

Cool Vendors in Cloud Computing System and Application Infrastructure, 2009

David W. Cearley, Ted Chamberlin, Lydia Leong, Daryl C. Plummer, Adam W. Couture, David Mitchell Smith

A wide range of infrastructure services is emerging to allow companies to either run existing applications on shared cloud-based infrastructure, or to create new applications built specifically for cloud environments. Companies evaluating the use of the cloud as a platform for custom applications face new offerings from existing players, and an expanding set of new vendors specifically targeting the cloud.

Key Findings

- Vendors providing off-site infrastructure in a utilitylike fashion are increasingly embracing cloud computing as the model to deliver these services.
- Cloud infrastructure services will be used by many packaged software vendors to enter the cloud application services market during the next five years.
- Expect the major enterprise vendors to enter the market aggressively, providing increased competition for smaller, niche and startup companies.
- Vendors offering cloud computing infrastructure services will proliferate through 2012, with market consolidation occurring from 2013 through 2015.

Recommendations

- Experiment with external, cloud-based services in 2009, and establish governance and management models. Look for low-risk, non-mission-critical opportunities to exploit cloud-based development and/or execution services.
- Evaluate cloud system and application infrastructure service offerings on elasticity (that is, how fast services can be brought online, repurposed or decommissioned), and the degree to which they support the creation of cloud-optimized solutions.
- Consider the extent of proprietary lock-in to one vendor (for example, the use of proprietary scripting languages vs. standard languages), and develop contingencies to move to an alternate provider or to internal systems, if needed.

TABLE OF CONTENTS

Analysis	3
1.0 What You Need to Know	3
2.0 Cool Vendors.....	3
Recommended Reading.....	8

ANALYSIS

This research does not constitute an exhaustive list of vendors in any given technology area, but rather is designed to highlight interesting, new and innovative vendors, products and services. Gartner disclaims all warranties, express or implied, with respect to this research, including any warranties of merchantability or fitness for a particular purpose.

1.0 What You Need to Know

Cloud services can be separated into two broad categories — infrastructure and applications — which parallel internal IT infrastructure and applications. Cloud applications refer to end-user-focused services, such as office applications, information feeds or business processes (for example, Amazon fulfillment services). Cloud infrastructure services refer to cloud-based services used as a foundation to build and deploy applications. This foundation can be further divided into basic system infrastructure (for example, virtualized server/computing and storage) services, and a range of higher-level application infrastructure (for example, application development, integration and business process management) services.

This research examines a select group of innovative vendors offering various types of cloud infrastructure services. They represent a variety of services and architectural approaches to building and running applications "in the cloud." System and application services that are in the cloud are relatively immature, but their use has been growing rapidly. For most enterprises, 2009 will be a year of experimentation and of the use of cloud infrastructure to power leading-edge and targeted applications. Basic cloud system infrastructure services (computing and storage) will begin to mature from 2010 through 2011, while application infrastructure services reach mainstream critical mass from 2012 through 2015. As independent software vendors (ISVs) look to enter the cloud market and deliver cloud services in addition to or instead of packaged software, many will turn to cloud system and application infrastructure services as a platform to build their cloud application services during the next five years.

2.0 Cool Vendors

AppZero, Boston, Massachusetts, USA (www.appzero.com)

Analysis by Daryl Plummer

Why Cool: Automated application provisioning is a growing area of concern in cloud computing, and with AppZero, an IT organization can move applications to and from the cloud and across servers, or spread them across virtual images quickly. AppZero provides automated application provisioning by deploying applications to the cloud using a virtual application appliance (VAA) model that supports multiple operating-system (OS) images and virtual-machine (VM) interfaces. The requirement to move applications into the cloud comes with the problem that many virtual appliance solutions only allow the use of one kind of OS image. The VAA encapsulates the application and the necessary software stack needed to support it, but with zero OS components. This reduces the options an IT organization has for moving applications. AppZero provides software to create, run and maintain VAAs in Linux, Solaris and Windows environments.

Challenges: Although the ability to move applications unchanged to a cloud system infrastructure has appeal, this approach will diminish as applications are designed specifically as cloud-based services. This would leave the company struggling to get noticed. Application virtualization is a topic of debatable long-term market differentiation. AppZero does server side virtualization — which is rare — and has the imprimatur today of a different approach. Long term, however, AppZero may find that there is little room to make money in this paradigm as larger players

garner customer attention and revenue. AppZero is a small company that is relatively unknown, and may not have the market reach to gain as much visibility or as many clients as are necessary to build a groundswell of customers.

Who Should Care: IT managers seeking to reduce the cost of application hosting and who wish to move from on-premises to cloud computing will find AppZero of interest. System architects and managers seeking to rationalize and consolidate computing resources by running applications in virtual appliances also will find AppZero of interest, as will ISVs seeking a partner that can help them deliver their applications to the cloud, or that can help them deliver implementations for proofs-of-concept more efficiently.

Engine Yard, San Francisco, California, USA (www.engineyard.com)

Analysis by Ted Chamberlin and Lydia Leong

Why Cool: Engine Yard is a managed hosting provider specializing exclusively in Ruby on Rails, with deep ties to the Ruby development community. It is the result of Quality Humans, a Ruby on Rails consultancy, spinning out its operations team in 2006, and its primary funding has come from Benchmark Capital. There are many possible stacks for Ruby on Rails deployment; Engine Yard's chosen configuration uses Mongrel plus MySQL or PostgreSQL, within a clustered environment. Engine Yard's elastic offerings allow users to buy "slices" of a cluster with all components on shared infrastructure, fractional clusters with all components other than the database shared, or dedicated clusters. It also sells a service called "Solo," which allows users to deploy Engine Yard's stack on one Amazon Elastic Compute Cloud (Amazon EC2) instance, for a monthly fee.

Engine Yard is extensively involved in the open-source Ruby community, and is responsible for three of the most significant large-scale Ruby projects — the Merb framework (which will be merged with Rails to create Rails 3.0), the Rubinius VM, and the Vertebra cloud orchestration framework. Vertebra is intended to be more than just a tool to deploy and manage applications and infrastructures; it also provides a message bus and an integration mechanism built around Extensible Messaging and Presence Protocol (XMPP), a protocol typically used for instant messaging systems. Engine Yard hopes that Vertebra will also be used as a framework for orchestrating distributed computing applications across multiple clouds; as such, portions of the framework are written in Erlang, a language designed for distributed computing.

Challenges: Engine Yard is trying to transform itself into a software-as-a-service (SaaS) provider, with software that deploys and manages Ruby applications. Not only does it have to improve its service delivery capabilities, but it also needs to make the SaaS transformation without losing the ability to drive adoption of its technologies within the Ruby and broader open-source community. As Engine Yard expands beyond developer-driven adoption to reach the enterprise customers that are increasingly embracing Rails as a platform for agile development, it also must develop the processes and tools necessary to deliver the kind of support enterprises expect. Moreover, it must rally other organizations to adopt Vertebra, and make its voice heard over all the other vendors currently championing orchestration platforms and interoperability standards. Finally, as a specialist, its success is closely tied to Ruby's popularity.

Who Should Care: Organizations that have a commitment to Ruby on Rails and that are looking for concentrated expertise will benefit from the scalability, pricing and service model Engine Yard has created.

Enomaly, Toronto, Ontario, Canada (www.enomaly.com)

Analysis by Daryl Plummer and David Cearley

Why Cool: Enomaly's product, the Elastic Computing Platform (ECP), is not a cloud computing service, but is an enabling technology that is used to create global-class cloud computing platforms and deliver cloud computing services. A carrier version of ECP enables carriers and hosting providers to offer cloud computing services to their customers. The ECP can also be delivered as a platform to be hosted, thus creating a public cloud presence by many infrastructure providers. Finally, ECP can be used by end companies as a means of introducing the private cloud to the enterprise. This multitrack approach fosters an ecosystem of providers delivering cloud capabilities in pure, hybrid and private models. The hybrid enterprise cloud enables enterprise customers to create and operate an internal cloud computing environment, and provides a bridge between the internal enterprise cloud and public cloud providers, such as Amazon EC2. ECP allows enterprise IT to efficiently plan deployments, automate VM scaling and load-balancing, and analyze, configure and optimize the use of capacity, while providing direct self-service access to capacity for business users, supporting a higher level of business agility.

Enomaly introduced cloud-oriented infrastructure with its ECP in 2004 before the term "cloud" was commonly used. The company has continued to focus on elasticity, which is one of the most important issues related to the innovation inherent in the cloud. The open-source version of ECP was SourceForge's Project of the Month (August 2008), and version 2.2 was released in January 2009. The company's founder, Reuven Cohen, is an established cloud innovator who is often credited with being an originator of the topic. He has also been serving as program chair of the Institute of Electrical and Electronics Engineers' cloud computing standardization efforts and as chair of the Cloud Computing Interoperability Forum.

Challenges: Enomaly competes with other companies that also claim to offer elastic computing infrastructure, so customers will have difficulty discerning when one company is offering differentiated value. Large enterprises remain cautious about the reliability, availability and security of public cloud environments limiting the use of public cloud infrastructure services, including for "cloudbursting" — a situation where excess demand is shifted to a public cloud platform for other uses of public cloud capacity. In addition, the use of ECP by carriers and other system infrastructure service providers will provide competition to the sale of the ECP to create private clouds if more customers move to public offerings built on ECP. The possibility also exists for a carrier or service provider using ECP to switch to a different cloud infrastructure, and to keep the customer revenue that previously had been funneled to Enomaly. These issues make support for pure, hybrid and private cloud models key to Enomaly's future success, and makes them interesting test cases for cloud computing.

Who Should Care: Service providers or carriers that are seeking to become cloud providers in the public cloud will find that Enomaly provides a path to them that does not require them to build their own platform, but simply to implement this platform and focus on delivering a high-quality service around it. This can foster competitors to Amazon and others with relatively low barriers to entry. Enterprises seeking to enter cloud computing by acquiring infrastructures that will give them more than just virtualization and create internal cloud environments will find that Enomaly's ECP offers dynamic elasticity to reduce capacity planning efforts. Through this vendor, they will also find a potential path from private to public cloud offerings, so that the path will not end only in the private context.

LongJump, Sunnyvale, California, USA (www.longjump.com)

Analysis by David Mitchell Smith and David Cearley

Why Cool: LongJump is a cloud application infrastructure service provider delivering an application-platform-as-a-service (APaaS) offering. The company also has a history as a CRM vendor and offers a solution called "Relationals" (also the name of the parent company) that is delivered as a cloud application service. Its Business Application Platform, branded as LongJump, is similar in concept and origin to Force.com's underlying salesforce.com CRM

application. LongJump is a APaaS cloud platform provider that stands out. It offers a runtime environment and a multitenant development environment, and it supports the creation of multitenant applications. It also offers multitenancy configuration options that allow the user to store data in a common multitenant database, one tenant database instance that is dedicated to the user or in a mixed multitenancy model. The service offers a visual development environment and uses Java as its programming language. LongJump also offers a range of delivery options. In addition to the cloud computing service, the company supports hosting of its environment on third-party infrastructure (for example, Terremark Worldwide data center, Amazon EC2/Amazon Simple Storage Service [Amazon S3] or installation as software in an enterprise data center).

Challenges: LongJump is a fairly small private venture competing in a category that is increasingly attracting the bigger vendors, such as Microsoft and Google. Given difficult economic times and the likelihood of vendors being acquired or ceasing business, LongJump's ability to monetize its platform as an application makes it less of an issue than it is for those that don't have that ability.

Who Should Care: IT managers who are looking for flexible on-premises and off-premises deployment options using a business-oriented platform will find LongJump's approach useful.

ServePath, San Francisco, California, USA (www.servepath.com)

Analysis by Ted Chamberlin and Lydia Leong

Why Cool: Each of ServePath's product divisions does business under different brands; ServePath offers dedicated server hosting, ColoServe offers colocation, Upstream Networks offers content delivery network (CDN) services and GoGrid offers cloud infrastructure as a service. Together, these divisions allow ServePath to offer a comprehensive suite of hosting services, including offerings that hybridize cloud services with dedicated infrastructures.

GoGrid is positioned as a competitor to Amazon EC2, with pay-by-the-hour, Xen-based virtual servers, but differentiates itself from EC2 with strong customer service plus the type of options common to traditional hosters, such as managed services, private networking, storage area network (SAN) services, hardware load balancing and cooperation with compliance audits. GoGrid offers a friendly graphical user interface (GUI) based on a Google Web Toolkit (GWT) framework, as well as application programming interface (API)-based control. It is offering the specification for its representational state transfer (REST)-based API under a Creative Commons license, but is not open sourcing any source code; in essence, this simply allows other organizations that are creating clouds to share a common API call structure with GoGrid, thus letting them benefit from the tools ecosystem supporting GoGrid's API. As a benefit to GoGrid, the more clouds that share the GoGrid API, the greater the likelihood that any given tool will support it.

Challenges: Although an early leader in cloud infrastructure services, ServePath faces intensifying competition, and will be challenged to build market awareness against Amazon's aggressive developer evangelism. ServePath's ability to offer hybridized services, once a strong differentiator, is becoming more common as other hosting providers build or acquire cloud services. However, ServePath is also potentially a very attractive acquisition candidate to carriers, traditional hosters and IT services companies that want a turnkey cloud infrastructure service.

Who Should Care: Development and operations groups to which Amazon's pay-as-you-go infrastructure has appeal, or that have a need for hybrid cloud environments, should look at ServePath as one of several alternatives to Amazon EC2.

Vaultscape, San Diego, California, USA (www.vaultscape.com)

Analysis by Adam Couture

Why Cool: Vaultscape is a 2008 startup dedicated to providing cloud storage as an infrastructure utility. Unlike competitors Amazon and Nirvanix, Vaultscape's business model does not target Web 2.0 developers. Instead, the company is building out a tier of low-cost storage over the Internet for enterprises with a terabyte or more of data to archive. Vaultscape's infrastructure is just a bunch of disks (JBOD) serial ATA (SATA) drives with redundant array of independent disks (RAID) 6 protection and a self-healing file system (not global name space). Vaultscape's API is private, not public, decreasing the likelihood that customer data can be hacked. Anticipating the archival requirements of customers with sensitive data, Vaultscape has built a data center that is Statement on Auditing Standards (SAS) No. 70 certified and Payment Card Industry (PCI) Data Security Standard (DSS) compliant. The combination of commodity hardware, internally developed software and RAID 6 protection gives Vaultscape the confidence to offer a 100% data availability guarantee option.

Challenges: Vaultscape's technological challenges will be superseded by the business hurdles the company must overcome. Because it targets enterprise-class customers, a primary challenge will be attracting a name brand customer willing to entrust its archiving data to a startup and just as willing to be a reference. Vaultscape will also be challenged by new competition. Rackspace currently has a cloud storage offering in beta, and major IT vendors will be rolling out cloud storage services in 2009.

Who Should Care: Vaultscape should be of interest to customers with massive amounts of inactive or rarely accessed archiving data, such as archived e-mail, check images, MP3 archives of phone conversations or surveillance video. The ultralow cost of Vaultscape's storage infrastructure should make it especially intriguing to IT departments facing corporate mandates to trim budgets.

Voxel, New York City, New York, USA (www.voxel.net)

Analysis by Lydia Leong

Why Cool: Founded in 1999, self-funded and profitable, Voxel is a provider of Linux-based managed hosting and CDN services. It offer a Xen-based cloud hosting service that can be optionally hybridized with dedicated servers, and that has the range of options typical of traditional hosting solutions. Web GUI- and API-based dynamic provisioning and virtual-server resizing allows customers to cloudburst as suits their capacity needs. Voxel has data centers in the U.S., Europe and Asia, and offers zero-downtime, XenMotion-based live recovery for virtual servers across data centers located in the same metropolitan area.

Voxel owns a network, and is unique in offering Internet bandwidth through simple metering of outbound bytes transferred, with no additional charge for transfer over Voxel's CDN. To improve the efficiency of resource use and to encourage its customers to adopt the CDN service, Voxel has created an open-source Apache module that, when the Web server is under high load, automatically rewrites URLs to deliver content off a CDN.

Challenges: Voxel is a small company in a rapidly evolving and highly competitive industry. It currently appeals primarily to Web 2.0 companies, and has limited brand recognition beyond that core base. Its technology stack is likely to make it an attractive acquisition candidate for another provider seeking a turnkey cloud solution.

Who Should Care: Operations groups with "Web 2.0" rich applications and content, that want a managed utility hosting solution integrated with CDN services, or that need highly economical multiregion hosting, will find Voxel of interest.

RECOMMENDED READING

"Key Issues for Cloud Computing, 2009"

"Key Attributes Distinguish Cloud Computing Services"

"Cloud Computing Services: A Model for Categorizing and Characterizing Capabilities Delivered From the Cloud"

"Cloud Computing: Defining and Describing an Emerging Phenomenon"

"Contrasting Perspectives on Cloud Computing"

"Global Class: The Inspiration for Cloud Computing"

This research is part of a set of related research pieces. See "Cool Vendors 2009: Changing Models and Changing Times" for an overview.

REGIONAL HEADQUARTERS

Corporate Headquarters

56 Top Gallant Road
Stamford, CT 06902-7700
U.S.A.
+1 203 964 0096

European Headquarters

Tamesis
The Glanty
Egham
Surrey, TW20 9AW
UNITED KINGDOM
+44 1784 431611

Asia/Pacific Headquarters

Gartner Australasia Pty. Ltd.
Level 9, 141 Walker Street
North Sydney
New South Wales 2060
AUSTRALIA
+61 2 9459 4600

Japan Headquarters

Gartner Japan Ltd.
Aobadai Hills, 6F
7-7, Aobadai, 4-chome
Meguro-ku, Tokyo 153-0042
JAPAN
+81 3 3481 3670

Latin America Headquarters

Gartner do Brazil
Av. das Nações Unidas, 12551
9º andar—World Trade Center
04578-903—São Paulo SP
BRAZIL
+55 11 3443 1509