDataCenter CMP software (including via OEM partners such as Dell), which is fully compatible with its own public cloud IaaS offering. It has launched the Global Cloud Network, a cloud federation alliance focused on mobile carriers; its strategic carrier partners include Telefonica and Bharti Airtel.

Cautions

Joyent focuses on developing its own technology, and has a track record of releasing innovative capabilities. However, it faces a long-term challenge to compete against providers with greater development resources. That said, it has chosen to focus its efforts on particular areas, leaving certain capabilities, notably enterprise management features, to partners such as Dell (via its acquisition of Enstratius).

Joyent's feature set is strongly oriented toward cloud-native use cases, and it is highly developer-centric. It emphasizes API capabilities and the enablement of third-party tools, rather than portal capabilities of its own.

The unique nature of Joyent's offering makes it crucial for the company to develop an ecosystem around its platform; it must attract ISVs and third-party tool vendors, along with MSPs and SIs that can provide managed and professional services. Joyent has begun to build this ecosystem, focused on cloud-native vendors, but it is not yet mature.

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 previously 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.

Locations: Azure Infrastructure Services are available in multiple data centers in the U.S., as well as in Ireland, the Netherlands, Hong Kong, Japan, Singapore, China and (in preview) Brazil. Microsoft has global sales, and Azure support is provided during local business hours in English, French, German, Italian, Spanish, Japanese, Korean, Mandarin and Portuguese; 24/7 support is provided in English and Japanese only. The portal and documentation are available in those languages, as well as Russian.

Compute: Azure VMs are fixed-size, paid-by-the-VM, and Hyper-V-virtualized; they are metered by the minute.

Storage: Block storage ("virtual hard disk") is persistent and VM-independent. Object-based cloud storage is integrated with a CDN.

Network: There is no support for complex network topologies. Third-party connectivity is via partner exchange (Azure ExpressRoute).

Security: Virtual network topology limitations prevent useful deployment of most security-related virtual appliances, such as a perimeter intrusion detection/prevention system (IDS/IPS). RBAC uses Azure Active Directory, but permissions are whole-account.

Other notes: Enterprise-grade support costs extra. The SLA is multi-fault-domain, but does not have any exclusion for maintenance. Additional significant capabilities include orchestration (Azure Automation), scheduling, autoscaling and Hadoop as a service (HDInsight). Microsoft does not offer colocation; a partner exchange must be used instead. We provide a detailed technical evaluation in "Microsoft Windows Azure Infrastructure Services: In-Depth Assessment."

Recommended uses: General business applications and development environments for Microsoft-

centric organizations; cloud-native applications; use as part of an overall Microsoft Azure solution.

Strengths

Microsoft has a vision of infrastructure and platform services that are not only leading stand-alone offerings, but that also seamlessly extend and interoperate with on-premises Microsoft infrastructure (rooted in Hyper-V, Windows Server, Active Directory and System Center) and applications, as well as Microsoft's SaaS offerings. Its vision is global, and it is aggressively expanding into multiple international markets.

Microsoft's brand, existing customer relationships, history of running global-class consumer Internet properties, deep investments in engineering, and aggressive road map have enabled it rapidly to attain the status of strategic cloud IaaS provider. It is second in terms of cloud IaaS market share — albeit a distant second — but far ahead of its smaller competitors. Microsoft has pledged to maintain AWS-comparable pricing for the general public, and Microsoft customers who sign a contract can receive their enterprise discount on the service, making it highly cost-competitive. Microsoft is also extending special pricing to Microsoft Developer Network (MSDN) subscribers.

The broader Microsoft Azure service is a full-featured PaaS offering with significant complementary capabilities; the Virtual Machines are integrated into the overall offering. The IaaS and PaaS components within Microsoft Azure feel and operate like part of a unified whole, and Microsoft is making an effort to integrate them with Visual Studio, Team Foundation Server, Active Directory, System Center and PowerShell. Conversely, Windows Azure Pack offers an Azure-like user experience for on-premises infrastructure. Microsoft has built an attractive, modern, easy-to-use UI that will appeal to Windows administrators and developers. The integration with existing Microsoft tools is particularly attractive to customers who want hybrid cloud solutions.

Cautions

Azure Infrastructure Services are relatively new. Although Microsoft rolled out an array of new features throughout 2013, and continues rapidly to release both "catch-up" capabilities and fresh innovations, many enterprises will still find significant capability gaps, especially in networking and security. Many features are in "preview" (beta) or "coming soon," and it is not always obvious to customers which features are still in preview. Customers who intend to adopt Azure strategically and migrate applications over a period of two years or more (finishing in 2016 or later) can begin to deploy some workloads now, but those with a broad range of immediate enterprise needs are likely to encounter challenges.

Microsoft is in the midst of a multiyear initiative to make its on-premises software "cloud first," rather than trying to scale software originally built for on-premises single-enterprise use. It now faces the challenges of getting its core infrastructure technology to operate at cloud scale, managing that infrastructure at cloud scale, and facilitating the ability of customers to move toward more highly automated infrastructure.

Microsoft has just begun to build an ecosystem of partners around Azure Infrastructure Services, and it does not yet have a software licensing marketplace. Furthermore, it has little in the way of enterprise Linux options. Consequently, the offering is currently very Microsoft-centric and appeals primarily to .NET developers, although customers do run heterogeneous environments in Azure.

Rackspace

Rackspace is an independent Web hoster with a long track record of leadership in the managed hosting market. It is one of the founders of OpenStack. Its Rackspace Private Cloud business provides traditional commercial open-source support and professional services around it. In addition to its public cloud IaaS offerings, Rackspace can support custom OpenStack-based private clouds in its own data centers or in customer data centers. It also owns numerous related businesses; some, such as SaaS email, are part of Rackspace itself, while others, such as Jungle Disk, are subsidiaries.

Locations: Rackspace Public Cloud is located 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 in the U.S., along with the U.K., the Netherlands, Switzerland, Hong Kong and Australia. Support is provided in English only. The portal and documentation are available in English and Spanish.

Compute: Rackspace Public Cloud is a multitenant, fixed-size, Citrix Xen-virtualized, OpenStack-based public cloud IaaS offering. It lacks autorestart.

Storage: VM storage is persistent, but there is also optional persistent VM-independent block storage. There is an option for SSDs. There is object-based storage (Cloud Files) with an integrated CDN (via a partnership with Akamai).

Network: There is no self-service network security. Private connectivity requires use of the RackConnect service.

Security: RBAC permissions are whole-account. There is no MFA. There are no audit logs. There is no support for most common compliance requirements without the use of hybrid hosting.

Other notes: In-depth support requires managed services, which are optional. Additional capabilities include autoscaling and database-as-a-service offerings (MySQL, MongoDB and Redis). We provide a detailed technical evaluation in "Rackspace Public Cloud: In-Depth Assessment."

Recommended uses: 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; development environments where simplicity and ease of use are crucial.

Strengths

Rackspace has a coherent vision of cloud-enabled managed services that utilize automation and a DevOps philosophy, and that blend bare-metal and virtualized environments. Although it sells self-managed cloud IaaS, Rackspace does not focus on customers who want to exclusively self-manage.

Rackspace has a large base of existing managed hosting customers to which it can sell cloud services. Its high-touch Fanatical Support approach is attractive to customers who want the flexibility of a self-managed offering but who may not want to do day-to-day management themselves. Rackspace is increasingly focused on selling solutions, not just infrastructure.

Rackspace has been instrumental in evangelizing OpenStack as a key future infrastructure ecosystem (competing against Amazon, Microsoft and VMware). Its "open cloud" message resonates with buyers. However, its Public Cloud has limited interoperability with other OpenStack-based clouds, and customers may not be able to transfer images from other OpenStack-based clouds without modification. Although its cloud IaaS offerings are currently OpenStack-based, Rackspace is introducing more VMware-based offerings to its portfolio, in response to customer demand.

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. 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, nor maintain a competitive price. Rackspace is refocusing its business upon customers that need expert managed services for mission-critical needs, rather than trying to compete directly for self-managed cloud IaaS against hyperscale 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 is focused on a hybrid cloud strategy, for customers who want managed cloud infrastructure both in their internal data centers and in Rackspace data centers. Increasingly, it will compete against large IT outsourcers that are moving down-market with lighter-weight managed services offerings that use the customer's choice of a best-in-class cloud IaaS offering and are facilitated by inexpensive, offshore labor.

Rackspace has made many cloud-related acquisitions, in order to enhance its cloud capabilities and rapidly expand the number of developers it employs. However, Rackspace has not integrated these acquisitions into a cohesive whole. Many of these acquisitions actually can manage or operate with multiple cloud IaaS providers. While this potentially positions Rackspace for future multicloud management, and enables it to take advantage of the growth of competitors, it does not create a compelling value proposition for using Rackspace's own cloud IaaS offerings.

Verizon Terremark

Verizon Terremark encompasses Verizon's data center, cloud and security businesses. Its Enterprise Cloud brand encompasses multiple VMware-virtualized offerings: the standard Enterprise Cloud (public cloud from the original Terremark) and its Private Edition (public cloud with single-tenant compute) and Public Sector Edition (U.S. federal government community cloud) variants, along with three other platforms, namely Enterprise Cloud Managed Edition ([ECME], formerly the Verizon Computing as a Service public cloud offering), vCloud Express (paid-by-the-VM public cloud, which is being retired) and the new Verizon Cloud (in beta).

Locations: Enterprise Cloud is available in multiple data centers in the U.S., as well as the Netherlands and Brazil; support, the portal and documentation are provided in English, Dutch and Portuguese. ECME is available in East and West Coast data centers in the U.S., along with the U.K., the Netherlands and Hong Kong; the service is provided in English only. Verizon Terremark has global sales.

Compute: The Enterprise Cloud is available in both single-tenant and multitenant variants, including single-tenant compute with a multitenant back end. It is VMware-virtualized. Both paid-by-the-VM and SRP billing models are available. Bare-metal servers, on daily metering, are available in ECME. ECME lacks autorestart capability. Provisioning is nonsimultaneous.

Storage: Enterprise Cloud storage is persistent and VM-independent. Although storage snapshots are supported, they cannot be used as VM images (but VM images can be copied directly into the image catalog). ECME lacks image customization capabilities, and customers cannot import their own images. Object-based cloud storage is in beta.

Network: There is no back-end load balancing. ECME does not fully support complex network topologies and cannot use customer-provided IP addresses; additionally, VMs cannot have only private IP addresses.

Security: Enterprise Cloud RBAC permissions are per-group. ECME has neither RBAC nor MFA. There is no DDoS mitigation. Only Enterprise Cloud supports most common compliance requirements.

Other notes: The monitoring service only collects metrics. Managed services are normally bundled with ECME, but are not otherwise an option.

Recommended uses: Development environments and general business applications.

Strengths

Verizon Terremark, via its standard Enterprise Cloud service, has the longest track record in the market for VMware-virtualized enterprise-class public cloud IaaS. It can address hybrid hosting use cases via ECME, which bundles in managed services but has fewer self-service capabilities. However, Verizon Terremark's competitive differentiators have eroded over time as competitors have caught up and, in some cases, surpassed it in feature development.

Verizon Terremark introduced the beta of a next-generation, unified platform, called Verizon Cloud, in late 2013. This new platform will enable it to address a much broader range of use cases, and consolidates its development efforts onto a single hypervisor-neutral platform. However, although the Verizon Cloud has an innovative architecture, it needs to deliver more than a capable cloud IaaS platform and associated network services in order to realize its cloud ambitions.

Verizon Terremark has begun to build an ecosystem around Verizon Cloud, primarily with ISVs and hardware vendors whose technology can be delivered as a virtual appliance. However, it does not have a software marketplace.

Cautions

Until Verizon Terremark launches its new unified platform, customers must be careful to match the service they choose to their particular use case. They should also be aware that while Verizon Terremark is continuing to enhance existing cloud offerings, its engineering focus has shifted to the new platform.

Although Verizon Terremark has always done a significant amount of software development for its cloud offerings, rather than being wholly reliant on VMware, it is staking its future success on rapid innovation driven by agile development. This is an unusual strategy for a company owned by a

telecom carrier, and it is highly dependent on Verizon's willingness to interfere minimally with management. It is also competing with providers that have much greater engineering resources and a much faster pace of innovation.

Verizon's sales organization does not often have a relationship with the decision makers and influencers who select cloud IaaS providers, since they are different from the people who control network procurement. Verizon needs to become relevant to business managers and application development leaders, in order for Verizon Terremark to achieve its potential.

Virtustream

Virtustream is a small, independent service provider focused solely on cloud services. In addition to its cloud IaaS offering, it sells the software for its platform, which is called xStream.

Locations: Virtustream has data centers in the eastern and western U.S., and in the U.K. and Netherlands. It has sales in the U.S., along with London and Dubai sales offices. The service is provided in English only.

Compute: xStream is hypervisor-neutral but typically supports VMware and KVM. It is offered in both single-tenant and multitenant variants; it can support single-tenant VMs in its public cloud, 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.

Storage: Block storage is VM-independent, with SSD and encryption options. While storage snapshots are supported, they cannot be used as VM images. There is no object-based storage.

Network: Virtustream supports a full range of networking options, although it does not have back-end load balancing.

Security: RBAC is per-element and very configurable. Virtustream embeds a tool for governance, risk management and compliance (GRC). In February 2014, Virtustream acquired ViewTrust Technology, and it is integrating this company's capabilities as a service.

Other notes: There is no self-service monitoring, although Virtustream provides monitoring as a service. Managed services are optional.

Recommended uses: Enterprise applications, general business applications, e-business hosting and cloud-native applications.

Strengths

Virtustream's founders have backgrounds in VMware and SAP consultancies, as well as system integration, and the company has a strongly consultative approach, as well as particular expertise in SAP; in 2013, Virtustream received a direct investment from SAP America. Its cloud is targeted primarily at production applications, but, in an unusual approach, it is targeting both traditional enterprise workloads, including ERP applications, and cloud-native applications. It has been successful at winning large-scale enterprise deals, particularly those focused on SAP and that require managed services capabilities.

Virtustream has developed its own cloud platform technology, and uses a single unified architecture across public and private offerings, within both its own data centers and customers' data centers, enabling federation across multiple clouds. Although much of its infrastructure is VMware-virtualized, it can also support other hypervisors. It has divided its portal into an administrative UI and an end-user UI, to address some ease-of-use issues.

Virtustream's 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.

Cautions

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. It does not have the resources to compete for all workloads against providers whose greater resources allow development of much broader product portfolios. Rather, it provides deep and differentiated capabilities in its focus areas.

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.

Virtustream is a small but innovative service provider, and may be an attractive target for acquisition. Its strategy will require it to attract and retain significant engineering talent as well as application expertise. It will be challenged to grow its brand awareness and to manage the lengthy sales cycles that will be common in its targeted use cases.

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) into general availability in September 2013. It is a subsidiary of EMC.

Locations: vCHS is available in multiple data centers in the U.S., as well as in the U.K. VMware has global sales. Support is available in English, French, German, Portuguese, Spanish, Hindi, Japanese, and Mandarin. The portal and documentation are available in English only.

Compute: vCHS is a public cloud IaaS offering available in two variants, Virtual Private Cloud (multitenant compute and storage) and Dedicated Cloud (single-tenant compute and multitenant storage). It most closely resembles a vCloud Datacenter Service. Both offerings use SRP pricing; only Dedicated Cloud allows the customer to oversubscribe the resources.

Storage: VM storage is persistent, but there is also VM-independent block storage.

Network: vCHS supports a full range of networking options. Third-party connectivity is via exchange (VMware Direct Connect). Traffic between vCHS data centers transits the Internet, as VMware lacks an inter-data-center private WAN.

Security: RBAC permissions are per group. There is no MFA. Some logs are only available via customer service request. VMware is still in the process of obtaining compliance-related audits and certifications.

Other notes: The monitoring service does not generate alerts to customers. There is no colocation; a partner exchange must be used instead.

Recommended uses: Development environments, general business applications, supplementing existing VMware-virtualized environments, 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 strong brand, global sales reach, and a broad base of existing customers that are deeply committed to its technologies. Its strategy for vCHS is to offer hybrid cloud options to existing VMware customers, 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.

vCHS has the solid set of basic features and pure infrastructure focus that is typical of vCloud Datacenter Service offerings. It does not use the vCD portal, although vCD is available to customers if they want to access it. The notion of "single pane of glass" manageability appeals to many IT administrators, although vCHS does not yet deliver on this promise. However, the vCHS Disaster Recovery service may be attractive to customers with on-premises VMware. VMware also possesses a suite of ITOM tools that could, in the future, be used to enhance vCHS manageability.

VMware possesses deep engineering expertise, which could potentially allow it to innovate at the infrastructure level; vCHS development is not tied to the general VMware release cycles, but VMware is changing the way it develops its core software to deliver into vCHS first. However, VMware lacks service operations experience and vCHS lacks an operational track record; VMware might find it challenging to scale the service. Furthermore, VMware focuses primarily on infrastructure, and lacks clear plans for offering additional value-added services.

Cautions

vCHS 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 vCHS within a business, but they often prefer to build an internal private cloud, and they are also often the people that the business is trying to bypass by going to cloud IaaS. VMware needs to win over these administrators with regard to vCHS, but it also needs to develop a compelling value proposition for developers. Although EMC also owns Pivotal, and Cloud Foundry-based PaaS can be hosted on vCHS, VMware is also in the process of independently building developer-oriented capabilities.

VMware has previously used service providers as its channel to market, but none of those service providers attained true scale and they were not able to maintain the level of innovation necessary in this market. VMware is now directly competing with this channel, but it faces the same technical challenges, and many of the same business challenges, as those service providers. vCHS is a new offering, and has not yet established an operational track record.

VMware needs to build an ecosystem around vCHS. It has launched a software marketplace, but the software available via the vCHS service catalog currently only includes operating systems and Microsoft SQL Server. VMware does have a strong channel in its broader business, but SIs, MSPs and value-added resellers need to learn to sell vCHS effectively and deliver value on top.

Vendors Added and Dropped

We review and adjust our inclusion criteria for Magic Quadrants and MarketScopes as markets change. As a result of these adjustments, the mix of vendors in any Magic Quadrant or MarketScope may change over time. A vendor's appearance in a Magic Quadrant or MarketScope 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

Google
VMware

Dropped

SoftLayer. SoftLayer was acquired by IBM in July 2013, and is now evaluated under IBM (SoftLayer).

Savvis and Tier 3. Tier 3 was acquired by CenturyLink in November 2013. It was merged into Savvis (already owned by CenturyLink), and the combined companies were rebranded as CenturyLink Technology Solutions. Savvis and Tier 3 are now one entity, evaluated under CenturyLink.

Inclusion and Exclusion Criteria

To be included in this 2014 Magic Quadrant, vendors had to demonstrate the following, as of February 2014:

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

Market traction and momentum. They must be among the top 15 providers for the relevant segments (public and standardized private cloud IaaS, excluding small deployments of one or two VMs), based on Gartner-estimated market shares and client inquiry service data.

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).

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 10 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 32GB 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.
- Access to a Web services API.

Vendors Considered, but Not Included

This Magic Quadrant is global in scope, but most of the providers are based in the U.S. This is a reflection of the way the market is evolving. The market has matured more quickly in the U.S. and the bulk of revenue comes from U.S.-based customers and flows to U.S.-based companies — U.S.-based IaaS providers typically derive 20% or more of their revenue from customers outside the U.S. However, all the providers in this Magic Quadrant offer their services on a global basis, and most have at least one data center in North America, Western Europe and Asia/Pacific.

Significant Europe-based providers not in this Magic Quadrant include arsys, CloudSigma, Colt, Gigas, Orange Business Services, OVH and Skyscape. Providers with significant presence in the Asia/Pacific region that are not in this Magic Quadrant include Datapipe, NTT and Tata Communications.

In the evaluations for this Magic Quadrant, we considered a variety of interesting cloud IaaS providers that did not meet the criteria for inclusion. The more distinctive ones, 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: Hosting and iLand, which offer disaster recovery as a service on their cloud IaaS platforms.

Enterprise applications: NaviSite, a Time Warner Cable Company, whose cloud IaaS platform is suitable for hosting Oracle e-Business Suite and other complex applications.

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

There are also many cloud IaaS providers that specialize in serving small businesses that typically use just one or two VMs; examples are DigitalOcean, ElasticHosts and Linode. These low-cost providers are often used as an alternative to mass-market hosting. Gartner clients typically have needs beyond what such providers can fulfill.

Even though some businesses may use PaaS in a very IaaS-like manner, we 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" and "Magic Quadrant for Enterprise Integration Platform as a Service."

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 the four 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 road maps; 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, or who are concerned about their long-term strategic investment in a particular vendor.

Sales Execution/Pricing: Providers were evaluated on their ability to address the range of buyers for IaaS, including developers and business managers, as well as IT operations organizations; 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;" 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; 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 2012 and 2013 Magic Quadrants. We have continued to raise our expectations of a provider's feature set, and we have further increased the weighting of Overall Viability, reflecting our belief that even though some providers can accomplish great things with relatively few resources, long-term success in this market will require substantial investment, as well as the ability to attract an ecosystem.

Table 1. Ability to Execute Evaluation Criteria

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 (May 2014)

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 technology companies of all sizes — 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 road maps, as relevant to the four 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 2012 and 2013 Magic Quadrants. However, we have continued to increase our expectations for the breadth and depth of a provider's vision. 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 self-service all of the necessary compute, storage, network and management capabilities — in cooperation with an ecosystem of supporting partners.

Table 2. Completeness of Vision

Evaluation Criteria

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

Source: Gartner (May 2014)

Quadrant Descriptions

Leaders

Leaders distinguish themselves by offering a service suitable for strategic adoption and having an ambitious road map. 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

There are no Challengers in this Magic Quadrant. 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. Visionaries may be new market entrants, or they may be existing providers who are reinventing their business. 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

Niche Players may be excellent providers for the use cases in which they specialize, but may not serve a broad range of use cases well, or have a broadly ambitious road map. They may be relatively new entrants to this market, or may not yet have gained significant market share. Some may have solid leadership positions in markets adjacent to this market, but are still in the relatively early stages of developing capabilities in cloud IaaS. Providers that specialize in managed services on top of a "good enough" IaaS platform may be in this category. The more highly targeted your needs, the more likely it is that there will be a Niche Player ideal for your needs.

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. There has been tremendous hype about these services, but there are also a number of use cases for which cloud IaaS delivers excellent business value. Although the market is immature, it is evolving rapidly; it has begun its journey up the Slope of Enlightenment on Gartner's "Hype Cycle for Cloud Computing, 2013." Unfortunately, there is a great deal of market confusion and many providers articulate their offerings poorly. Therefore, care should be taken when sourcing these services.

The common use cases for cloud IaaS are development and testing environments; HPC and batch processing; Internet-facing websites and Web-based applications (which may or may not have architectures specifically designed for the cloud); and non-mission-critical internal business applications. An increasing number of organizations now run mission-critical business applications on cloud IaaS, and a significant number of organizations are in the midst of migrating most or all of their infrastructure to cloud IaaS. Migrations are most frequently done to avoid major capital expenditure, such as a hardware refresh or the construction of a data center.

Initially, most businesses choose use cases that are peripheral to their organization's IT needs, but, over time, they adopt cloud IaaS for mainstream business applications as well, including mission-critical

applications, mirroring the past decade's adoption pattern of virtualization in the data center. Many businesses, especially in the midmarket, will eventually migrate away from running their own data centers in favor of relying primarily on infrastructure in the cloud. Gartner's 2013 CIO Priorities Survey indicates that 28% of CIOs expect to source all critical applications and operations via the cloud by 2016, and 55% expect to do so by 2020 (see "Hunting and Harvesting in a Digital World: The 2013 CIO Agenda").

Although some organizations still source cloud IaaS in a tactical, per-project fashion, most organizations are now looking for long-term strategic partners. This 2014 Magic Quadrant focuses on evaluating providers through the lens of their suitability for strategic adoption. We believe that while the market is still relatively immature, customers may reasonably begin making strategic choices, based on their own speed of adoption. Customers who will not have the majority of their workloads on cloud IaaS until 2016 or later may choose strategic providers whose offerings are still substantively incomplete, if they are confident that those providers will have the necessary capabilities by the time they need them. We recommend that prospective customers with immediate needs focus on finding the cloud provider that matches their anticipated use cases for the next year. In some cases, businesses may have to use multiple cloud IaaS providers to meet the needs of diverse use cases.

Market Overview

Cloud IaaS is a computing resource, along with associated storage and network resources, offered to the customer via self-service in a highly automated way, 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 PaaS and SaaS.

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. See "Evaluating Cloud Infrastructure as a Service" and its related reports to understand the range of options available in this market.

What Types of Workload Are Being Placed on Cloud IaaS?

There are three broad categories of customer needs in cloud IaaS:

- The hosting of a single application, or a closely related group of applications
- A VDC that will serve a broad range of different workloads
- Batch computing

Hosting is the most common need. For instance, a media company with a marketing microsite for a movie, a software company offering SaaS and a retailer needing a lightweight version of its e-commerce site for disaster-recovery purposes are examples of customers with hosting needs that can be fulfilled by IaaS. These are generally production applications, although there is some test and development as well. Some of these customers have mission-critical needs, while others do not.

Customers with a broad range of unrelated workloads are less common, but are growing in importance, particularly in the midmarket, where IaaS is gradually replacing or supplementing traditional data center infrastructure. The VDC is typically used very similarly to the organization's internal virtualization environment — primarily for less mission-critical production applications, or test and development environments — but is increasingly being used to run more mission-critical applications.

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. Other than the need to access large amounts of commodity compute at the lowest possible price, with little concern for infrastructure reliability, these customers typically have needs very similar to those of VDC customers, although some HPC use cases benefit from specialized hardware such as 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 both "cloud native" applications that have been architected with cloud transaction processing principles in mind (see "From OLTP to Cloud TP: The Third Era of Transaction Processing Aims to the Cloud"), 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?

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, 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.

The market is in a transition phase. The first phase of the market's development, from 2006 to 2013, was marked by the gradual maturation of the ability to offer infrastructure resources on a self-service basis. By the end of this phase, customers could expect reasonably good delivery of these capabilities from a large number of providers. But the most visionary providers tried to provide customers with capabilities beyond mere infrastructure rental, reinventing the relationship between infrastructure, applications, and management. In 2013, incumbent IT vendors began acting aggressively to counter the threat of cloud computing to their businesses. We have now entered a second phase in which providers will invest tremendous resources in developing cloud services. There are promising new entrants. Furthermore, providers that did well in the market's early years but have not kept up with the pace of innovation are getting a second chance to pursue new strategies and reinvest in engineering. Most such providers are likely to need two years to develop truly competitive capabilities. During this transition phase, market share is likely to continue to consolidate, but by the third phase of the market, in 2016 and beyond, new strong competitors may emerge.

There are many providers, but they vary widely in quality and capabilities. There are many competitors in this market; new entrants continue to launch offerings and existing providers are expanding the market segments they serve. Many of the newer market entrants are very large IT companies with considerable sales reach, which could potentially accelerate their growth; however, burdened by their legacy offerings, they might not develop superior offerings. Broadly, providers can be divided into two categories — those that are investing deeply in engineering in order to provide a rich suite of features and extensive automation for self-service enablement, and those that will provide only a basic set of IaaS features but intend to differentiate in some other way, such as via managed services, PaaS or SaaS capabilities.

The market's growth disproportionately favors the market leaders, and the market is consolidating. Many providers have solid offerings that encompass the most fundamental capability in this market — the ability to provision VMs rapidly on-demand, coupled with storage and an Internet connection. But most are finding it challenging to move beyond this point, and are finding it increasingly difficult to grow or even maintain their customer base. Customers' expectations are increasing, use cases are broadening, and many providers have neither the ambition nor the resources to compete across the full breadth of the addressable market.

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.

Customers are choosing IaaS platforms first, managed services second. To deliver greater value to customers, cloud IaaS providers must improve the quality and efficiency with which customers can manage their infrastructure. They must find ways to reduce the burden of operational chores such as patch management and backups. While manual managed services are frequently used to substitute for automated offerings, efficiency demands automation instead of operators. Consequently, the choice of an IaaS offering ultimately impacts the quality of the customer's IT operations. Customers who want to outsource the management of their infrastructure will increasingly adopt a best-in-class IaaS offering, and then seek a managed services provider to manage it, rather than choosing to adopt a "managed cloud" offering from a managed services provider that can offer only basic IaaS capabilities on its own platform. Customers may extend existing outsourcing relationships to include management of a third-party cloud IaaS offering.

Public and private cloud IaaS are converging. Service providers are increasingly using dynamic physical and logical isolation mechanisms to create "private" infrastructure within a shared, multitenant capacity pool. This allows for economies of scale, while enabling customers to meet a broader range of security and compliance requirements. See "Best Practice: Evaluate Isolation Mechanisms in Public and Private Cloud IaaS" for details on this convergence and how to choose the level of isolation you need. We believe that, over time, the leading providers will offer a single, highly flexible platform across both their own data centers and customers' data centers. As a result, this Magic Quadrant covers not only public cloud IaaS, but standardized private cloud IaaS as well.

Hybrid cloud is not yet a reality. While it is relatively straightforward to move VM images from one cloud to another, truly hybrid multicloud scenarios are rare. The tools to enable true "single pane of glass" management and seamless movement across infrastructure platforms are not mature, and there are significant differences in cloud IaaS implementations, even between providers using the same underlying CMP. Note that the claim that an ecosystem is "open" has nothing to do with actual portability. The organizations that use cloud IaaS most effectively will embrace cloud-native management, rather than allow the legacy enterprise environment to dictate their choices.

IaaS and PaaS capabilities will overlap. Cloud IaaS providers are increasingly offering middleware capabilities as a service, and are likely to add capabilities such as the provisioning and orchestration of application containers. Many leading providers will offer both IaaS and PaaS, and in many cases will blend IaaS and PaaS capabilities. The spectrum of services allows customers to decide on a trade-off between control and convenience. Customers want to develop, deploy and manage applications efficiently, and will choose the combination of capabilities that best suits their needs.

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.

Buying centers for IaaS are diverse. The early adopters in the IaaS market were developers. As the market matures, developers remain an important audience, because a great deal of IaaS adoption is 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. This is particularly true for the single-application, "hosting" side of the market. IT operations is, however, increasingly involved in IaaS sourcing, and is likely to be the primary buying center for multiple-application needs. IaaS providers vary in their ability to target these different buying centers. Furthermore, most providers focus on either a developer audience or an IT operations audience, and their feature set and style of service are oriented accordingly, although leading providers will increasingly offer capabilities attractive to both audiences.

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. 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.

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 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 (see "Cloud Computing Can Be the Singular Solution for at Least Five Use Cases"). 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.

© 2014 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](#) 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. 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."](#)