



Moving your Infrastructure to the Cloud

How to Maximize Benefits and Avoid Pitfalls

Table of Contents

Executive Summary	1
Introduction	2
Definitions	3
IaaS	4
Cost Savings of the Cloud	4
Marginal cost - The Economies of Scale	4
Burstability - The Utility Model	6
From Capex to Opex	9
Focusing on Core Business	9
Realizing the Savings Through Proper	
Vendor Selection	10
Billing capabilities	10
Fixed costs	11
Minimum charges	11
Performance and Support	12
Performance	12
<i>Service Level Agreements</i>	
<i>Transparency in performance</i>	
<i>Persistent vs. Ephemeral Storage, CPU Power and more</i>	
Support	13
Service through Standards	14
<i>Performance through open standards</i>	
<i>Flexibility through open standards</i>	
Security and Compliance	15
Conclusion	17
About Diversity Analysis	18
About Rackspace	19
About the author	19
Ben Kepes	19

Executive Summary

With Cloud Computing becoming more widely utilized, it is important for organizations to understand ways to maximize benefits and minimize risks of a move to the cloud. This paper details the significant benefits that Cloud Computing brings and provides guidance to IT decision makers to help their decision making process. This is especially important given the plethora of vendors in the marketplace today. Buyers need to appreciate that assessing individual providers is critical to the success of Cloud Computing programs.

Introduction

Cloud Computing is a hot topic these days, with interest shown across all levels of the organization - from the C-suite, to Corporate IT through to end users. These various groups are all interested to see the potential benefits and possible pitfalls from Cloud Computing. Analyst firms have been quick to document both the current acceleration in Cloud Computing adoption and the potential market size in the years ahead. Daryl Plummer, analyst at Gartner says ¹ that “Adoption of the cloud is rising rapidly - there’s no sign that it’s going back. We have to think about how [we will] deal with this, more than likely, you will be dealing with hundreds of cloud services [by 2015]”. Other analysts concur with a recent IDC survey for example finding that “savvy CIOs now see the cloud as being an extension of their sourcing strategies”²

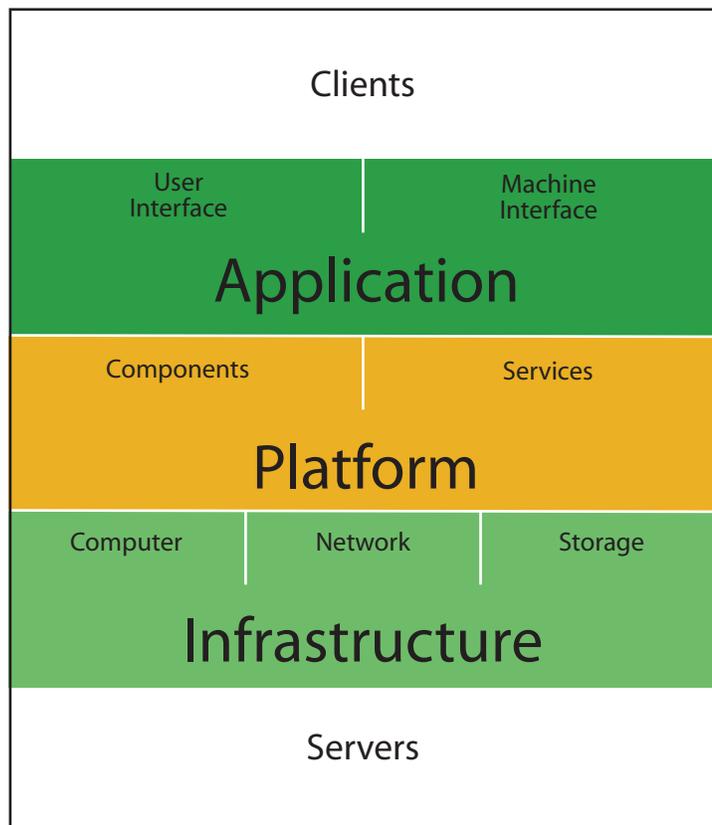
This interest has tended to create two distinct camps however. On the one hand cloud evangelists who seek to push cloud adoption immediately despite potential risks. On the other hand the conservative traditional technologists who seek to dismiss Cloud Computing as either the same as what has gone before, grossly over-hyped or even potentially dangerous to the organization itself.

We believe that Cloud Computing brings unprecedented benefits to organizations and will be the dominant way of acquiring technology in the future. In keeping with the highly important place Cloud Computing will have for organizations, we further believe that it is vital that leaders within business who are tasked with making technology acquisition decisions should be well appraised of both the benefits and the risks of Cloud Computing and should therefore be well placed to make good buying decisions.

This paper seeks to give decision makers some guidance so that they are well informed to ensure maximum cost savings, highest performance levels and strictest security from the infrastructure they ultimately invest in. This is important because, as we will

show, the barriers to entry for a Cloud Computing vendor are relatively low. While this may mean there is more choice in the marketplace, it also introduces the potential for underfunded and low quality providers to enter the market and therefore organizations looking at moving to the cloud need to be especially cognisant of the various issues outlined in the following pages.

Definitions



Cloud Computing is often depicted as a layered stack, with infrastructure, platform and application components

Before looking at the different issues buyers need to be aware of, it's worth setting the scene and defining Cloud Computing. We broadly accept the Wikipedia definition of Cloud Computing as follows: "Cloud Computing is Internet-based computing, whereby shared resources, software, and information are provided to computers and other devices on-demand, like the electricity grid."³ Generally Cloud Computing is divided into three distinct areas

– SaaS (Software-as-a-Service), PaaS (Platform-as-a-Service) and IaaS (Infrastructure-as-a-Service). This paper will focus on the IaaS level of the stack.

IaaS

Infrastructure as a Service, delivers computer infrastructure - typically a platform virtualization environment - as a service. Rather than purchasing servers, software, datacenter space or network equipment, clients instead buy those resources as a fully outsourced service on demand.

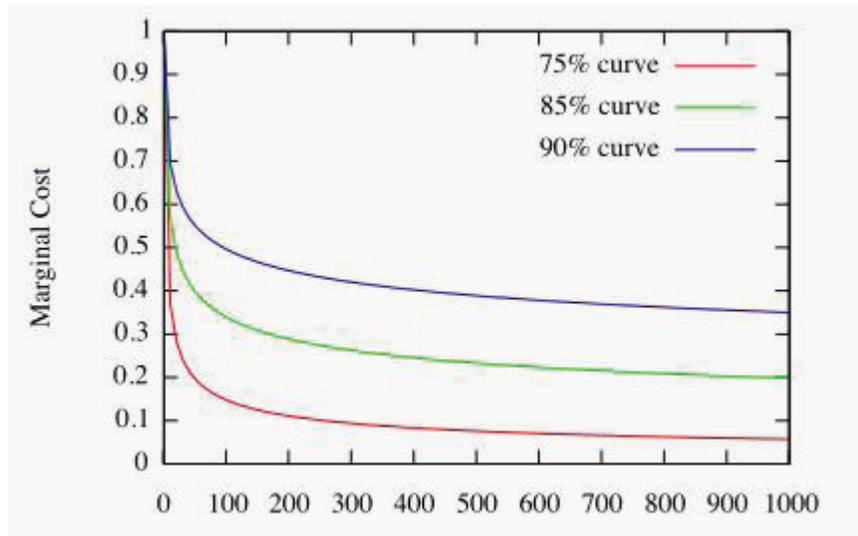
Within IaaS, there are some sub-categories that, while outside the scope of this paper, are worth noting. Generally IaaS can be obtained as public or private infrastructure or a combination of the two. “Public cloud” is considered infrastructure that consists of shared resources, deployed on a self-service basis over the Internet. By contrast, “private cloud” is infrastructure that emulates Cloud Computing but does so on a private network. Additionally, some hosting providers are beginning to offer a combination of traditional dedicated hosting alongside public and/or private cloud networks.

Cost Savings of the Cloud

While every situation will be different, it is useful to articulate the potential areas for cost savings that Cloud Computing can bring to an organization. The two biggest savings that are gained from a move to Cloud Computing come from economies of scale and a utility-based pricing model. We will look at each of these in turn.

Marginal cost - The Economies of Scale

The Boston Consulting Group (BCG) developed the concept of experience curves to describe how production costs tend to fall in a predictable fashion as the number of units produced increase. In other words, the more you produce (or the higher number of units you provide to the marketplace) the lower the unit cost will be.



This experience curve graph shows why private organizations are unlikely to beat the marginal cost of Cloud Computing infrastructure.

Using some analysis of known datacenter and Cloud Computing costs, Jon Moore found⁴ that an organization would have to deploy between 5000 and 15000 servers to get the per-server marginal cost to a similar level enjoyed by Cloud Computing vendors.

A recent report by O’Reilly Media, demonstrated the relative costs of internal IT⁵, Managed Services and cloud infrastructure. It is important to realize that, if a reputable cloud vendor is chosen, the quality and availability of the cloud service should be significantly higher than both internal IT and managed services.

	Internal IT	Managed Services	The Cloud
Capital Investment	\$40000	\$0	\$0
Setup Costs	\$10000	\$5000	\$1000
Monthly Services	\$0	\$4000	\$2400
Monthly Labor	\$3200	\$0	\$1000
Cost over Three Years	\$149000	\$129000	\$106000
Savings gained	0%	13%	29%

Estimated costs of infrastructure for two application servers, two database servers and a load balancer across internal, managed and cloud deployment models. Source O’Reilly Media, George Reese⁶.

Adding further weight to the marginal cost argument is the fact that Cloud Computing vendors, with their massive economies of scale, can amortize the cost of world class engineers and operations staff across many thousands of individual servers. Part of the marginal cost argument suggests then that even when an organization is able to achieve hardware costs approaching those enjoyed by Cloud Computing vendors, they would be unlikely to obtain staff of the calibre employed by those vendors.

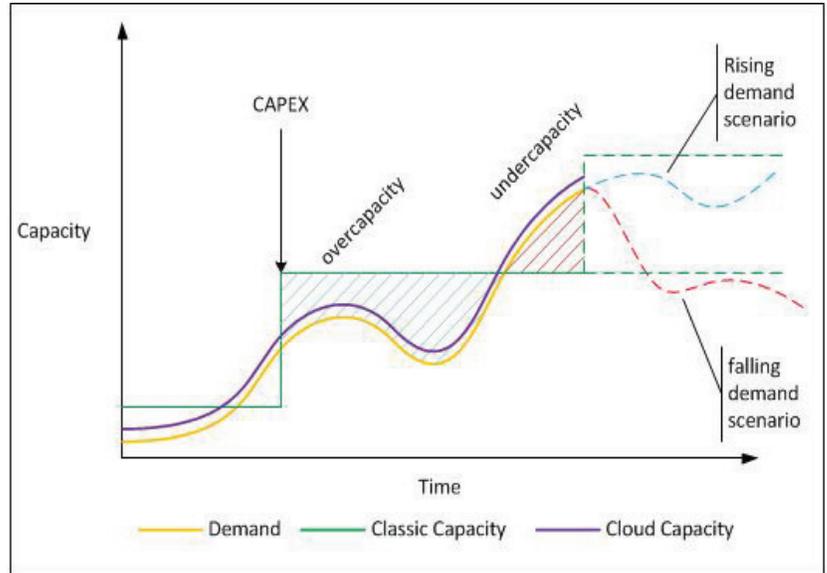
Burstability - The Utility Model

Much has been written likening Cloud Computing to a traditional utility service, like water or electricity. Nicholas Carr's book *The Big Switch*⁷ is an excellent introductory read to this subject. In the book Carr says that:

...fragmentation [of IT into separately installed and managed systems that individual corporations own] is wasteful. It imposes large capital investments and heavy fixed costs on firms, and it leads to redundant expenditures and high levels of overcapacity, both in the technology itself and in the labor force operating it...Once it becomes possible to provide the technology centrally, large-scale utility suppliers arise to displace the private providers...in the end the savings offered by utilities become too compelling to resist, even for the largest enterprises. The grid wins.

When planning infrastructure requirements, organizations are forced to use peak load criteria when sizing their resource needs. What this means is that resources need to be provided for the highest possible load situation. The graph below illustrates this point. The green line is on-premise infrastructure and shows how large, capital-intensive onsite infrastructure either leaves an organization with excess capacity that is under-utilized or with an excess in demand leading to service degradation.

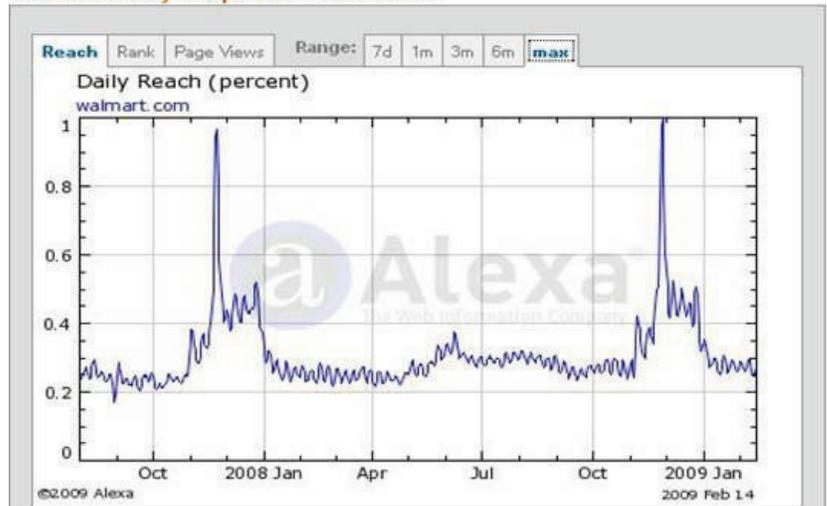
The purple line, closely following actual demand shows the benefits of Cloud Computing infrastructure to scale up and down readily along with actual demand.



Capacity vs Utilization curves⁸

In another example, looking at some real demand curves shows why fixed capacity infrastructure is problematic. The image below indicates an actual traffic history for Walmart.com and indicates that off-peak load is in the vicinity of 20% of peak load.

Traffic History Graph for walmart.com



Traffic history of Walmart.com⁹

Cloud Computing, with its elastic capacity, far reduces the need to forecast website traffic for capacity planning. Vendors are able to aggregate demand for multiple websites and in doing so smoothen the total demand for infrastructure. In this way individual organizations can enjoy the ability to meet their peak demands, without having to provide infrastructure that will often remain idle.

A more detailed description of the laws of Cloud Computing economics, while outside the ambit of this report, has been published by Business Week.¹⁰ It is worth noting that, despite the per unit price from a utility service provider potentially being higher than an owned resource through the addition of a utility margin, aggregate cost will be reduced. As with other utilities, cloud service providers should only charge for utilized resources, so although utilities may cost more when they are used, they cost nothing when they are not and customers can save money - particularly when workloads are spiky as in the above example.

From Capex to Opex

A consistent trend, accelerated by the global financial crisis, is a trend for financial decision makers to pare back on new capital investments (CapEx), opting wherever possible to fund projects from operating expenditures (OpEx) instead.

The convergence of rapid technological changes, fiscal pressures and imperatives and the move to a highly dynamic organizational structure heighten the demand for a move

The table below provides a simple differentiation between CapEx and OpEx:

	Capital Expenses (CapEx)	Operating Expenses (OpEx)
Purpose	To buy assets with a useful life beyond the current year	Ongoing costs to run a business
When paid	Lump sum (or financed, with extra charges)	Monthly
When accounted for	Over 3 to 10 years, as an asset depreciates	In the current month or year
Listed as	Property or equipment, depreciation	Operating cost
Tax treatment	Deducted over time as asset depreciates	Deducted in current tax year
Example	Buying a laser printer	Buying paper and toner for printer

*Differentiating CapEx from OpEx*¹¹

Since Cloud Computing vendors often bill on a pay-as-you-go basis, acquiring IT infrastructure can become more like buying paper and toner for a laser printer, than buying the laser printer itself.

Focusing on Core Business

Cloud Computing allows organizations to focus on their core business and abstract responsibility for what are essentially commodity services to a third party. It is the contention of the author that IT departments should be performing highly strategic work - basic maintenance of infrastructure is both a negative drain on their time but also arguably something that is better achieved by highly skilled specialists rather than the generalists

that IT staffers tend to be. It has even been estimated ¹² that 80% of IT expenditure goes to simply “keeping the lights on” for an organization.

Realizing the Savings Through Proper Vendor Selection

The two reasons for reduced total cost of ownership (TCO) from Cloud Computing, economies of scale and utility pricing, presuppose that customers will have these benefits passed on to them by the vendors. For this reason it is important to assess a number of issues relating to Cloud Computing vendor pricing, billing capabilities, fixed costs and contract minimums.

Billing capabilities

In order to enjoy cost benefits from Cloud Computing, billing needs to be as granular as possible. That is, the finer the utility cost is divided, the higher the savings are likely to be. In the example of highly variable websites, often the biggest peaks are very short lived, lasting only a few hours. Cloud Computing infrastructure should therefore be able to be billed on an hourly basis. The example below, while simplistic in its approach, illustrates this fact.

Example.com is a “daily deals” site that experiences peak loading for two hours a day of roughly 10 times its base load. Base load for example.com is two individual servers. In this example both vendors charge \$1 per server hour.

Vendor A charges on a daily rate - as example.com has a peak of 20 servers during the day; their costs would be calculated as:

$$((24 \text{ hours} * 20 \text{ servers})) * \$1 = \$480$$

Vendor B charges per hour - in this case example.com usage is calculated simply as:

$$((22 \text{ hours} * 2 \text{ servers}) + (2 \text{ hours} * 20 \text{ servers})) * \$1 = \$84$$

The difference in these two scenarios equates to \$396 a day because Vendor A bases its fee off of maximum daily usage, rather than the more flexible hourly usage.

Fixed costs

Many Cloud Computing vendors seek to offset their initial costs by charging users a setup fee. This setup fee is problematic for two reasons. Firstly it increases the average cost of the services themselves. More seriously however it reduces the likelihood that organizations will be able to truly enjoy the scalability of the cloud. If fees are charged to provision (or even worse, both provision and de-provision) infrastructure, organizations will tend to reduce their provisioning events by averaging demand. For this reason services from providers who charge setup fees are sub-optimal and should be avoided.

Minimum charges

It is important to assess whether the prospective vendor has a minimum charge for cloud compute and storage services, and whether they demand a minimum contract period. Part of the benefit of Cloud Computing is the ability to both scale infrastructure to peak demand, and also move work-loads between providers at will. Any minimum charges or contract periods for compute or storage services are a barrier to this flexibility and should be avoided. Fixed fees for other services, such as advanced monitoring or professional services might make sense when they are combined with utility compute and storage. These fees should be assessed on a case-by-case basis.

Performance and Support

As stated previously, Cloud Computing has the tendency to democratize technology and reduce the barriers to entry. While this is a positive trait, it does introduce some risks as operators can launch with less than excellent performance and levels of support. It is therefore very important to understand and assess the performance and service factors of Cloud Computing.

Performance

With increasing demands, and reduced resourcing for IT departments, it is attractive to avoid relying on already over-stretched resources to maintain in-house infrastructure, if an equal or greater level of performance can be obtained through a Cloud Computing vendor. The following section introduces some of the most important aspects to consider when evaluating vendor performance.

Service Level Agreements

Strong Service Level Agreements (SLAs) from Cloud Computing vendors are a must to ensure performance. Without these agreements, and penalties for failing to meet them, vendors have less of an incentive to maintain performance at the highest levels. SLAs can include factors such as network availability, datacenter uptime, host failure and migration.

Transparency in performance

While we maintain that Cloud Computing is generally more reliable than running infrastructure in-house, no Cloud Computing vendor is immune from service outages, even with strong SLAs in place. Reputable Cloud Computing vendors are transparent about performance and care should be taken in selecting a vendor who is open about things like service outages. Many Cloud Computing vendors will have real time dashboards to indicate service status and are likely to have support staff available 24x7 in the event of any problems.

Persistent vs. Ephemeral Storage, CPU Power and more

In addition to evaluating vendor SLAs and transparency policies, it is important for buyers to be aware of the technical facets of Cloud Computing that affect performance and the difference between vendors. For example, it is important to be aware that some cloud services do not provide “persistence”, meaning that in the event of a failure, local data will be lost. Depending on the type of workloads that will be moved to the cloud, it may be appropriate for the organization to ensure that server images are persistent in the event of a failure.

Similarly, dependent on the organization’s needs, it may be important to ensure server size and CPU power are appropriate for particular use cases.

Support

When assessing a move from in-house to cloud infrastructure, but also when assessing different cloud vendors, it is important to ascertain the quality of support that the vendor can provide.

In-house IT is unlikely to have the scale to be able to offer dedicated and round the clock support. There is a significant risk then for service degradations and even outages to occur in an in-house IT situation. This can be contrasted with Cloud Computing where the very core business of the vendor is to deliver high quality and highly available infrastructure services.

We believe therefore that robust support should be an included feature of cloud infrastructure. We advise prospective Cloud Computing customers to ensure that support is built into the base price of the service. While total cost for a service-included solution may be slightly higher than for one that comes without service, outsourcing technology requires an ongoing relationship and hence service should be considered integral to the solution.

One indicator of high-quality support is the ability to get in contact with support staff 24x7x365 through multiple channels such as telephone, chat or email. Another factor to consider is whether or not the vendor takes responsibility for upgrades and security patches.

Service through Standards

Cloud Computing infrastructure can either be available in propriety formats and protocols, or in broadly accepted open standards. It is our contention that open standards increase the overall quality of performance as well as give users flexibility and security over their infrastructure.

Performance through open standards

While outside of the scope of this report, we contend that open standards encourage the creation of a network of related and ancillary products and services. This ecosystem will include infrastructure monitoring and management tools which can help increase the overall performance of the solution.

Flexibility through open standards

We strongly believe that organizations should have the ability to move their infrastructure between different providers at will. This ability gives them the flexibility to react to a changing business environment, allows them to move between vendors as price and service offerings change, and gives them a degree of control by removing vendor lock-in. We advise organizations to think about solution providers that embrace open standards.

Security and Compliance

Security should be of general concern for technologists - Cloud Computing does nothing to change that. Any organization embarking on infrastructure spend needs to ensure a thorough due-diligence occurs in terms of security - whether the infrastructure be on-premise or cloud based.

Notwithstanding this we believe that cloud infrastructure provides a level of security beyond the capability of most on-premise deployments. There are a number of specialist services that cloud providers can provide that, due to their cost or complexity, would generally be unavailable in traditional infrastructure situations.

Unfortunately there is a significant level of doubt created by traditional vendors regarding the security of Cloud Computing and it is important to understand how security is handled in the cloud. In order to ensure the highest standards of security for their customer data, cloud providers employ world standard security professionals and follow industry best practices to keep data safe.

It has to also be stressed however that security is a shared responsibility between vendors and customers. While vendors have an obligation to ensure datacenters are built to the highest level of reliability and security, it is also incumbent upon users to ensure they use best practices in areas they have control over – the utilization of firewalls, strong passwords and employee vetting are some of the areas that are beyond the control of the Cloud Computing vendor.

The other important issue for businesses using a cloud vendor is compliance. Many businesses are required to be compliant with a host of different regulations. Relevant acts and standards that businesses may need to comply with include:

- Health Insurance Portability and Accountability Act (HIPAA) relating to the security and privacy of health data
- Statement on Auditing Standards No. 70 (SAS70) relating to the internal control of service organizations
- SSAE16, a new standard that will apply from 2011

- Payment Card Industry Data Security Standard (PCI DSS) relating to the security and privacy of credit card information

When an organization puts its data in third party datacenters, it is important that the outsourcing vendors are also compliant with these requirements. Even though the regulatory requirements vary depending on the area of business and legal jurisdiction, most cloud vendors are compliant with SAS 70 Type II audit, a global auditing standard designed to evaluate and issue an opinion on a service organization's controls (SAS70 will be superseded by SSAE16 in 2011).

Cloud vendors hosting data of clients from the financial sector also need to be compliant with the higher standards of PCI DSS, while those hosting data from the medical sector need to comply with HIPAA. For these customers it may be appropriate to ascertain whether the vendor supports a dedicated/hybrid capability for PCI and HIPAA Compliance standards.

While by no means an exhaustive list, we would advise customers to ask prospective vendors the following questions:

- What is your approach to service security? Can you offer an overview of your general security approach?
- What security procedures are in place at the datacenter?
How many technicians have access to my data and how well are those technicians vetted before they are given access?
How do you maintain physical security of your datacenter?
- What are the security measures you use to authenticate users?
- What level of encryption do you offer to protect my data?
- How secure is your application and do you work with any independent security vendors to vet the overall security of your product?
- Are you compliant with the particular regulations applicable to my business?

Conclusion

Cloud Computing is undoubtedly a growing revolution within IT. However, as with any paradigm shift, there are risks and rewards which need to be finely balanced.

While we believe that Cloud Computing brings unprecedented benefits to organizations and will be the dominant way of acquiring technology in the future, we also realize that IT decision makes are confused given the sheer number of vendors in the marketplace. For this reason IT decision makers need to be well appraised of the issues around price, performance and security of their chosen supplier.

We recommend that a full due diligence process be undertaken when assessing vendors and that, where appropriate, independent advice should be obtained that matches the product offering with the business requirements.

About Diversity Analysis

Diversity Analysis is a broad spectrum consultancy specialising in SaaS, Cloud Computing and business strategy. Our research focuses on the trends in these areas with greater emphasis on technology, business strategies, mergers and acquisitions. The extensive experience of our analysts in the field and our closer interactions with both vendors and users of these technologies puts us in a unique position to understand their perspectives perfectly and, also, to offer our analysis to match their needs. Our Analysts take a deep dive into the latest technological developments in the above mentioned areas. This, in turn, helps our clients stay ahead of the competition by taking advantage of these newer technologies and, also, by understanding any pitfalls they have to avoid.

Our Offerings: We offer both analysis and consultancy in the areas related to SaaS and Cloud Computing. Our focus is on technology, business strategy, mergers and acquisitions. Our methodology is structured as follows:

- Research Alerts
- Research Briefings
- Whitepapers
- Case Studies

We also participate in various conferences and are available for vendor briefings through Telephone and/or Voice Over IP.



About Rackspace

Rackspace Hosting is the world's leading specialist in hosting and Cloud Computing. The San Antonio-based company provides Fanatical Support® to its customers, across a portfolio of IT services, including Managed Hosting and Cloud Computing. Rackspace is also the founder of OpenStack™, an open source cloud platform with broad industry support, designed to offer cloud consumers greater choice. For more information, visit www.rackspace.com.



About the author

Ben Kepes

Ben is the founder and managing director of Diversity Limited, a consultancy specializing in Cloud Computing/SaaS, Collaboration, Business strategy and user-centric design. More information on Ben and Diversity Limited can be found at <http://diversity.net.nz>

- [1] http://blogs.computerworld.com/17008/single_largest_opportunity_in_cloud_computing
- [2] <http://www.networkworld.com/news/2010/090710-more-private-cloud-adoption-expected.html>
- [3] http://en.wikipedia.org/wiki/Cloud_computing
- [4] <http://www.cloudcomputingeconomics.com/2009/01/experience-curves-for-data-center.html>
- [5] <http://broadcast.oreilly.com/2008/10/the-economics-of-cloud-c.html>
- [6] <http://broadcast.oreilly.com/2008/10/the-economics-of-cloud-c.html> for more information about the economics of Cloud Computing see also <http://gigaom.com/2010/06/06/lazy-hazy-crazy-the-10-laws-of-behavioral-cloudonomics/>
- [7] <http://www.amazon.com/Big-Switch-Rewiring-Edison-Google/dp/0393062287>
- [8] <http://www.chades.net/>
- [9] Alexa
- [10] http://www.businessweek.com/technology/content/sep2008/tc2008095_942690.htm
- [11] http://broadcast.rackspace.com/hosting_knowledge/whitepapers/GettingOnTheRightSideOfTheCapexVsOpexDivide.pdf
- [12] Gartner Report
<http://www.gartner.com/it/page.jsp?id=497088>