



CLOUD UNIVERSITY

TEXTBOOK VOLUME I

A PRIMER FOR 'THE CLOUD'

TABLE OF CONTENTS

■	■ Preface	2
	■ How to Use This Book	3
■	Chapter 1: Revolution Not Evolution	4
	■ Virtualization	5
	■ Democratization of Computing	6
	■ Scalability and Fast Provisioning	9
	■ Commoditization of infrastructure	11
■	Chapter 2: Clouconomics	13
	■ The Problem with the 80-20 Rule	14
	■ Remember the Opportunity Cost	16
	■ OpEx is the new CapEx	18
	■ Total Cost of Ownership	21
	■ Time is Money	23
■	Chapter 3: Understanding the Cloud Computing Stack	26
	■ Software as a Service	27
	■ Platform as a Service	31
	■ Infrastructure as a Service	35
■	Final Words...	39

PREFACE

AS A CLOUD COMPUTING BLOGGER, ANALYST AND EVANGELIST,



I spend much of my time talking to vendors, customers and the general public about Cloud Computing. I run CloudCamps, non-profit Cloud-focused unconferences, all across Australia and New Zealand. I also follow the market closely and cover the latest developments on my own blog.

All of these varied activities, and my own history owning and running business in varied sectors, has shown me two things; firstly just how powerful the move to Cloud Computing can be and secondly, that small and mid-size businesses are so busy keeping up with day to day activities that they really don't have time to educate themselves on what the Cloud means for them.

With this realization, I decided to do something about the gap, and launched CloudU, a comprehensive series of reports, webinars, events and now books, aimed at increasing Cloud awareness among businesses. To create the program, I collaborated with Rackspace, a company that has a point of view that matches my own regarding Cloud computing. The content of CloudU isn't vendor-specific however, anyone looking at moving to the Cloud can benefit from the content – regardless of which vendor they're considering.


Hence, the first CloudU textbook, A Primer to The Cloud. In this textbook, I share details from the first three CloudU whitepapers along with highlights from our CloudU Notebooks series of weekly blog posts. I hope you get as much out of reading this book as I've gotten out of creating it. Thanks for your time.

A handwritten signature in black ink, appearing to be 'Ben Kepes'.

Ben Kepes, CloudU Curator

HOW TO USE THIS BOOK

A Primer for 'The Cloud'



The first volume of the CloudU textbooks, A Primer to The Cloud, is two things. It is both an instruction and learning manual for business owners and technical folks wanting to learn about the Cloud, and also a record of some great stories that illustrate the theoretical insights we're trying to get across. You'll notice as you go through the book that we've included pop-out pages that tell the story about different organizations that are taking advantage of various aspects of Cloud Computing. These snippets are by no means required reading, but they do illustrate the points we're making and help readers to understand just how empowering Cloud Computing is.

CHAPTER 1: REVOLUTION NOT EVOLUTION

How Cloud Computing differs from traditional IT and why it matters



Cloud Computing is a revolution that will define IT beginning in the second decade of the 21st Century. This new form of computing is perfectly poised to provide solutions to a host of business problems within organizations large and small. Cloud Computing will be the catalyst for the long predicted notion of “ubiquitous computing.” So just what “is” Cloud Computing, and why it is so different from what has come before? The following pages will detail four main areas in which Cloud Computing allows businesses to break from the past:

- Virtualization – The ability to increase computing efficiency
- Democratization of Computing – Bringing enterprise scale infrastructure to small and mid-size businesses
- Scalability and Fast Provisioning – Bringing web scale IT at a rapid pace
- Commoditization of Infrastructure – Enabling IT to focus on the strategic aspects of its role

VIRTUALIZATION

THE ABILITY TO INCREASE COMPUTING EFFICIENCY

Picture if you will a traditional server. Housed in a noisy cabinet somewhere, a server is a computer that can generally do one thing at a time. It may fill the role of email server, database server or web server but running multiple processes concurrently risks reliability and efficiency so typically servers are operated as single use machines. While this may be a robust way of providing a service, it is inefficient since many times servers have excess processing capacity above what is used by a single application.

Virtualization was developed to overcome this limitation of physical hardware as it enables multiple pseudo-servers to be run on one physical device. This division of a single physical server into multiple “virtual” servers containing multiple sets of segregated data is the backbone of Cloud Computing as it allows for far greater flexibility and resource utilization.

Virtualization not only brings efficiency gains in terms of processing power but also saves electric power, space and cooling since the number of physical machines running is greatly reduced. To illustrate this point, studies¹ have found that Cloud applications consume 90% less energy than on-premise ones.

While virtualization is a Cloud Computing enabler, Cloud Computing itself enables some major shifts within organizations. The first of these is the tendency for Cloud Computing to democratize technology in a way not previously seen.

DEMOCRATIZATION OF COMPUTING

BRINGING ENTERPRISE SCALE INFRASTRUCTURE TO SMALL AND MID-SIZE BUSINESSES

Cloud Computing is facilitating a seismic shift in terms of business development. Formerly entrepreneurs who wished to start a business had to invest significant capital into hardware and software licenses. Even the simplest of businesses required expensive software licenses, a server or two and the associated administration cost of keeping it all running.

The availability of Cloud Computing solutions has led to a massive shift in the availability of computing power. It is now almost effortless for an entrepreneur to set themselves up with some infrastructure and applications upon which to run their business. For many, Cloud Computing provides a server capable of running many of the most common web or business applications. A recent study by Github² indicates that less than 25% of Y Combinator start-ups are self hosting their web infrastructure. No longer is enterprise scale infrastructure the exclusive domain of enterprises.

This democratization is analogous to the widespread availability of the word processor. Formerly the creation of documents was the sole preserve of the typing pool, an overworked shared resource which required specialized, skilled workers to function. Office productivity applications enabled even the least dexterous of executives to create professional reports and letters, all from the comfort of their own PC. This move, while arguably detrimental to those who made their living working in a typing pool, greatly increased the efficiencies and timeliness of document creation.

² <http://jpf.github.com/domain-profiler/ycombinator.html?2010>

The ease, economics and speed of provisioning Cloud Computing resources is enabling an entire generation of businesses to be founded. One needs only look at the meteoric growth of question and answer site Quora.com which in December 2009 began to experience usage spikes of 5 to 10 times its normal load. By utilizing Cloud Computing for their infrastructure needs, Quora was able to handle the load with relatively few issues.³ It is our contention that in a few years time, the need to wait for the provisioning of infrastructure, the need for skilled engineers to perform the tasks and the need for administrative staff to “keep the lights on” will be but a distant memory.

CASE STUDY

Recently an announcement caught my eye that really speaks to the contention that Cloud Computing democratizes IT. Cloud Computing is facilitating a seismic shift in terms of business development. Case in point, I was excited by the acquisition announcement made recently by some of my friends who started up a great little Learning Management System (LMS), Litmos. They announced that they were being acquired by the massive software company Callidus.

Litmos grew their business to the point of acquisition by using democratized technology, from economical internal development tools, to Cloud storage, and from web based email platforms to SaaS accounting. While it is impossible to argue that startups didn't grow before the advent of the Cloud, I believe that Cloud has enabled these startups to grow more quickly, more easily and more flexibly than before. As mentioned previously, a recent study of the services that technology incubator Y Combinator companies use reveals that, indeed, startups are embracing the Cloud over and above running technology in house.




SCALABILITY AND FAST PROVISIONING

BRINGING WEB SCALE IT AT A RAPID PACE

Traditionally infrastructure has suffered from a “boom and bust” cycle where organizations swing between under-provisioning and over-provisioning. For any organization where workload is erratic, there will always be one of two situations;

- 1 An over provision of servers creating unused capacity
- 2 An under provision of servers causing service degradation

Cloud Computing enables organizations to maintain infrastructure to required levels at all times, as such it enables cost savings to be gained by virtue of the fact that, despite the per unit price from a utility service provider potentially being higher than an owned resource, aggregate cost can be reduced by paying only for what is required, when it is required.



CASE STUDY

One of my favorite examples of how to avoid this “boom and bust” cycle can be illustrated by the much talked about deal-of-the-day website Groupon, which is enjoying staggering growth and buying competitor companies all around the world. This sort of accelerated acquisition strategy is great, but it makes life really difficult for those in the back offices trying to manage capacity. Wanting to utilize an accounting and back office solution that was consistent across all their different locations, Groupon decided that a Cloud solution would give them the ultimate levels of scalability, flexibility and most importantly agility.

As it happened Groupon chose to go with NetSuite, a well-known Cloud Enterprise Resource Planning (ERP) vendor. What is really important here is that it was a Cloud Computing solution that gave Groupon the ability to bring new locations on quickly while making sure capacity met demand. The time, hassle and process involved with doing that on traditional software would have put a serious dampener on their growth. The numbers are amazing! Groupon went live with their Cloud solution in six weeks, across five international markets and 26 different countries. All of this was possible with solution that, due to the fact that it lives “in the Clouds” inherently allows for collaboration, scale and flexibility.



COMMODITIZATION OF INFRASTRUCTURE

ENABLING IT TO FOCUS ON THE STRATEGIC ASPECTS OF IT'S ROLE

It's hard to overstate just how much the IT role is changing in the face of Cloud Computing. Cloud Computing vendors are often quick to use cost reduction as their main selling point for the Cloud, but it would appear that users are more thoughtful than this and perceive the business agility gains to be the number one benefit of a move to the Cloud. A recent SandHill report⁴ found that around 50% of respondents consider agility as their primary reason for adopting the Cloud.

It's not hard to believe these statistics when one remembers the estimates that put IT maintenance at around 80% of total IT expenditure.⁵ When one considers that Cloud infrastructure is still nascent and the vast majority of servers are still managed in-house, this presents a significant opportunity and a significant change agent for traditional IT. The underlying trend here is pressure upon IT departments to produce greater outputs, with less resourcing. Cloud Computing offers the ability for IT departments to apply resources as, and where they are needed.

Clearly the savings to be gained from a move to the Cloud free up IT resources for adding business value rather than simply maintaining the status quo. This change however will require IT personnel to embrace the new world order and learn a new set of skills that the organization will require. IT departments, and individual IT personnel, will need to move from being primarily technologists with a modicum of business knowledge, to being truly balanced professionals who can equally mix technical ability with an understanding of the business drivers.

⁴ <http://www.reuters.com/article/idUS144602338520100818>

⁵ Gartner Report <http://www.gartner.com/it/page.jsp?id=497088>

It's hard to not resort to hyperbole when discussing just how much Cloud Computing allows organizations to focus on their core business. The very fact that they are able to abstract responsibility for what are essentially commodity services to a third party drives significantly more value to the organization than any mere financial benefit through cost reductions that Cloud Computing can bring. Few people would argue that IT departments should be focusing on high-level strategic work. We contend that Cloud Computing enables this to occur more readily than ever before.

THE PROBLEM WITH THE 80-20 RULE

THE 80-20 RULE IS OFTEN USED WITHIN ORGANIZATIONS...

to illustrate the large effects that small variables can have. It was first suggested by business management thinker Joseph Juran and originally called the Pareto principle⁶ after Italian economist Vilfredo Pareto. Rather than an absolute measure, it tends to be a generalization that is intended to make a point about distribution curves. The most well known use of the rule is the 80-20 rule in sales which says that 80% of revenue for a business is derived from 20% of customers.

Information Technology has its own series of 80-20 rules. Gartner estimates that IT maintenance accounts for around 80% of total IT expenditure.⁷

⁶ http://en.wikipedia.org/wiki/Pareto_principle

⁷ Gartner Report <http://www.gartner.com/it/page.jsp?id=497088>

However we contend that the 80-20 rule occurs elsewhere within IT and relates to time, just as much as it relates to monetary costs. When we look at organizations running their own data center infrastructure, and extend Gartner's findings, we hypothesize that only 20% of the time and effort that goes into running applications, where all business value is concentrated, is actually concerned with running those applications themselves. The diagram illustrates the extent that routine and non-core tasks, like patching operating systems and performing backups, impact upon the time of IT departments.

Cloud Computing is a force that helps flip this ratio and gives IT departments the ability to spend 80% of their time on core business processes, like business application design. It's for this reason, the ability to go from 20% of time and money dedicated to core business processes to 80%, that the economics of Cloud Computing is so compelling. Nowhere is the current model's inefficiency more evident than in the opportunity costs that organizations pay to manage their own computing needs.

IDEAL DISTRIBUTION OF TIME & RESOURCES

80%

Application

20%

Core Technology Servers
Operating System Data Center

REMEMBER THE OPPORTUNITY COST

OPPORTUNITY COST, A CONCEPT FIRST DEVELOPED BY BRITISH PHILOSOPHER JOHN STUART MILL...

is a basic economic premise that is concerned with the costs related to the choices NOT made by someone.

Opportunity cost is:

“...the cost related to the next-best choice available to someone who has picked among several mutually exclusive choices. It is a key concept in economics... opportunity costs are not restricted to monetary or financial costs: the real cost of output forgone, lost time, pleasure or any other benefit that provides utility should also be considered opportunity costs...⁸”

Opportunity cost is an important concept when discussing the economics of Cloud Computing because it allows one to assess the true cost of any potential action. When choosing a particular direction for IT spend, for example, there may be no “direct cost” attached to maintaining the status quo. Data centers have already been built, software purchased. However by including opportunity costs in any calculation, an organization allows for a truer comparison between the various choices to be made.

⁸ http://en.wikipedia.org/wiki/Opportunity_cost

We can now apply the concept to a decision to either retain on-premise IT or move to the Cloud. As we've already seen, roughly 80% of IT time and expenditure is wasted on processes that don't create any value for the organization. The opportunity cost of not choosing the Cloud is therefore the benefit that can accrue to the organization through optimal utilization of that 80%. To put it simplistically, a move to the Cloud can make the difference between an organization being 20% efficient, and one being 80% efficient.

While opportunity cost, and the value to be gained by reducing that cost, is a compelling benefit of moving to Cloud Computing, many critical readers will want to see more concrete examples of the economics at work. To this end it is important to understand the gains to be made from a move away from capital expenditure, and over to operating expenditure.

OPEX IS THE NEW CAPEX

TRADITIONAL IT EXPENDITURE HAS BEEN VERY CAPITAL INTENSIVE.

Hardware had to be bought outright and software licenses were generally an expenditure that appeared on the balance sheet. For this reason the decision making process for technology spend became very drawn out.

One of the core tenets of Cloud Computing is that it is a recurring expenditure model much like telephone or electricity expenditure in that it is accounted for as a standard operating expense. There are several distinct reasons that operating expenditure (OpEx), is preferred to capital expenditure (CapEx).⁹

FINANCIAL CONSIDERATIONS

OpEx is beneficial for the organization, as it gives the flexibility to terminate costs at will. With a capital purchase, the server or software being acquired is fully committed to. Regardless of whether it is being utilized, the ongoing costs, by way of depreciation or financing costs, still need to be borne. Contrast this with OpEx where, in the event that the item is no longer required, payments can cease rapidly. It is for this reason that many companies prefer leasing vehicles in place of purchasing them outright.

While it is true that organizations pay a premium per unit for the flexibility to be able to suspend service, the total cost of ownership of owned assets is much higher as demonstrated in the chart.¹⁰

ALLOWS BUSINESS UNITS TO DECIDE

Most organizations have relatively strict rules in place for all but the most simple of capital expenditures. Operating expenditure however tends to be more frequently delegated to individual business units. In this way, and in keeping with the democratization that is attendant with Cloud Computing¹¹, individual business units have the ability to acquire technology that answers their particular business needs. This flexibility at the business unit, and even individual level, is in fact one of the major forces that is contributing to the growth of Cloud Computing. While sometimes

	Internal IT	Managed Services	The Cloud
Capital Investment	\$40,000	\$0	\$0
Setup Costs	\$1,000	\$5,000	\$1,000
Monthly Services	\$0	\$4,000	\$2,400
Monthly Labor	\$3,200	\$0	\$1,000
Cost over three years	\$149,000	\$129,000	\$106,000
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 Reese^a

¹⁰ <http://broadcast.oreilly.com/2008/10/the-economics-of-cloud-c.html>

¹¹ http://broadcast.rackspace.com/hosting_knowledge/whitepapers/Revolution_Not_Evolution-Whitepaper.pdf

^a <http://fyi.oreilly.com/2009/04/cloud-computing--an-excerpt-fr.html>

lauded, and sometimes lamented in the press,¹² “rogue” departments in larger organizations are often the first to experiment with Cloud Computing and the first to experience the economic benefits that go along with it.

OVERCOMES EXPENDITURE LIMITATIONS

Acquiring capital for large purchases is difficult, for all sizes of organization. This is especially true for smaller organizations for which finance companies apply rigorous debt to equity ratios and thus the amount of capital they can acquire. For this reason it has historically been difficult for organizations to sufficiently justify capital expenditure to get approval for many projects. Moving to an OpEx model removes this limitation and allows small scale projects to be undertaken, unconstrained by capital considerations. While a move away from CapEx is undoubtedly attractive to organizations, it is via Total Cost of Ownership (TOC) that the economic benefits of Cloud Computing become most clear.

TOTAL COST OF OWNERSHIP

WHEN COMPARING COSTS BETWEEN ON-PREMISE OPTIONS...

and Cloud Computing,¹³ it is important to accurately assess the true costs of both options. It's important to remember that, with the Cloud, most costs are up front and readily calculated – this is due to a number of factors:

- Cloud providers give transparent pricing based on different usage metrics - RAM, storage, bandwidth, among others
- Pricing is frequently fixed per unit of time. Customers gain certainty over pricing and are then able to readily calculate costs based on several different usage estimates

Compare this to on-premise technology. In a recent article for CIO.com,¹⁴ Bernard Golden discussed why direct cost-comparisons between the Cloud and on-premise are difficult. As he points out, calculations of in-house costs fail to take into account:

- The direct costs that accompany running a server: power, floor space, storage, and IT operations to manage those resources
- The indirect costs of running a server: network and storage infrastructure and IT operations to manage the general infrastructure.
- The overhead costs of owning a server: procurement and accounting personnel, not to mention a critical resource in short supply: IT management and its attention.

¹³ See <http://diversity.net.nz/wp-content/uploads/2011/01/Moving-to-the-Clouds.pdf>

¹⁴ http://www.cio.com/article/484429/Capex_vs._Opex_Most_People_Miss_the_Point_About_Cloud_Economics

All of these hidden costs make a direct cost comparison difficult. However to help in calculating true TCO, Jonathan Koomey from Stanford University wrote “A Simple Model for Determining the True Total Cost of Ownership for Data Centers”.¹⁵ As part of the paper, Koomey developed a spreadsheet to aid in the calculation of true TCO for data centers which can be used in costs comparisons.¹⁶

¹⁵ <http://www.slideshare.net/data-centers/a-simple-model-for-determining-true-total-cost-of-ownership>

¹⁶ <http://www.itbusinessedge.com/cm/docs/DOC-1212>

TIME IS MONEY

FOCUS ON WHAT MATTERS

A recurring theme among Cloud proponents is the fact that Cloud Computing enables organizations to focus on their core business. In the same way that we consider it bizarre that, given widespread availability of electricity on tap, an organization might create their own electricity plant to power their factory, so too is it becoming more bizarre to host one's own software or buy one's own hardware.

Recently Netflix, the nearly \$10 billion online video rental and streaming service, posted¹⁷ an article detailing their decisions to move to Cloud Computing infrastructure. It is worth reviewing their rationale, in particular a deciding factor that directly relates to focusing on core activities. As Netflix says:

“The problems (the Cloud hosting companies) are trying to solve are incredibly difficult ones, but they aren't specific to our business. Every successful internet company has to figure out great storage solutions, hardware failover, networking infrastructure, etc. We want our engineers to focus as much of their time as possible on product innovation for the Netflix customer experience; that is what differentiates us from our competitors... We chose to be pioneers in this transition (to the Cloud) so we could leverage our investment as we grow, rather than to double down on a model we expect will decline in the industry.”

While it can be difficult to assign a direct economic benefit to the ability to focus on core business processes and abstract responsibility for non-core activities to a third party, we consider it analogous to specialization of labor. In the 18th Century, Economist Adam Smith described in *The Wealth of Nations*,¹⁸ the production of a pin in the following way;¹⁹

“One man draws out the wire, another straightens it, a third cuts it, a fourth points it, a fifth grinds it at the top for receiving the head: to make the head requires two or three distinct operations: to put it on is a particular business, to whiten the pins is another ... and the important business of making a pin is, in this manner, divided into about eighteen distinct operations.”

Smith goes on to estimate somewhere between a 240 and 4800 fold increase in productivity by a specialization of labor of this sort. It is our view that Cloud Computing will drive benefits for organizational focus akin to those that Smith witnessed in the pin factory.

¹⁸ <http://www.econlib.org/library/Smith/smWN.html>

¹⁹ http://en.wikipedia.org/wiki/Business_process

CASE STUDY

I've been reading a copy of Do More Faster recently. Do More Faster is a kind of an entrepreneurs guide book written by two well-known authors; Brad Feld, founder and investor of the Foundry and co-founder of the TechStars startup accelerator and David Cohen, co-founder and CEO of TechStars.

Do More Faster is a series of contributions, mostly from entrepreneurs who have been through the TechStars program. Some of these entrepreneurs have gone on to sell their businesses; many have obtained venture funding for their startups, while some have shut down their businesses. The book then is a great lesson from the coalface of startup life.

One aspect of business that the contributors seem to keep coming back to again and again is the need to be focused. This take many different forms; from the need to be able to pivot the business idea if and when necessary, to the need to measure every metric. No matter what micro aspect of business the chapter was looking at, it seems it could be summarized by the word "focus".

One of the benefits of Cloud Computing is that it allows organizations to focus on what is important. Rather than spending time simply "keeping the lights on", performing routine activities that add nothing to the business itself, Cloud Computing gives the opportunity to abstract away from the business all that is non-core (at least in the IT area) and instead focus solely on strategic value.

That's the value that Cloud can bring to organizations – and every extra ounce of focus in these difficult times is valuable.



A handwritten signature in black ink, appearing to be 'R' followed by a flourish.

SOFTWARE AS A SERVICE

SOFTWARE AS A SERVICE (SaaS) IS DEFINED AS;

“...software that is deployed over the internet... With SaaS, a provider licenses an application to customers either as a service on demand, through a subscription, in a “pay-as-you-go” model, or at no charge when there is opportunity to generate revenue from streams other than the user, such as from advertisement or user list sales.”²⁰”

SaaS is a rapidly growing market as indicated in recent reports that predict ongoing double digit growth.^{21 22} This rapid growth indicates that SaaS will soon become commonplace within every organization and hence it is important that buyers and users of technology understand what SaaS is and where it is suitable.

²⁰ http://en.wikipedia.org/wiki/Software_as_a_service

²¹ <http://www.readwriteweb.com/cloud/2010/07/sass-providers-challenge-the-k.php>

²² <http://www.networkworld.com/news/2010/101810-saas-on-a-tear-says.html>

CHARACTERISTICS OF SaaS

Some defining characteristics of SaaS include;

- Web access to commercial software
- Software managed from a central location
- Software delivered in a “one to many” model
- Users not required to handle software upgrades and patches
- Application Programming Interfaces (APIs) allow for integration between different pieces of software

WHERE SaaS MAKES SENSE

Cloud Computing generally and SaaS in particular, is a rapidly growing method of delivering technology. That said, organizations considering a move to the Cloud will want to consider which applications they move to SaaS. As such there are particular solutions we consider prime candidate for an initial move to SaaS;

- “Vanilla” offerings where the solution is largely undifferentiated. A good example of a vanilla offering would include email where many times competitors use the same software precisely because this fundamental technology is a requirement for doing business, but does not itself confer a competitive advantage.
- Applications where there is significant interplay between the organization and the outside world. For example, email newsletter campaign software.
- Applications that have a significant need for web or mobile access. An example would be mobile sales management software.
- Software that is only to be used for a short term need. An example would be collaboration software for a specific project.
- Software where demand spikes significantly. An example tax or billing software used once a month.

SaaS is widely accepted to have been introduced to the business world by the Salesforce.com²³ Customer Relationship Management (CRM) product. As one of the earliest entrants it is not surprising that CRM is the most popular SaaS application area,²⁴ however e-mail, financial management, customer service and expense management have also gotten good uptake via SaaS.

²³ <http://www.salesforce.com>

²⁴ <http://www.networkworld.com/news/2010/050610-gartner-saas-adoption-on-the.html>

CASE STUDY

“Is SaaS simply the old-school ASP (Application Service Provider) model by a new name?” I get this question a lot. ASP’s, which were big in the 1990s, were essentially third parties that bundled up traditional desktop software and delivered it to customers remotely, via the Internet. Ultimately, the ASP model failed for a number of reasons, not least of which was delivering clunky desktop software via a pipe that is of indeterminate and variable size and capacity. ASP’s also failed because they were trying to shoehorn software that had a business model of one-to-one, into a one-to-many delivery method. Desktop software simply isn’t licensed in a way that makes online delivery work.

Without going into a long diatribe about why ASP’s failed, here are a few reasons that SaaS is a more compelling model;

- Cost of deployment – ASP deployed regular desktop software and hence had high cost structures, SaaS is developed from the ground up to be small and efficient
- Usability – ASP customizations were hard, user interfaces were traditional. SaaS brings a modern design aesthetic to software which encourages customization on the client side
- Upgrades – ASPs were reliant on legacy vendors generally glacial development pace. SaaS is all about iteration
- Integrations – ASPs were monolithic whereas SaaS is designed around integration via APIs
- Support costs – ASPs generally added support on to the price of its software, with SaaS it tends to be all-inclusive
- Hardware compliance – ASP often demands particular hardware, true SaaS runs in any modern browser on any operating system

SaaS
≠
ASP

So my advice for anyone who gets told that SaaS is just ASP by another name? Don’t believe it



PLATFORM AS A SERVICE

PLATFORM AS A SERVICE (PaaS) BRINGS THE BENEFITS...

...that SaaS brought for applications to the software development world. PaaS can be defined as a computing platform that allows the creation of web applications quickly and easily and without the complexity of buying and maintaining the software and infrastructure underneath it.

PaaS is analogous to SaaS except that, rather than being software delivered over the web, it is a platform for the creation of software, delivered over the web.

CHARACTERISTICS OF PaaS

There are a number of different takes on what constitutes PaaS but some basic characteristics include;²⁵

- Services to develop, test, deploy, host and maintain applications in the same integrated development environment. All the varying services needed to fulfill the application development process
- Web based user interface creation tools to create, modify, test and deploy different UI scenarios
- Multi-tenant architecture where multiple concurrent users utilize the same development application
- Built in scalability of deployed software including load balancing and failover
- Integration with web services and databases via common standards
- Support for development team collaboration. Some PaaS solutions include project planning and communication tools
- Tools to handle billing and subscription management

²⁵ http://en.wikipedia.org/wiki/Platform_as_a_service and <http://java.dzone.com/articles/what-platform-service-paas>

PaaS, which is similar in many ways to IaaS is differentiated from IaaS by the addition of value added services and comes in two distinct flavors;

- 1 A collaborative platform for software development, focused on workflow management regardless of the data source being used for the application. An example of this approach would be Heroku, a PaaS that utilizes the Ruby on Rails development language.
- 2 A platform that allows for the creation of software utilizing proprietary data from an application. This sort of PaaS can be seen as a method to create applications with a common data form or type. An example of this sort of platform would be the Force.com PaaS from Salesforce.com which is used almost exclusively to develop applications that work with the Salesforce.com CRM.

WHERE PaaS MAKES SENSE

PaaS is especially useful in any situation where multiple developers will be working on a development project or where other external parties need to interact with the development process. Likewise, building on a PaaS can be useful when the application being developed needs to be able to access a proprietary data source (such as data in a particular CRM vendor's application) and the vendor offers easy access to that data through the PaaS. Finally PaaS is useful where developers wish to automate testing and deployment services.

The popularity of agile software development, a group of software development methodologies based on iterative and incremental development, will also increase the uptake of PaaS as it eases the difficulties around rapid development and iteration of software.

Some examples of PaaS include;

- Google App Engine²⁶
- Microsoft Azure Services²⁷
- Force.com²⁸

26 <http://code.google.com/appengine/>

27 <http://www.microsoft.com/windowsazure/>

28 <http://www.salesforce.com/platform/>

CASE STUDY

When talking about the Cloud, I make an effort to explain the Cloud Computing stack to non-technical, Mom and Pop businesses and many have commented that the triangle metaphor, with IaaS on the bottom, SaaS on top and PaaS in the middle, is slightly imperfect for describing what is in the Cloud. One of the reasons for this is the high degree of breadth that PaaS has come to enjoy. Because PaaS is a lot of different things to different people, it can be hard to really get a clear picture of what it is, and what it is not. So this is how I've come to think about it.

When talking to people about how they use PaaS, there tends to be two main groups of users: on one hand there are the stereotypical developers, at the other extreme there are the folks that I'll call business developers. It's fair to say that development is a little segmented across these lines.

The PaaS tools available in the marketplace tend to reflect these groups. On the one hand there are tools that focus on automating some of the infrastructure requirements that developers need to face. Think of these PaaS's as offerings to automate application deployment. On the other hand there are the development tools that allow developers to take business data and create processes and work flows based upon that data. The first approach is typified by tools such as Heroku and EngineYard whereas the later approach is typified by Force.com. A developer wanting to create a new mobile application with zero-connection to the Salesforce.com CRM, would not use the Force.com PaaS. It just wouldn't make sense. Likewise, a business developer wanting to create a Salesforce.com plugin would most likely not use a PaaS like Heroku when Force.com is available and makes simple the access to critical Salesforce.com data.



This breakdown of PaaS use cases into two camps might not clear up all the confusion around PaaS, but I hope it helps.

A handwritten signature in black ink, appearing to be the initials 'AM' followed by a stylized flourish.

INFRASTRUCTURE AS A SERVICE

INFRASTRUCTURE AS A SERVICE (IaaS) IS A WAY OF DELIVERING...

Cloud Computing infrastructure, servers, storage, network and operating systems as an on-demand service. Rather than purchasing servers, software, datacenter space or network equipment, businesses instead buy those resources as a fully outsourced service on demand.²⁹

Within IaaS, there are some sub-categories that 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 some of Cloud Computing features, like virtualization, but does so on a private, dedicated hardware.

Additionally, some hosting providers are beginning to offer a combination of traditional dedicated hosting alongside Public and/ or Private Cloud networks. This combination approach is generally called “Hybrid Cloud”.

CHARACTERISTICS OF IaaS

As with the two previous sections, SaaS and PaaS, IaaS is a rapidly developing field. That said there are some core characteristics which describe what IaaS is. IaaS is generally accepted to comply with the following;

- Resources are distributed as a service
- Allows for dynamic scaling
- Has a variable cost, utility pricing model
- Generally includes multiple users on a single piece of hardware

There are a plethora of IaaS providers out there from the largest Cloud players like Amazon Web Services³⁰ and Rackspace³¹ to more boutique regional players.

As mentioned previously, the line between PaaS and IaaS is becoming more blurred as vendors introduce tools as part of IaaS that help with deployment including the ability to deploy multiple types of Clouds.³²

³⁰ <http://aws.amazon.com/>

³¹ <http://www.rackspacecloud.com/index.php>

³² <http://m.zdnet.com/blog/forrester/is-the-iaaspaas-line-beginning-to-blur/583>

WHERE IaaS MAKES SENSE

IaaS makes sense in a number of situations and these are closely related to the benefits that Cloud Computing bring. Situations that are particularly suitable for Cloud infrastructure include;

- Where demand is very volatile – any time there are significant spikes and troughs in terms of demand on the infrastructure
- For new organizations without the capital to invest in hardware
- Where the organization is growing rapidly and scaling hardware would be problematic
- Where there is pressure on the organization to limit capital expenditure and to move to operating expenditure
- For specific line of business, trial or temporary infrastructure needs

WHERE IaaS MAY NOT BE THE BEST OPTION

While IaaS provides massive advantages for situations where scalability and quick provisioning are beneficial, there are situations where its limitations may be problematic. Examples of situations where we would advise caution with regards IaaS include;

- Where regulatory compliance makes the offshoring or outsourcing of data storage and processing difficult
- Where the highest levels of performance are required, and on-premise or dedicated hosted infrastructure has the capacity to meet the organization's needs

CASE STUDY

In my travels running CloudCamps, I'm always surprised how often the 'how do you define Cloud' question comes up. Spending so much time in the rarefied atmosphere of the twitterverse, it's easy to just assume that everyone else "gets it". The truth is somewhat different.

This fact was driven home to me at a function when I bumped into a Cabinet Minister from my home country of New Zealand. Now this is a powerful guy who one would have expected would have his fingers well and truly on the pulse of innovation. When I was introduced to him as a well-respected Cloud Computing commentator, said Minister asked me if that had something to do with Meteorology! Trying hard to not choke on my croissant, I gave him a quick definition of Cloud, and went home pondering the need for simplification in this area.

You see the definition most people use currently is that provided by NIST, the National Institute of Science and technology. The NIST definition runs several hundred words and while undeniably comprehensive, no one could ever accuse it of being brief or easily read. It's for this reason that I've recently started using a definition devised by Dave Nielsen, one of the founders of CloudCamp and a really nice guy who does a huge amount for the Cloud community. Dave has the following pithy definition that does much to remove the mystery surrounding Cloud and articulates it in ways that anyone can understand.

Under Dave's definition, Cloud is OSSM (pronounced 'awesome'), meaning that Cloud Computing is a computing resource that is;

- On-demand: the server is already setup and ready to be deployed
- Self-service: customer chooses what they want, when they want it
- Scalable: customer can choose how much they want and ramp up if necessary
- Measurable: there's metering/reporting so you know you are getting what you pay for

In testing the acronym with a bunch of people - both from a technology and a non-technology background, the response has been the same across the board. People really connect with the definition and, most importantly, see it as something understandable and non-threatening. That's a good thing for all of us, after all the Cloud truly is awesome, and OSSM helps people understand that.

OSSM
=
AWESOME



CONCLUSION

FINAL WORDS...

This first CloudU book is really just the beginning – both because we’ve got a heap of content still to share, but mainly because Cloud is still in its infancy – like the proverbial snowball rolling down a hill, Cloud will gain momentum and size in the months and years ahead.

That’s why I’m personally so passionate about the Cloud, and why I believe it’s important for anyone involved in technology to understand the implications Cloud has for them. We’re in the midst of a changing paradigm – mobile devices, ubiquitous connectivity and the demand for flexibility at all levels are three distinct demands that Cloud is uniquely positioned to answer.

I hope you enjoyed and gained some insight from this book, and encourage you to look out for the next edition sometime soon.

Thanks!

A handwritten signature in black ink, appearing to read 'Ben Kepes', with a stylized, flowing script.

Ben Kepes, CloudU Curator

ABOUT

DIVERSITY ANALYSIS



Diversity Analysis is a broad spectrum consultancy specializing 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 service leader in cloud computing, and a founder of OpenStack™, an open source cloud platform. The San Antonio-based company provides Fanatical Support® to its customers, across a portfolio of IT services, including Managed Hosting and Cloud Computing. 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>.

